IEEE Std 1003.1™, 2003 Edition

The Open Group Technical Standard Base Specifications, Issue 6

Includes IEEE Std 1003.1™-2001 and IEEE Std 1003.1™-2001/Cor 1-2002

Information Technology — Portable Operating System Interface (POSIX®)

Shell and Utilities

Sponsor

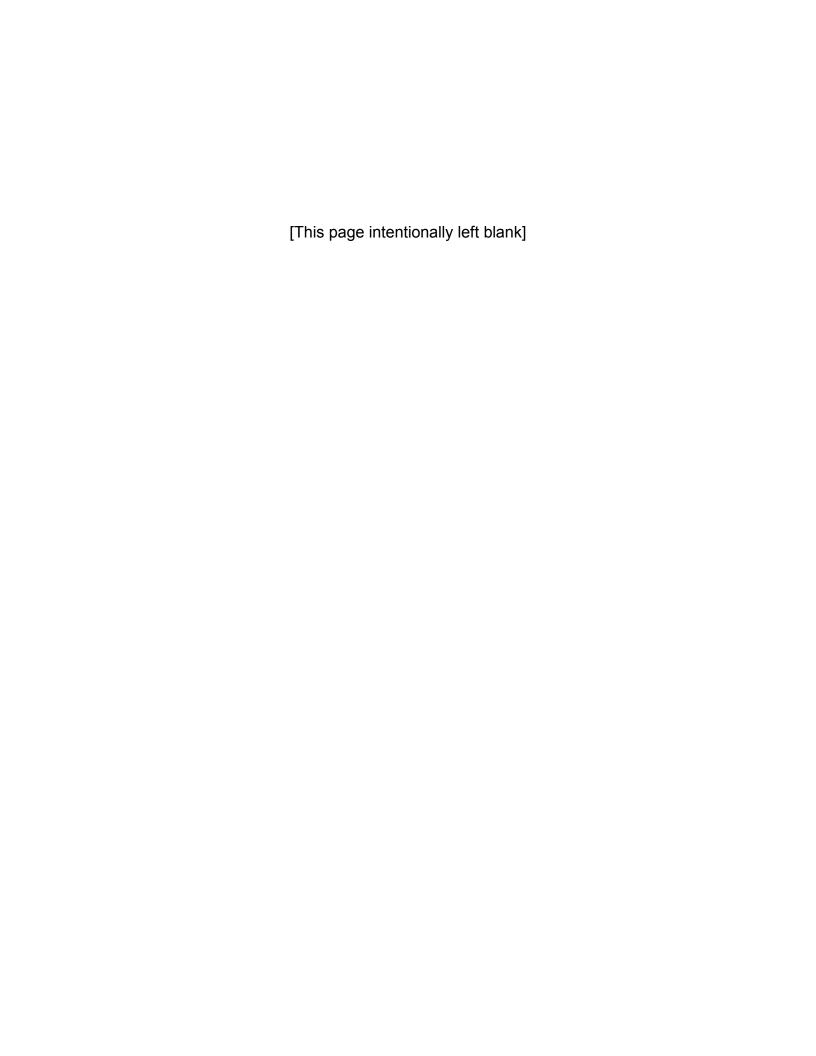
Portable Applications Standards Committee IEEE Computer Society

and

The Open Group







Abstract

This standard is simultaneously ISO/IEC 9945:2002, IEEE Std 1003.1-2001, and forms the core of the Single UNIX Specification, Version 3

The IEEE Std 1003.1, 2003 Edition includes IEEE Std 1003.1-2001/Cor 1-2002 incorporated into IEEE Std 1003.1-2001 (base document). The Corrigendum addresses problems discovered since the approval of IEEE Std 1003.1-2001. These changes are mainly due to resolving integration issues raised by the merger of the base documents that were incorporated into IEEE Std 1003.1-2001, which is the single common revision to IEEE Std 1003.1 $^{\text{TM}}$ -1996, IEEE Std 1003.2 $^{\text{TM}}$ -1992, ISO/IEC 9945-1:1996, ISO/IEC 9945-2:1993, and the Base Specifications of The Open Group Single UNIX Specification, Version 2.

This standard defines a standard operating system interface and environment, including a command interpreter (or "shell"), and common utility programs to support applications portability at the source code level. This standard is intended to be used by both applications developers and system implementors and comprises four major components (each in an associated volume):

- General terms, concepts, and interfaces common to all volumes of this standard, including utility conventions and C-language header definitions, are included in the Base Definitions volume.
- Definitions for system service functions and subroutines, language-specific system services for the C programming language, function issues, including portability, error handling, and error recovery, are included in the System Interfaces volume.
- Definitions for a standard source code-level interface to command interpretation services (a "shell") and common utility programs for application programs are included in the Shell and Utilities volume.
- Extended rationale that did not fit well into the rest of the document structure, which contains historical information concerning the contents of this standard and why features were included or discarded by the standard developers, is included in the Rationale (Informative) volume.

The following areas are outside the scope of this standard:

- Graphics interfaces
- · Database management system interfaces
- · Record I/O considerations
- · Object or binary code portability
- · System configuration and resource availability

This standard describes the external characteristics and facilities that are of importance to applications developers, rather than the internal construction techniques employed to achieve these capabilities. Special emphasis is placed on those functions and facilities that are needed in a wide variety of commercial applications.

Keywords

application program interface (API), argument, asynchronous, basic regular expression (BRE), batch job, batch system, built-in utility, byte, child, command language interpreter, CPU, extended regular expression (ERE), FIFO, file access control mechanism, input/output (I/O), job control, network, portable operating system interface (POSIX $^{\textcircled{\$}}$), parent, shell, stream, string, synchronous, system, thread, X/Open System Interface (XSI)

Copyright © 2001-2003 by the Institute of Electrical and Electronics Engineers, Inc. and The Open Group. All rights reserved. This printing is by the International Organization for Standardization with special permission of the Institute of Electrical and Electronics Engineers, Inc. and The Open Group. Published in Switzerland.

Shell and Utilities, Issue 6

Published 31 March 2003 by the Institute of Electrical and Electronics Engineers, Inc. 3 Park Avenue, New York, NY 10016-5997, U.S.A. ISBN: 0-7381-3437-6 PDF 0-7381-3564-X/SS95078 CD-ROM 0-7381-3563-1/SE95078 Printed in the United States of America by the IEEE.

Published 31 March 2003 by The Open Group Apex Plaza, Forbury Road, Reading, Berkshire RG1 1AX, U.K. Document Number: C033 ISBN: 1-931624-25-9

Printed in the U.K. by The Open Group.

All rights reserved. No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without prior written permission from both the IEEE and The Open Group.

Portions of this standard are derived with permission from copyrighted material owned by Hewlett-Packard Company, International Business Machines Corporation, Novell Inc., The Open Software Foundation, and Sun Microsystems, Inc.

Permissions

Authorization to photocopy portions of this standard for internal or personal use is granted provided that the appropriate fee is paid to the Copyright Clearance Center or the equivalent body outside of the U.S. Permission to make multiple copies for educational purposes in the U.S. requires agreement and a license fee to be paid to the Copyright Clearance Center.

Beyond these provisions, permission to reproduce all or any part of this standard must be with the consent of both copyright holders and may be subject to a license fee. Both copyright holders will need to be satisfied that the other has granted permission. Requests to the copyright holders should be sent by email to austin-group-permissions@opengroup.org.

Feedback

This standard has been prepared by the Austin Group. Feedback relating to the material contained in this standard may be submitted using the Austin Group web site at http://www.opengroup.org/austin/defectform.html.

IEEE Standards documents are developed within the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (IEEE-SA) Standards Board. The IEEE develops its standards through a consensus development process, approved by the American National Standards Institute, which brings together volunteers representing varied viewpoints and interests to achieve the final product. Volunteers are not necessarily members of the Institute and serve without compensation. While the IEEE administers the process and establishes rules to promote fairness in the consensus development process, the IEEE does not independently evaluate, test, or verify the accuracy of any of the information contained in its standards.

Use of an IEEE Standard is wholly voluntary. The IEEE disclaims liability for any personal injury, property, or other damage, of any nature whatsoever, whether special, indirect, consequential, or compensatory, directly or indirectly resulting from the publication, use of, or reliance upon this, or any other IEEE Standard document.

The IEEE does not warrant or represent the accuracy or content of the material contained herein, and expressly disclaims any express or implied warranty, including any implied warranty of merchantability or fitness for a specific purpose, or that the use of the material contained herein is free from patent infringement. IEEE Standards documents are supplied "AS IS".

The existence of an IEEE Standard does not imply that there are no other ways to produce, test, measure, purchase, market, or provide other goods and services related to the scope of the IEEE Standard. Furthermore, the viewpoint expressed at the time a standard is approved and issued is subject to change brought about through developments in the state of the art and comments received from users of the standard. Every IEEE Standard is subjected to review at least every five years for revision or reaffirmation. When a document is more than five years old and has not been reaffirmed, it is reasonable to conclude that its contents, although still of some value, do not wholly reflect the present state of the art. Users are cautioned to check to determine that they have the latest edition of any IEEE Standard.

In publishing and making this document available, the IEEE is not suggesting or rendering professional or other services for, or on behalf of, any person or entity. Nor is the IEEE undertaking to perform any duty owed by any other person or entity to another. Any person utilizing this, and any other IEEE Standards document, should rely upon the advice of a competent professional in determining the exercise of reasonable care in any given circumstances.

Interpretations: Occasionally questions may arise regarding the meaning of portions of standards as they relate to specific applications. When the need for interpretations is brought to the attention of the IEEE, the Institute will initiate action to prepare appropriate responses. Since IEEE Standards represent a consensus of concerned interests, it is important to ensure that any interpretation has also received the concurrence of a balance of interests. For this reason, IEEE and the members of its societies and Standards Coordinating Committees are not able to provide an instant response to interpretation requests except in those cases where the matter has previously received formal consideration.

Comments for revision of IEEE Standards are welcome from any interested party, regardless of membership affiliation with the IEEE.¹ Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments. Comments on standards and requests for interpretations should be addressed to:

Secretary, IEEE-SA Standards Board, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, U.S.A.

Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken with respect to the existence or validity of any patent rights in connection therewith. The IEEE shall not be responsible for identifying patents for which a license may be required by an IEEE Standard or for conducting inquiries into the legal validity or scope of those patents that are brought to its attention.

A patent holder has filed a statement of assurance that it will grant licenses under these rights without compensation or under reasonable rates and non-discriminatory, reasonable terms and conditions to all applicants desiring to obtain such licenses. The IEEE makes no representation as to the reasonableness of rates and/or terms and conditions of the license agreements offered by patent holders. Further information may be obtained from the IEEE Standards Department.

Authorization to photocopy portions of any individual standard for internal or personal use is granted in the U.S. by the Institute of Electrical and Electronics Engineers, Inc., provided that the appropriate fee is paid to the Copyright Clearance Center. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained through the Copyright Clearance Center. To arrange for payment of the licensing fee, please contact:

Copyright Clearance Center, Customer Service, 222 Rosewood Drive, Danvers, MA 01923, U.S.A., Tel.: +1 978 750 8400

Amendments, corrigenda, and interpretations for this standard, or information about the IEEE standards development process, may be found at http://standards.ieee.org.

Full catalog and ordering information on all IEEE publications is available from the IEEE Online Catalog & Store at http://shop.ieee.org/store.

^{1.} For this standard, please send comments via the Austin Group as requested on page iii.

^{2.} Please refer to the special provisions for this standard on page iii concerning permissions from both copyright holders and arrangements to cover photocopying and reproduction across the world, as well as by commercial organizations wishing to license the material for use in product documentation.

The Open Group

The Open Group, a vendor and technology-neutral consortium, is committed to delivering greater business efficiency by bringing together buyers and suppliers of information technology to lower the time, cost, and risks associated with integrating new technology across the enterprise.

The Open Group's mission is to offer all organizations concerned with open information infrastructures a forum to share knowledge, integrate open initiatives, and certify approved products and processes in a manner in which they continue to trust our impartiality.

In the global eCommerce world of today, no single economic entity can achieve independence while still ensuring interoperability. The assurance that products will interoperate with each other across differing systems and platforms is essential to the success of eCommerce and business workflow. The Open Group, with its proven testing and certification program, is the international guarantor of interoperability in the new century.

The Open Group provides opportunities to exchange information and shape the future of IT. The Open Group's members include some of the largest and most influential organizations in the world. The flexible structure of The Open Groups membership allows for almost any organization, no matter what their size, to join and have a voice in shaping the future of the IT world.

More information is available on The Open Group web site at http://www.opengroup.org.

The Open Group has over 15 years' experience in developing and operating certification programs and has extensive experience developing and facilitating industry adoption of test suites used to validate conformance to an open standard or specification. The Open Group portfolio of test suites includes the *Westwood* family of tests for this standard and the associated certification program for Version 3 of the Single UNIX Specification, as well tests for CDE, CORBA, Motif, Linux, LDAP, POSIX.1, POSIX.2, POSIX Realtime, Sockets, UNIX, XPG4, XNFS, XTI, and X11. The Open Group test tools are essential for proper development and maintenance of standards-based products, ensuring conformance of products to industry-standard APIs, applications portability, and interoperability. In-depth testing identifies defects at the earliest possible point in the development cycle, saving costs in development and quality assurance.

More information is available at http://www.opengroup.org/testing.

The Open Group publishes a wide range of technical documentation, the main part of which is focused on development of Technical and Product Standards and Guides, but which also includes white papers, technical studies, branding and testing documentation, and business titles. Full details and a catalog are available at http://www.opengroup.org/pubs.

As with all *live* documents, Technical Standards and Specifications require revision to align with new developments and associated international standards. To distinguish between revised specifications which are fully backwards compatible and those which are not:

- A new *Version* indicates there is no change to the definitive information contained in the previous publication of that title, but additions/extensions are included. As such, it *replaces* the previous publication.
- A new *Issue* indicates there is substantive change to the definitive information contained in the previous publication of that title, and there may also be additions/extensions. As such, both previous and new documents are maintained as current publications.

Readers should note that Corrigenda may apply to any publication. Corrigenda information is published at http://www.opengroup.org/corrigenda.

Full catalog and ordering information on all Open Group publications is available at http://www.opengroup.org/pubs.

Chapter	1	Introduction
-	1.1	Scope
	1.2	Conformance
	1.3	Normative References
	1.4	Change History
	1.5	Terminology
	1.6	Definitions
	1.7	Relationship to Other Documents
	1.7.1	System Interfaces
	1.7.1.1	Process Attributes
	1.7.1.2	Concurrent Execution of Processes
	1.7.1.3	File Access Permissions
	1.7.1.4	File Read, Write, and Creation
	1.7.1.5	File Removal
	1.7.1.6	File Time Values
	1.7.1.7	File Contents
	1.7.1.8	Pathname Resolution
	1.7.1.9	Changing the Current Working Directory
	1.7.1.10	Establish the Locale
	1.7.1.11	Actions Equivalent to Functions
	1.7.2	Concepts Derived from the ISO C Standard
	1.7.2.1	Arithmetic Precision and Operations
	1.7.2.2	Mathematical Functions
	1.8	Portability
	1.8.1	Codes
	1.9	Utility Limits
	1.10	Grammar Conventions
	1.11	Utility Description Defaults
	1.12	Considerations for Utilities in Support of Files of Arbitrary Size 2
	1.13	Built-In Utilities
Chapter	2	Shell Command Language2
-	2.1	Shell Introduction
	2.2	Quoting
	2.2.1	Escape Character (Backslash)
	2.2.2	Single-Quotes3
	2.2.3	Double-Quotes
	2.3	Token Recognition
	2.3.1	Alias Substitution
	2.4	Reserved Words
	2.5	Parameters and Variables
	251	Positional Parameters 3

2.5.2	Special Parameters
2.5.3	Shell Variables
2.6	Word Expansions
2.6.1	Tilde Expansion
2.6.2	Parameter Expansion
2.6.3	Command Substitution
2.6.4	Arithmetic Expansion41
2.6.5	Field Splitting42
2.6.6	Pathname Expansion
2.6.7	Quote Removal
2.7	Redirection
2.7.1	Redirecting Input44
2.7.2	Redirecting Output44
2.7.3	Appending Redirected Output44
2.7.4	Here-Document44
2.7.5	Duplicating an Input File Descriptor45
2.7.6	Duplicating an Output File Descriptor45
2.7.7	Open File Descriptors for Reading and Writing
2.8	Exit Status and Errors
2.8.1	Consequences of Shell Errors
2.8.2	Exit Status for Commands 46
2.9	Shell Commands 47
2.9.1	Simple Commands
2.9.1.1	Command Search and Execution
2.9.2	Pipelines
2.9.3	Lists 50
2.9.3.1	Asynchronous Lists 50
2.9.3.2	Sequential Lists
2.9.3.3	AND Lists 51
2.9.3.4	OR Lists
2.9.4	Compound Commands
2.9.4.1	Grouping Commands
2.9.4.2	The for Loop
2.9.4.3	Case Conditional Construct
2.9.4.4	The if Conditional Construct
2.9.4.5	The while Loop
2.9.4.6	The until Loop
2.9.5	Function Definition Command 54
2.10	
2.10.1	Shell Grammar Lexical Conventions 55 Shell Grammar Lexical Conventions 55
2.10.1	Shell Grammar Rules
2.10.2	
	Signals and Error Handling 61
2.12	Shell Execution Environment 61 Pattern Metabing Notation 65
2.13	Pattern Matching Notation 62
2.13.1	Patterns Matching a Single Character
2.13.2	Patterns Matching Multiple Characters
2.13.3	Patterns Used for Filename Expansion
2.14	Special Built-In Utilities

		break	65
		colon	67
		continue	69
		dot	71
		eval	73
		exec	75
		exit	77
		export	79
		readonly	82
		return	84
		set	86
		shift	92
		times	94
		trap	96
		unset	99
Chapter	3	Batch Environment Services	101
•	3.1	General Concepts	101
	3.1.1		101
	3.1.2	Batch Queues	101
	3.1.3		102
	3.1.4		102
	3.1.5		102
	3.1.6		103
	3.1.7		103
	3.1.8		103
	3.1.9	Batch Authorization	103
	3.1.10		104
	3.1.11	Batch Notification	104
	3.2	Batch Services	104
	3.2.1	Batch Job States	105
	3.2.2	Deferred Batch Services	106
	3.2.2.1	Batch Job Execution	106
	3.2.2.2	Batch Job Routing	113
	3.2.2.3		113
	3.2.2.4	Batch Server Restart	114
	3.2.2.5	Batch Job Abort	114
	3.2.3	Requested Batch Services	115
	3.2.3.1	Delete Batch Job Request	115
	3.2.3.2	Hold Batch Job Request	116
	3.2.3.3	Batch Job Message Request	116
	3.2.3.4		117
	3.2.3.5	Locate Batch Job Request	117
	3.2.3.6	Modify Batch Job Request	117
	3.2.3.7	Move Batch Job Request	118
	3.2.3.8	Queue Batch Job Request	118
	3.2.3.9		119
	3.2.3.10		119

	3.2.3.11	Rerun Batch Job Request	120
	3.2.3.12	Select Batch Jobs Request	
	3.2.3.13	Server Shutdown Request	
	3.2.3.14	Server Status Request	
	3.2.3.15	Signal Batch Job Request	
	3.2.3.16	Track Batch Job Request	121
	3.3	Common Behavior for Batch Environment Utilities	122
	3.3.1	Batch Job Identifier	122
	3.3.2	Destination	123
	3.3.3	Multiple Keyword-Value Pairs	123
Chapter	4	Utilities	125
		Index	1081
List of Fig	gures		
	4-1	pax Format Archive Example	713
List of Ta	bles		
	1-1	Actions when Creating a File that Already Exists	5
	1-2	Selected ISO C Standard Operators and Control Flow Keywords	8
	1-3	Utility Limit Minimum Values	17
	1-4	Symbolic Utility Limits	18
	1-5	Regular Built-In Utilities	28
	3-1	Batch Utilities	101
	3-2	Environment Variable Summary	105
	3-3	Next State Table	107
	3-4	Results/Output Table	108
	3-5	Batch Services Summary	115
	4-1	Expressions in Decreasing Precedence in awk	156
	4-2	Escape Sequences in awk	162
	4-3	Operators in <i>bc</i>	198
	4-4	Programming Environments: Type Sizes	215
	4-5	Programming Environments: c99 and cc Arguments	
	4-6	ASCII to EBCDIC Conversion	
	4-7	ASCII to IBM EBCDIC Conversion	306
	4-8	File Utility Output Strings	446
	4-9	Table Size Declarations in <i>lex</i>	538
	4-10	Escape Sequences in <i>lex</i>	540
	4-11	ERE Precedence in <i>lex</i>	541
	4-12	Named Characters in od	682
	4-13	ustar Header Block	718
	4-14	ustar mode Field	719
	4-15	Octet-Oriented cpio Archive Entry	
	4-16	Values for cpio c_mode Field	
	4-17	Variable Names and Default Headers in ps	755

4-18	Environment Variable Values (Utilities)	804
4-19	Control Character Names in stty	889
4-20	Circumflex Control Characters in stty	
4-21	uuencode Base64 Values	973
4-22	Internal Limits in vacc	1075



Structure of the Standard

This standard was originally developed by the Austin Group, a joint working group of members of the IEEE, members of The Open Group, and members of ISO/IEC Joint Technical Committee 1, as one of the four volumes of IEEE Std 1003.1-2001. The standard was approved by ISO and IEC and published in four parts, correlating to the original volumes.

A mapping of the parts to the volumes is shown below:

ISO/IEC 9945 Part	IEEE Std 1003.1 Volume	Description
9945-1	Base Definitions	Includes general terms, concepts, and interfaces common to all parts of ISO/IEC 9945, including utility conventions and C-language header definitions.
9945-2	System Interfaces	Includes definitions for system service functions and subroutines, language-specific system services for the C programming language, function issues, including portability, error handling, and error recovery.
9945-3	Shell and Utilities	Includes definitions for a standard source code-level interface to command interpretation services (a "shell") and common utility programs for application programs.
9945-4	Rationale	Includes extended rationale that did not fit well into the rest of the document structure, containing historical information concerning the contents of ISO/IEC 9945 and why features were included or discarded by the standard developers.

All four parts comprise the entire standard, and are intended to be used together to accommodate significant internal referencing among them. POSIX-conforming systems are required to support all four parts.

Introduction

Note: This introduction is not part of IEEE Std 1003.1-2001, Standard for Information Technology — Portable Operating System Interface (POSIX).

This standard has been jointly developed by the IEEE and The Open Group. It is simultaneously an IEEE Standard, an ISO/IEC Standard, and an Open Group Technical Standard.

The Austin Group

This standard was developed, and is maintained, by a joint working group of members of the IEEE Portable Applications Standards Committee, members of The Open Group, and members of ISO/IEC Joint Technical Committee 1. This joint working group is known as the Austin Group.³ The Austin Group arose out of discussions amongst the parties which started in early 1998, leading to an initial meeting and formation of the group in September 1998. The purpose of the Austin Group has been to revise, combine, and update the following standards: ISO/IEC 9945-1, ISO/IEC 9945-2, IEEE Std 1003.1, IEEE Std 1003.2, and the Base Specifications of The Open Group Single UNIX Specification.

After two initial meetings, an agreement was signed in July 1999 between The Open Group and the Institute of Electrical and Electronics Engineers (IEEE), Inc., to formalize the project with the first draft of the revised specifications being made available at the same time. Under this agreement, The Open Group and IEEE agreed to share joint copyright of the resulting work. The Open Group has provided the chair and secretariat for the Austin Group.

The base document for the revision was The Open Group's Base volumes of its Single UNIX Specification, Version 2. These were selected since they were a superset of the existing POSIX.1 and POSIX.2 specifications and had some organizational aspects that would benefit the audience for the new revision.

The approach to specification development has been one of "write once, adopt everywhere", with the deliverables being a set of specifications that carry the IEEE POSIX designation, The Open Group's Technical Standard designation, and an ISO/IEC designation. This set of specifications forms the core of the Single UNIX Specification, Version 3.

This unique development has combined both the industry-led efforts and the formal standardization activities into a single initiative, and included a wide spectrum of participants. The Austin Group continues as the maintenance body for this document.

Anyone wishing to participate in the Austin Group should contact the chair with their request. There are no fees for participation or membership. You may participate as an observer or as a contributor. You do not have to attend face-to-face meetings to participate; electronic participation is most welcome. For more information on the Austin Group and how to participate, see http://www.opengroup.org/austin.

The Austin Group is named after the location of the inaugural meeting held at the IBM facility in Austin, Texas in September 1998.

Background

The developers of this standard represent a cross section of hardware manufacturers, vendors of operating systems and other software development tools, software designers, consultants, academics, authors, applications programmers, and others.

Conceptually, this standard describes a set of fundamental services needed for the efficient construction of application programs. Access to these services has been provided by defining an interface, using the C programming language, a command interpreter, and common utility programs that establish standard semantics and syntax. Since this interface enables application writers to write portable applications—it was developed with that goal in mind—it has been designated POSIX, an acronym for Portable Operating System Interface.

Although originated to refer to the original IEEE Std 1003.1-1988, the name POSIX more correctly refers to a *family* of related standards: IEEE Std 1003.*n* and the parts of ISO/IEC 9945. In earlier editions of the IEEE standard, the term POSIX was used as a synonym for IEEE Std 1003.1-1988. A preferred term, POSIX.1, emerged. This maintained the advantages of readability of the symbol "POSIX" without being ambiguous with the POSIX family of standards.

Audience

The intended audience for this standard is all persons concerned with an industry-wide standard operating system based on the UNIX system. This includes at least four groups of people:

- 1. Persons buying hardware and software systems
- 2. Persons managing companies that are deciding on future corporate computing directions
- 3. Persons implementing operating systems, and especially
- 4. Persons developing applications where portability is an objective

Purpose

Several principles guided the development of this standard:

Application-Oriented

The basic goal was to promote portability of application programs across UNIX system environments by developing a clear, consistent, and unambiguous standard for the interface specification of a portable operating system based on the UNIX system documentation. This standard codifies the common, existing definition of the UNIX system.

• Interface, Not Implementation

This standard defines an interface, not an implementation. No distinction is made between library functions and system calls; both are referred to as functions. No details of the implementation of any function are given (although historical practice is sometimes indicated in the RATIONALE section). Symbolic names are given for constants (such as signals and error numbers) rather than numbers.

^{4.} The name POSIX was suggested by Richard Stallman. It is expected to be pronounced *pahz-icks*, as in *positive*, not *poh-six*, or other variations. The pronunciation has been published in an attempt to promulgate a standardized way of referring to a standard operating system interface.

Source, Not Object, Portability

This standard has been written so that a program written and translated for execution on one conforming implementation may also be translated for execution on another conforming implementation. This standard does not guarantee that executable (object or binary) code will execute under a different conforming implementation than that for which it was translated, even if the underlying hardware is identical.

• The C Language

The system interfaces and header definitions are written in terms of the standard C language as specified in the ISO C standard.

• No Superuser, No System Administration

There was no intention to specify all aspects of an operating system. System administration facilities and functions are excluded from this standard, and functions usable only by the superuser have not been included. Still, an implementation of the standard interface may also implement features not in this standard. This standard is also not concerned with hardware constraints or system maintenance.

• Minimal Interface, Minimally Defined

In keeping with the historical design principles of the UNIX system, the mandatory core facilities of this standard have been kept as minimal as possible. Additional capabilities have been added as optional extensions.

• Broadly Implementable

The developers of this standard endeavored to make all specified functions implementable across a wide range of existing and potential systems, including:

- 1. All of the current major systems that are ultimately derived from the original UNIX system code (Version 7 or later)
- 2. Compatible systems that are not derived from the original UNIX system code
- 3. Emulations hosted on entirely different operating systems
- 4. Networked systems
- 5. Distributed systems
- 6. Systems running on a broad range of hardware

No direct references to this goal appear in this standard, but some results of it are mentioned in the Rationale (Informative) volume.

• Minimal Changes to Historical Implementations

When the original version of IEEE Std 1003.1 was published, there were no known historical implementations that did not have to change. However, there was a broad consensus on a set of functions, types, definitions, and concepts that formed an interface that was common to most historical implementations.

The adoption of the 1988 and 1990 IEEE system interface standards, the 1992 IEEE shell and utilities standard, the various Open Group (formerly X/Open) specifications, and the subsequent revisions and addenda to all of them have consolidated this consensus, and this revision reflects the significantly increased level of consensus arrived at since the original versions. The earlier standards and their modifications specified a number of areas where consensus had not been reached before, and these are now reflected in this revision. The authors of the original versions tried, as much as possible, to follow the principles below

when creating new specifications:

- 1. By standardizing an interface like one in an historical implementation; for example, directories
- 2. By specifying an interface that is readily implementable in terms of, and backwards-compatible with, historical implementations, such as the extended *tar* format defined in the *pax* utility
- 3. By specifying an interface that, when added to an historical implementation, will not conflict with it; for example, the *sigaction()* function

This revision tries to minimize the number of changes required to implementations which conform to the earlier versions of the approved standards to bring them into conformance with the current standard. Specifically, the scope of this work excluded doing any "new" work, but rather collecting into a single document what had been spread across a number of documents, and presenting it in what had been proven in practice to be a more effective way. Some changes to prior conforming implementations were unavoidable, primarily as a consequence of resolving conflicts found in prior revisions, or which became apparent when bringing the various pieces together.

However, since it references the 1999 version of the ISO C standard, and no longer supports "Common Usage C", there are a number of unavoidable changes. Applications portability is similarly affected.

This standard is specifically not a codification of a particular vendor's product.

It should be noted that implementations will have different kinds of extensions. Some will reflect "historical usage" and will be preserved for execution of pre-existing applications. These functions should be considered "obsolescent" and the standard functions used for new applications. Some extensions will represent functions beyond the scope of this standard. These need to be used with careful management to be able to adapt to future extensions of this standard and/or port to implementations that provide these services in a different manner.

• Minimal Changes to Existing Application Code

A goal of this standard was to minimize additional work for the developers of applications. However, because every known historical implementation will have to change at least slightly to conform, some applications will have to change.

This Standard

This standard defines the Portable Operating System Interface (POSIX) requirements and consists of the following volumes:

- Base Definitions
- Shell and Utilities (this volume)
- System Interfaces
- Rationale (Informative)

This Volume

The Shell and Utilities volume describes the commands and utilities offered to application programs on POSIX-conformant systems. Readers are expected to be familiar with the Base Definitions volume.

This volume is structured as follows:

- Chapter 1 explains the status of this volume and its relationship to other formal standards. It also describes the defaults used by the utility descriptions in Chapter 4.
- Chapter 2 describes the command language used in POSIX-conformant systems.
- Chapter 3 describes a set of services and utilities that are implemented on systems supporting the Batch Environment Services and Utilities option.
- Chapter 4 consists of reference pages for all utilities available on POSIX-conformant systems.

Comprehensive references are available in the index.

Typographical Conventions

The following typographical conventions are used throughout this standard. In the text, this standard is referred to as IEEE Std 1003.1-2001, which is technically identical to The Open Group Base Specifications, Issue 6.

The typographical conventions listed here are for ease of reading only. Editorial inconsistencies in the use of typography are unintentional and have no normative meaning in this standard.

Reference	Example	Notes
C-Language Data Structure	aiocb	
C-Language Data Structure Member	aio_lio_opcode	
C-Language Data Type	long	
C-Language External Variable	errno	
C-Language Function	system()	
C-Language Function Argument	arg1	
C-Language Function Family	exec	
C-Language Header	<sys stat.h=""></sys>	
C-Language Keyword	return	
C-Language Macro with Argument	assert()	
C-Language Macro with No Argument	INET_ADDRSTRLEN	
C-Language Preprocessing Directive	#define	
Commands within a Utility	a, c	
Conversion Specification, Specifier/Modifier Character	%A, g, E	1
Environment Variable	PATH	
Error Number	[EINTR]	
Example Output	Hello, World	
Filename	/tmp	
Literal Character	'c','\r','\'	2
Literal String	"abcde"	2
Optional Items in Utility Syntax	[]	
Parameter	<directory pathname=""></directory>	
Special Character	<newline></newline>	3

Reference	Example	Notes
Symbolic Constant	_POSIX_VDISABLE	
Symbolic Limit, Configuration Value	{LINE_MAX}	4
Syntax	<pre>#include <sys stat.h=""></sys></pre>	
User Input and Example Code	echo Hello, World	5
Utility Name	awk	
Utility Operand	file_name	
Utility Option	$-\mathbf{c}$	
Utility Option with Option-Argument	− w width	

Notes:

- 1. Conversion specifications, specifier characters, and modifier characters are used primarily in date-related functions and utilities and the *fprintf* and *fscanf* formatting functions.
- 2. Unless otherwise noted, the quotes shall not be used as input or output. When used in a list item, the quotes are omitted. For literal characters, $' \setminus '$ (or any of the other sequences such as $' \cdot ' \cdot '$) is the same as the C constant $' \setminus \setminus '$ (or $' \setminus ' \cdot '$).
- 3. The style selected for some of the special characters, such as <newline>, matches the form of the input given to the *localedef* utility. Generally, the characters selected for this special treatment are those that are not visually distinct, such as the control characters <tab> or <newline>.
- 4. Names surrounded by braces represent symbolic limits or configuration values which may be declared in appropriate headers by means of the C #define construct.
- 5. Brackets shown in this font, "[]", are part of the syntax and do *not* indicate optional items. In syntax the '|' symbol is used to separate alternatives, and ellipses ("...") are used to show that additional arguments are optional.

Shading is used to identify extensions and options; see Section 1.8.1 (on page 9).

Footnotes and notes within the body of the normative text are for information only (informative).

Informative sections (such as Rationale, Change History, Application Usage, and so on) are denoted by continuous shading bars in the margins.

Ranges of values are indicated with parentheses or brackets as follows:

- (a,b) means the range of all values from a to b, including neither a nor b.
- [a,b] means the range of all values from a to b, including a and b.
- [a,b) means the range of all values from a to b, including a, but not b.
- (a,b] means the range of all values from a to b, including b, but not a.

Participants

IEEE Std 1003.1-2001 was prepared by the Austin Group, sponsored by the Portable Applications Standards Committee of the IEEE Computer Society, The Open Group, and ISO/SC22 WG15.

The Austin Group

At the time of approval, the membership of the Austin Group was as follows:

Andrew Josey, Chair Donald W. Cragun, Organizational Representative, IEEE PASC Nicholas Stoughton, Organizational Representative, ISO/SC22 WG15 Mark Brown, Organizational Representative, The Open Group Cathy Hughes, Technical Editor

Austin Group Technical Reviewers

Peter Anvin Michael Gonzalez Sandra O'Donnell Bouazza Bachar Joseph M. Gwinn Frank Prindle Theodore P. Baker Jon Hitchcock Curtis Royster Jr. **Yvette Ho Sang** Walter Briscoe Glen Seeds Mark Brown Cathy Hughes Keld Jorn Simonsen **Dave Butenhof** Lowell G. Johnson Raja Srinivasan Nicholas Stoughton Geoff Clare **Andrew Josey** Donald W. Cragun Michael Kavanaugh Donn S. Terry David Korn Fred Tydeman Lee Damico Ulrich Drepper Marc Aurele La France Peter Van Der Veen James Youngman Paul Eggert Jim Meyering Jim Zepeda Joanna Farley Gary Miller Clive D.W. Feather Finnbarr P. Murphy Jason Zions Andrew Gollan Joseph S. Myers

Austin Group Working Group Members

Harold C. Adams
Peter Anvin
Pierre-Jean Arcos
Jay Ashford
Bouazza Bachar
Theodore P. Baker
Robert Barned
Joel Berman
David J. Blackwood

Shirley Bockstahler-Brandt

James Bottomley
Walter Briscoe
Andries Brouwer
Mark Brown
Eric W. Burger
Alan Burns
Andries Brouwer
Dave Butenhof
Keith Chow
Geoff Clare

Donald W. Cragun Lee Damico

Juan Antonio De La Puente

Ming De Zhou Steven J. Dovich Richard P. Draves Ulrich Drepper Paul Eggert Philip H. Enslow Joanna Farley Clive D.W. Feather Pete Forman

Mark Funkenhauser

Lois Goldthwaite Andrew Gollan Michael Gonzalez Karen D. Gordon Joseph M. Gwinn Steven A. Haaser Charles E. Hammons Chris J. Harding Barry Hedquist Vincent E. Henley Karl Heubaum

Jon Hitchcock
Yvette Ho Sang
Niklas Holsti
Thomas Hosmer
Cathy Hughes
Jim D. Isaak
Lowell G. Johnson
Michael B. Jones
Andrew Josey
Michael J. Karels

Michael Kavanaugh David Korn Steven Kramer Thomas M. Kurihara Marc Aurele La France C. Douglass Locke Nick Maclaren Roger J. Martin Craig H. Meyer Jim Meyering Gary Miller

Finnbarr P. Murphy Joseph S. Myers John Napier

Peter E. Obermayer James T. Oblinger Sandra O'Donnell Frank Prindle Francois Riche John D. Riley Andrew K. Roach Helmut Roth Jaideep Roy Curtis Royster Jr. Stephen C. Schwarm

Glen Seeds

Richard Seibel
David L. Shroads Jr.
W. Olin Sibert
Keld Jorn Simonsen
Curtis Smith
Raja Srinivasan
Nicholas Stoughton
Marc J. Teller
Donn S. Terry
Fred Tydeman
Mark-Rene Uchida
Scott A. Valcourt
Peter Van Der Veen
Michael W. Vannier

Eric Vought

Frederick N. Webb Paul A.T. Wolfgang Garrett A. Wollman James Youngman

Oren Yuen Janusz Zalewski Jim Zepeda Jason Zions

The Open Group

When The Open Group approved the Base Specifications, Issue 6 on 12 September 2001, the membership of The Open Group Base Working Group was as follows:

Andrew Josey, Chair Finnbarr P. Murphy, Vice-Chair Mark Brown, Austin Group Liaison Cathy Hughes, Technical Editor

Base Working Group Members

Bouazza Bachar Joanna Farley Frank Prindle
Mark Brown Andrew Gollan Andrew K. Roach
Dave Butenhof Karen D. Gordon Curtis Royster Jr.
Donald W. Cragun Gary Miller Nicholas Stoughton
Lawre Design P. Marreley Kariina Taviii

Larry Dwyer Finnbarr P. Murphy Kenjiro Tsuji

IEEE

When the IEEE Standards Board approved IEEE Std 1003.1-2001 on 6 December 2001, the membership of the committees was as follows:

Portable Applications Standards Committee (PASC)

Lowell G. Johnson, Chair Joseph M. Gwinn, Vice-Chair Jay Ashford, Functional Chair Andrew Josey, Functional Chair Curtis Royster Jr., Functional Chair Nicholas Stoughton, Secretary

Balloting Committee

The following members of the balloting committee voted on IEEE Std 1003.1-2001. Balloters may have voted for approval, disapproval, or abstention:

Harold C. Adams Steven A. Haaser Frank Prindle Charles E. Hammons Pierre-Jean Arcos Francois Riche Jav Ashford Chris J. Harding John D. Rilev Theodore P. Baker Andrew K. Roach **Barry Hedguist** Robert Barned Vincent E. Henley Helmut Roth David J. Blackwood Karl Heubaum Jaideep Roy Curtis Royster Jr. Shirley Bockstahler-Brandt Niklas Holsti James Bottomley Thomas Hosmer Stephen C. Schwarm Mark Brown Jim D. Isaak Richard Seibel David L. Shroads Jr. Eric W. Burger Lowell G. Johnson Alan Burns Michael B. Jones W. Olin Sibert **Dave Butenhof** Andrew Josev Keld Jorn Simonsen **Keith Chow** Michael J. Karels Nicholas Stoughton Donald W. Cragun Steven Kramer Donn S. Terry Juan Antonio De La Puente Thomas M. Kurihara Mark-Rene Uchida Ming De Zhou C. Douglass Locke Scott A. Valcourt Steven J. Dovich Roger J. Martin Michael W. Vannier Richard P. Draves Craig H. Meyer Frederick N. Webb Finnbarr P. Murphy Paul A.T. Wolfgang Philip H. Enslow Michael Gonzalez John Napier Oren Yuen Peter E. Obermayer Janusz Zalewski Karen D. Gordon Joseph M. Gwinn James T. Oblinger

The following organizational representative voted on this standard:

Andrew Josey, X/Open Company Ltd.

IEEE-SA Standards Board

When the IEEE-SA Standards Board approved IEEE Std 1003.1-2001 on 6 December 2001, it had the following membership:

Donald N. Heirman, Chair **James T. Carlo**, Vice-Chair **Judith Gorman**, Secretary

James H. Gurney James W. Moore Satish K. Aggarwal Mark D. Bowman Richard J. Holleman Robert F. Munzner Gary R. Engmann Lowell G. Johnson Ronald C. Petersen Gerald H. Peterson Harold E. Epstein Robert J. Kennelly H. Landis Floyd Joseph L. Koepfinger* John B. Posey Peter H. Lips Gary S. Robinson Jay Forster* Akio Tojo Howard M. Frazier L. Bruce McClung

Ruben D. Garzon Daleep C. Mohla Donald W. Zipse

Also included are the following non-voting IEEE-SA Standards Board liaisons:

Alan Cookson, NIST Representative **Donald R. Volzka**, TAB Representative

Yvette Ho Sang, Don Messina, Savoula Amanatidis, IEEE Project Editors

^{*} Member Emeritus

IEEE Std 1003.1-2001/Cor 1-2002 was prepared by the Austin Group, sponsored by the Portable Applications Standards Committee of the IEEE Computer Society, The Open Group, and ISO/IEC JTC 1/SC22/WG15.

The Austin Group

At the time of approval, the membership of the Austin Group was as follows:

Andrew Josey, Chair
Donald W. Cragun, Organizational Representative, IEEE PASC
Nicholas Stoughton, Organizational Representative, ISO/IEC JTC 1/SC22/WG15
Mark Brown, Organizational Representative, The Open Group
Cathy Fox, Technical Editor

Austin Group Technical Reviewers

Theodore P. Baker Mark Funkenhauser Frank Prindle Julian Blake Lois Goldthwaite Kenneth Raeburn Andries Brouwer Andrew Gollan Tim Robbins Michael Gonzalez Glen Seeds Mark Brown Dave Butenhof Bruno Haible Matthew Seitz **Geoff Clare** Ben Harris Keld Jorn Simonsen Donald W. Cragun Jon Hitchcock Nicholas Stoughton Ken Dawson Alexander Terekhov Andreas Jaeger Ulrich Drepper Andrew Josey Donn S. Terry Larry Dwyer Jonathan Lennox Mike Wilson Paul Eggert Nick Maclaren Garrett A. Wollman Joanna Farley Jack McCann Mark Ziegast Clive D.W. Feather Wilhelm Mueller Cathy Fox Joseph S. Myers

Austin Group Working Group Members

Harold C. Adams
Alejandro Alonso
Jay Ashford
Theodore P. Baker
David J. Blackwood
Julian Blake
Mitchell Bonnett
Andries Brouwer
Mark Brown
Eric W. Burger
Alan Burns
Dave Butenhof
Keith Chow
Geoff Clare
Luis Cordova

Donald W. Cragun Dragan Cvetkovic Lee Damico Ken Dawson Jeroen Dekkers

Juan Antonio De La Puente

Steven J. Dovich Ulrich Drepper Dr. Sourav Dutta Larry Dwyer Paul Eggert Joanna Farley Clive D.W. Feather Yaacov Fenster Cathy Fox

Mark Funkenhauser Lois Goldthwaite Andrew Gollan Michael Gonzalez Karen D. Gordon Scott Gudgel Joseph M. Gwinn Steven A. Haaser Bruno Haible

Charles E. Hammons

Bryan Harold Ben Harris

Barry Hedguist

Barry Hedquist
Karl Heubaum
Jon Hitchcock
Andreas Jaeger
Andrew Josey
Kenneth Lang
Pi-Cheng Law
Jonathan Lennox
Nick Maclaren

Roger J. Martin Jack McCann George Miao Wilhelm Mueller
Finnbarr P. Murphy
Joseph S. Myers
Alexey Neyman
Charles Ngethe
Peter Petrov
Frank Prindle
Vikram Punj
Kenneth Raeburn
Francois Riche
Tim Robbins
Curtis Royster Jr.

Gil Shultz

Stephen C. Schwarm

Diane Schleicher

Glen Seeds Matthew Seitz Keld Jorn Simonsen Doug Stevenson Nicholas Stoughton Alexander Terekhov Donn S. Terry

Donn S. Terry Mike Wilson

Garrett A. Wollman

Oren Yuen Mark Ziegast

The Open Group

When The Open Group approved the Base Specifications, Issue 6, Technical Corrigendum 1 on 7 February 2003, the membership of The Open Group Base Working Group was as follows:

Andrew Josey, Chair Finnbarr P. Murphy, Vice-Chair Mark Brown, Austin Group Liaison Cathy Fox, Technical Editor

Base Working Group Members

Mark Brown Joanna Farley
Dave Butenhof Andrew Gollan
Donald W. Cragun Finnbarr P. Murphy
Larry Dwyer Frank Prindle
Ulrich Drepper Andrew K. Roach

Curtis Royster Jr. Nicholas Stoughton Kenjiro Tsuji

IEEE

When the IEEE Standards Board approved IEEE Std 1003.1-2001/Cor 1-2002 on 11 December 2002, the membership of the committees was as follows:

Portable Applications Standards Committee (PASC)

Lowell G. Johnson, Chair Joseph M. Gwinn, Vice-Chair Jay Ashford, Functional Chair Andrew Josey, Functional Chair Curtis Royster Jr., Functional Chair Nicholas Stoughton, Secretary

Balloting Committee

The following members of the balloting committee voted on IEEE Std 1003.1-2001/Cor 1-2002. Balloters may have voted for approval, disapproval, or abstention:

Alejandro Alonso	Michael Gonzalez	Charles Ngethe
Jay Ashford	Scott Gudgel	Peter Petrov
David J. Blackwood	Charles E. Hammons	Frank Prindle
Julian Blake	Bryan Harold	Vikram Punj
Mitchell Bonnett	Barry Hedquist	Francois Riche
Mark Brown	Karl Heubaum	Curtis Royster Jr.
Dave Butenhof	Lowell G. Johnson	Diane Schleicher
Keith Chow	Andrew Josey	Stephen C. Schwarm
Luis Cordova	Kenneth Lang	Gil Shultz
Donald W. Cragun	Pi-Cheng Law	Nicholas Stoughton
Steven J. Dovich	George Miao	Donn S. Terry
Dr. Sourav Dutta	Roger J. Martin	Oren Yuen
Yaacov Fenster	Finnbarr P. Murphy	Juan A. de la Puente

IEEE-SA Standards Board

When the IEEE-SA Standards Board approved IEEE Std 1003.1-2001/Cor 1-2002 on 11 December 2002, the membership was as follows:

James T. Carlo, Chair James H. Gurney, Vice-Chair Judith Gorman, Secretary

Sid Bennett Arnold M. Greenspan Daleep C. Mohla H. Stephen Berger Raymond Hapeman William J. Moylan Clyde R. Camp Donald M. Heirman Malcolm V. Thaden Richard DeBlasio Richard H. Hulett Geoffrey O. Thompson Harold E. Epstein Lowell G. Johnson Howard L. Wolfman Julian Forster* Joseph L. Koepfinger* Don Wright

Howard M. Frazier Peter H. Lips
Toshio Fukuda Nader Mehravari

Also included are the following non-voting IEEE-SA Standards Board liaisons:

Alan Cookson, NIST Representative **Satish K. Aggarwal**, NRC Representative **Savoula Amanatidis**, IEEE Standards Managing Editor

^{*} Member Emeritus

Trademarks

The following information is given for the convenience of users of this standard and does not constitute endorsement of these products by The Open Group or the IEEE. There may be other products mentioned in the text that might be covered by trademark protection and readers are advised to verify them independently.

1003.1TM is a trademark of the Institute of Electrical and Electronic Engineers, Inc.

AIX[®] is a registered trademark of IBM Corporation.

AT&T® is a registered trademark of AT&T in the U.S.A. and other countries.

BSDTM is a trademark of the University of California, Berkeley, U.S.A.

Hewlett-Packard[®], HP[®], and HP-UX[®] are registered trademarks of Hewlett-Packard Company.

IBM[®] is a registered trademark of International Business Machines Corporation.

The Open Group and Boundaryless Information Flow are trademarks and UNIX is a registered trademark of The Open Group in the United States and other countries. All other trademarks are the property of their respective owners.

POSIX[®] is a registered trademark of the Institute of Electrical and Electronic Engineers, Inc.

Sun[®] and Sun Microsystems[®] are registered trademarks of Sun Microsystems, Inc.

 $/usr/group^{\textcircled{R}}$ is a registered trademark of UniForum, the International Network of UNIX System Users.

Acknowledgements

The contributions of the following organizations to the development of IEEE Std 1003.1-2001 are gratefully acknowledged:

- AT&T for permission to reproduce portions of its copyrighted System V Interface Definition (SVID) and material from the UNIX System V Release 2.0 documentation.
- The SC22 WG14 Committees.

This standard was prepared by the Austin Group, a joint working group of the IEEE, The Open Group, and ISO SC22 WG15.

Referenced Documents

Normative References

Normative references for this standard are defined in the Base Definitions volume.

Informative References

The following documents are referenced in this standard:

1984 /usr/group Standard

/usr/group Standards Committee, Santa Clara, CA, UniForum 1984.

Almasi and Gottlieb

George S. Almasi and Allan Gottlieb, *Highly Parallel Computing*, The Benjamin/Cummings Publishing Company, Inc., 1989, ISBN: 0-8053-0177-1.

ANSI C

American National Standard for Information Systems: Standard X3.159-1989, Programming Language C.

ANSI X3.226-1994

American National Standard for Information Systems: Standard X3.226-1994, Programming Language Common LISP.

Brawer

Steven Brawer, Introduction to Parallel Programming, Academic Press, 1989, ISBN: 0-12-128470-0.

DeRemer and Pennello Article

DeRemer, Frank and Pennello, Thomas J., *Efficient Computation of LALR(1) Look-Ahead Sets*, SigPlan Notices, Volume 15, No. 8, August 1979.

Draft ANSI X3J11.1

IEEE Floating Point draft report of ANSI X3J11.1 (NCEG).

FIPS 151-1

Federal Information Procurement Standard (FIPS) 151-1. Portable Operating System Interface (POSIX)—Part 1: System Application Program Interface (API) [C Language].

FIPS 151-2

Federal Information Procurement Standards (FIPS) 151-2, Portable Operating System Interface (POSIX)— Part 1: System Application Program Interface (API) [C Language].

HP-UX Manual

Hewlett-Packard HP-UX Release 9.0 Reference Manual, Third Edition, August 1992.

IEC 60559: 1989

IEC 60559: 1989, Binary Floating-Point Arithmetic for Microprocessor Systems (previously designated IEC 559: 1989).

IEEE Std 754-1985

IEEE Std 754-1985, IEEE Standard for Binary Floating-Point Arithmetic.

IEEE Std 854-1987

IEEE Std 854-1987, IEEE Standard for Radix-Independent Floating-Point Arithmetic.

IEEE Std 1003.9-1992

IEEE Std 1003.9-1992, IEEE Standard for Information Technology — POSIX FORTRAN 77 Language Interfaces — Part 1: Binding for System Application Program Interface API.

IETF RFC 791

Internet Protocol, Version 4 (IPv4), September 1981.

IETF RFC 819

The Domain Naming Convention for Internet User Applications, Z. Su, J. Postel, August 1982.

IETF RFC 822

Standard for the Format of ARPA Internet Text Messages, D.H. Crocker, August 1982.

IETF RFC 919

Broadcasting Internet Datagrams, J. Mogul, October 1984.

IETF RFC 920

Domain Requirements, J. Postel, J. Reynolds, October 1984.

IETF RFC 921

Domain Name System Implementation Schedule, J. Postel, October 1984.

IETF RFC 922

Broadcasting Internet Datagrams in the Presence of Subnets, J. Mogul, October 1984.

IETF RFC 1034

Domain Names — Concepts and Facilities, P. Mockapetris, November 1987.

IETF RFC 1035

Domain Names — Implementation and Specification, P. Mockapetris, November 1987.

IETF RFC 1123

Requirements for Internet Hosts — Application and Support, R. Braden, October 1989.

IETF RFC 1886

DNS Extensions to Support Internet Protocol, Version 6 (IPv6), C. Huitema, S. Thomson, December 1995.

IETF RFC 2045

Multipurpose Internet Mail Extensions (MIME), Part 1: Format of Internet Message Bodies, N. Freed, N. Borenstein, November 1996.

IETF RFC 2181

Clarifications to the DNS Specification, R. Elz, R. Bush, July 1997.

IETF RFC 2373

Internet Protocol, Version 6 (IPv6) Addressing Architecture, S. Deering, R. Hinden, July 1998.

IETF RFC 2460

Internet Protocol, Version 6 (IPv6), S. Deering, R. Hinden, December 1998.

Internationalisation Guide

Guide, July 1993, Internationalisation Guide, Version 2 (ISBN: 1-859120-02-4, G304), published by The Open Group.

ISO C (1990)

ISO/IEC 9899: 1990, Programming Languages — C, including Amendment 1: 1995 (E), C Integrity (Multibyte Support Extensions (MSE) for ISO C).

ISO 2375: 1985

ISO 2375: 1985, Data Processing — Procedure for Registration of Escape Sequences.

ISO 8652: 1987

ISO 8652:1987, Programming Languages — Ada (technically identical to ANSI standard 1815A-1983).

ISO/IEC 1539: 1990

ISO/IEC 1539:1990, Information Technology — Programming Languages — Fortran (technically identical to the ANSI X3.9-1978 standard [FORTRAN 77]).

ISO/IEC 4873: 1991

ISO/IEC 4873: 1991, Information Technology — ISO 8-bit Code for Information Interchange — Structure and Rules for Implementation.

ISO/IEC 6429: 1992

ISO/IEC 6429:1992, Information Technology — Control Functions for Coded Character Sets.

ISO/IEC 6937: 1994

ISO/IEC 6937:1994, Information Technology — Coded Character Set for Text Communication — Latin Alphabet.

ISO/IEC 8802-3:1996

ISO/IEC 8802-3: 1996, Information Technology — Telecommunications and Information Exchange Between Systems — Local and Metropolitan Area Networks — Specific Requirements — Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications.

ISO/IEC 8859

ISO/IEC 8859, Information Technology — 8-Bit Single-Byte Coded Graphic Character Sets:

Part 1: Latin Alphabet No. 1

Part 2: Latin Alphabet No. 2

Part 3: Latin Alphabet No. 3

Part 4: Latin Alphabet No. 4

Part 5: Latin/Cyrillic Alphabet

Part 6: Latin/Arabic Alphabet

Part 7: Latin/Greek Alphabet

Part 8: Latin/Hebrew Alphabet

Part 9: Latin Alphabet No. 5

Part 10: Latin Alphabet No. 6

Part 13: Latin Alphabet No. 7

Part 14: Latin Alphabet No. 8

Part 15: Latin Alphabet No. 9

ISO POSIX-1: 1996

ISO/IEC 9945-1:1996, Information Technology — Portable Operating System Interface (POSIX) — Part 1: System Application Program Interface (API) [C Language] (identical to ANSI/IEEE Std 1003.1-1996). Incorporating ANSI/IEEE Stds 1003.1-1990, 1003.1b-1993, 1003.1c-1995, and 1003.1i-1995.

ISO POSIX-2: 1993

ISO/IEC 9945-2:1993, Information Technology — Portable Operating System Interface (POSIX) — Part 2: Shell and Utilities (identical to ANSI/IEEE Std 1003.2-1992, as amended by ANSI/IEEE Std 1003.2a-1992).

Issue 1

X/Open Portability Guide, July 1985 (ISBN: 0-444-87839-4).

Issue 2

X/Open Portability Guide, January 1987:

- Volume 1: XVS Commands and Utilities (ISBN: 0-444-70174-5)
- Volume 2: XVS System Calls and Libraries (ISBN: 0-444-70175-3)

Issue 3

X/Open Specification, 1988, 1989, February 1992:

- Commands and Utilities, Issue 3 (ISBN: 1-872630-36-7, C211); this specification was formerly X/Open Portability Guide, Issue 3, Volume 1, January 1989, XSI Commands and Utilities (ISBN: 0-13-685835-X, XO/XPG/89/002)
- System Interfaces and Headers, Issue 3 (ISBN: 1-872630-37-5, C212); this specification was formerly X/Open Portability Guide, Issue 3, Volume 2, January 1989, XSI System Interface and Headers (ISBN: 0-13-685843-0, XO/XPG/89/003)
- Curses Interface, Issue 3, contained in Supplementary Definitions, Issue 3 (ISBN: 1-872630-38-3, C213), Chapters 9 to 14 inclusive; this specification was formerly X/Open Portability Guide, Issue 3, Volume 3, January 1989, XSI Supplementary Definitions (ISBN: 0-13-685850-3, XO/XPG/89/004)
- Headers Interface, Issue 3, contained in Supplementary Definitions, Issue 3 (ISBN: 1-872630-38-3, C213), Chapter 19, Cpio and Tar Headers; this specification was formerly X/Open Portability Guide Issue 3, Volume 3, January 1989, XSI Supplementary Definitions (ISBN: 0-13-685850-3, XO/XPG/89/004)

Issue 4

CAE Specification, July 1992, published by The Open Group:

- System Interface Definitions (XBD), Issue 4 (ISBN: 1-872630-46-4, C204)
- Commands and Utilities (XCU), Issue 4 (ISBN: 1-872630-48-0, C203)
- System Interfaces and Headers (XSH), Issue 4 (ISBN: 1-872630-47-2, C202)

Issue 4, Version 2

CAE Specification, August 1994, published by The Open Group:

- System Interface Definitions (XBD), Issue 4, Version 2 (ISBN: 1-85912-036-9, C434)
- Commands and Utilities (XCU), Issue 4, Version 2 (ISBN: 1-85912-034-2, C436)
- System Interfaces and Headers (XSH), Issue 4, Version 2 (ISBN: 1-85912-037-7, C435)

Issue 5

Technical Standard, February 1997, published by The Open Group:

- System Interface Definitions (XBD), Issue 5 (ISBN: 1-85912-186-1, C605)
- Commands and Utilities (XCU), Issue 5 (ISBN: 1-85912-191-8, C604)
- System Interfaces and Headers (XSH), Issue 5 (ISBN: 1-85912-181-0, C606)

Knuth Article

Knuth, Donald E., *On the Translation of Languages from Left to Right*, Information and Control, Volume 8, No. 6, October 1965.

KornShell

Bolsky, Morris I. and Korn, David G., *The New KornShell Command and Programming Language*, March 1995, Prentice Hall.

MSE Working Draft

Working draft of ISO/IEC 9899: 1990/Add3: Draft, Addendum 3 — Multibyte Support Extensions (MSE) as documented in the ISO Working Paper SC22/WG14/N205 dated 31 March 1992.

POSIX.0: 1995

IEEE Std 1003.0-1995, IEEE Guide to the POSIX Open System Environment (OSE) (identical to ISO/IEC TR 14252).

POSIX.1: 1988

IEEE Std 1003.1-1988, IEEE Standard for Information Technology — Portable Operating System Interface (POSIX) — Part 1: System Application Program Interface (API) [C Language].

POSIX.1: 1990

IEEE Std 1003.1-1990, IEEE Standard for Information Technology — Portable Operating System Interface (POSIX) — Part 1: System Application Program Interface (API) [C Language].

POSIX.1a

P1003.1a, Standard for Information Technology — Portable Operating System Interface (POSIX) — Part 1: System Application Program Interface (API) — (C Language) Amendment.

POSIX.1d: 1999

IEEE Std 1003.1d-1999, IEEE Standard for Information Technology — Portable Operating System Interface (POSIX) — Part 1: System Application Program Interface (API) — Amendment 4: Additional Realtime Extensions [C Language].

POSIX.1g: 2000

IEEE Std 1003.1g-2000, IEEE Standard for Information Technology — Portable Operating System Interface (POSIX) — Part 1: System Application Program Interface (API) — Amendment 6: Protocol-Independent Interfaces (PII).

POSIX.1i: 2000

IEEE Std 1003.1j-2000, IEEE Standard for Information Technology — Portable Operating System Interface (POSIX) — Part 1: System Application Program Interface (API) — Amendment 5: Advanced Realtime Extensions [C Language].

POSIX.1q: 2000

IEEE Std 1003.1q-2000, IEEE Standard for Information Technology — Portable Operating System Interface (POSIX) — Part 1: System Application Program Interface (API) — Amendment 7: Tracing [C Language].

POSIX.2b

P1003.2b, Standard for Information Technology — Portable Operating System Interface (POSIX) — Part 2: Shell and Utilities — Amendment.

POSIX.2d:-1994

IEEE Std 1003.2d-1994, IEEE Standard for Information Technology — Portable Operating System Interface (POSIX) — Part 2: Shell and Utilities — Amendment 1: Batch Environment.

POSIX.13:-1998

IEEE Std 1003.13: 1998, IEEE Standard for Information Technology — Standardized Application Environment Profile (AEP) — POSIX Realtime Application Support.

Sarwate Article

Sarwate, Dilip V., *Computation of Cyclic Redundancy Checks via Table Lookup*, Communications of the ACM, Volume 30, No. 8, August 1988.

Sprunt, Sha, and Lehoczky

Sprunt, B., Sha, L., and Lehoczky, J.P., *Aperiodic Task Scheduling for Hard Real-Time Systems*, The Journal of Real-Time Systems, Volume 1, 1989, Pages 27-60.

SVID. Issue 1

American Telephone and Telegraph Company, System V Interface Definition (SVID), Issue 1; Morristown, NJ, UNIX Press, 1985.

SVID, Issue 2

American Telephone and Telegraph Company, System V Interface Definition (SVID), Issue 2; Morristown, NJ, UNIX Press, 1986.

SVID. Issue 3

American Telephone and Telegraph Company, System V Interface Definition (SVID), Issue 3; Morristown, NJ, UNIX Press, 1989.

The AWK Programming Language

Aho, Alfred V., Kernighan, Brian W., and Weinberger, Peter J., *The AWK Programming Language*, Reading, MA, Addison-Wesley 1988.

UNIX Programmer's Manual

American Telephone and Telegraph Company, *UNIX Time-Sharing System: UNIX Programmer's Manual*, 7th Edition, Murray Hill, NJ, Bell Telephone Laboratories, January 1979.

XNS, Issue 4

CAE Specification, August 1994, Networking Services, Issue 4 (ISBN: 1-85912-049-0, C438), published by The Open Group.

XNS, Issue 5

CAE Specification, February 1997, Networking Services, Issue 5 (ISBN: 1-85912-165-9, C523), published by The Open Group.

XNS, Issue 5.2

Technical Standard, January 2000, Networking Services (XNS), Issue 5.2 (ISBN: 1-85912-241-8, C808), published by The Open Group.

X/Open Curses, Issue 4, Version 2

CAE Specification, May 1996, X/Open Curses, Issue 4, Version 2 (ISBN: 1-85912-171-3, C610), published by The Open Group.

Yacc

Yacc: Yet Another Compiler Compiler, Stephen C. Johnson, 1978.

Source Documents

Parts of the following documents were used to create the base documents for this standard:

AIX 3.2 Manual

AIX Version 3.2 For RISC System/6000, Technical Reference: Base Operating System and Extensions, 1990, 1992 (Part No. SC23-2382-00).

OSF/1

OSF/1 Programmer's Reference, Release 1.2 (ISBN: 0-13-020579-6).

OSF AES

Application Environment Specification (AES) Operating System Programming Interfaces Volume, Revision A (ISBN: 0-13-043522-8).

System V Release 2.0

- UNIX System V Release 2.0 Programmer's Reference Manual (April 1984 Issue 2).
- UNIX System V Release 2.0 Programming Guide (April 1984 Issue 2).

System V Release 4.2

Operating System API Reference, UNIX SVR4.2 (1992) (ISBN: 0-13-017658-3).

Chapter 1 Introduction

1.1 Scope

The scope of IEEE Std 1003.1-2001 is described in the Base Definitions volume of IEEE Std 1003.1-2001.

5 1.2 Conformance

Conformance requirements for IEEE Std 1003.1-2001 are defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 2, Conformance.

8 1.3 Normative References

Normative references for IEEE Std 1003.1-2001 are defined in the Base Definitions volume of IEEE Std 1003.1-2001.

1.4 Change History

Change history is described in the Rationale (Informative) volume of IEEE Std 1003.1-2001, and in the CHANGE HISTORY section of reference pages.

1.5 Terminology

This section appears in the Base Definitions volume of IEEE Std 1003.1-2001, but is repeated here for convenience:

For the purposes of IEEE Std 1003.1-2001, the following terminology definitions apply:

can

Describes a permissible optional feature or behavior available to the user or application. The feature or behavior is mandatory for an implementation that conforms to IEEE Std 1003.1-2001. An application can rely on the existence of the feature or behavior.

implementation-defined

Describes a value or behavior that is not defined by IEEE Std 1003.1-2001 but is selected by an implementor. The value or behavior may vary among implementations that conform to IEEE Std 1003.1-2001. An application should not rely on the existence of the value or behavior. An application that relies on such a value or behavior cannot be assured to be portable across conforming implementations.

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

legacy

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

Terminology Introduction

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

may

Describes a feature or behavior that is optional for an implementation that conforms to IEEE Std 1003.1-2001. An application should not rely on the existence of the feature or behavior. An application that relies on such a feature or behavior cannot be assured to be portable across conforming implementations.

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

shall

For an implementation that conforms to IEEE Std 1003.1-2001, describes a feature or behavior that is mandatory. An application can rely on the existence of the feature or behavior.

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

should

For an implementation that conforms to IEEE Std 1003.1-2001, describes a feature or behavior that is recommended but not mandatory. An application should not rely on the existence of the feature or behavior. An application that relies on such a feature or behavior cannot be assured to be portable across conforming implementations.

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

undefined

Describes the nature of a value or behavior not defined by IEEE Std 1003.1-2001 which results from use of an invalid program construct or invalid data input.

The value or behavior may vary among implementations that conform to IEEE Std 1003.1-2001. An application should not rely on the existence or validity of the value or behavior. An application that relies on any particular value or behavior cannot be assured to be portable across conforming implementations.

unspecified

Describes the nature of a value or behavior not specified by IEEE Std 1003.1-2001 which results from use of a valid program construct or valid data input.

The value or behavior may vary among implementations that conform to IEEE Std 1003.1-2001. An application should not rely on the existence or validity of the value or behavior. An application that relies on any particular value or behavior cannot be assured to be portable across conforming implementations.

Introduction Definitions

1.6 Definitions

Concepts and definitions are defined in the Base Definitions volume of IEEE Std 1003.1-2001.

1.7 Relationship to Other Documents

1.7.1 System Interfaces

This subsection describes some of the features provided by the System Interfaces volume of IEEE Std 1003.1-2001 that are assumed to be globally available on all systems conforming to this volume of IEEE Std 1003.1-2001. This subsection does not attempt to detail all of the features defined in the System Interfaces volume of IEEE Std 1003.1-2001 that are required by all of the utilities defined in this volume of IEEE Std 1003.1-2001; the utility and function descriptions point out additional functionality required to provide the corresponding specific features needed by each.

The following subsections describe frequently used concepts. Many of these concepts are described in the Base Definitions volume of IEEE Std 1003.1-2001. Utility and function description statements override these defaults when appropriate.

81 1.7.1.1 Process Attributes

The following process attributes, as described in the System Interfaces volume of IEEE Std 1003.1-2001, are assumed to be supported for all processes in this volume of IEEE Std 1003.1-2001:

Controlling Terminal
Current Working Directory
Effective Group ID
Effective User ID
File Descriptors
File Mode Creation Mask
Process Group ID
Process ID

Real Group ID
Root Directory
Saved Set-Group-ID
Saved Set-User-ID
Session Membership
Supplementary Group IDs
Process ID

riocess in

A conforming implementation may include additional process attributes.

1.7.1.2 Concurrent Execution of Processes

The following functionality of the *fork()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001 shall be available on all systems conforming to this volume of IEEE Std 1003.1-2001:

- 1. Independent processes shall be capable of executing independently without either process terminating.
- 2. A process shall be able to create a new process with all of the attributes referenced in Section 1.7.1.1, determined according to the semantics of a call to the *fork()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001 followed by a call in the child process to one of the *exec* functions defined in the System Interfaces volume of IEEE Std 1003.1-2001.

105 1.7.1.3 File Access Permissions

The file access control mechanism described by the Base Definitions volume of IEEE Std 1003.1-2001, Section 4.4, File Access Permissions shall apply to all files on an implementation conforming to this volume of IEEE Std 1003.1-2001.

109 1.7.1.4 File Read, Write, and Creation

110

111

112

113 114

116

117

118

119

120

121

122 123

124

125

126

127

128129

130

131

132

133 134

135

If a file that does not exist is to be written, it shall be created as described below, unless the utility description states otherwise.

When a file that does not exist is created, the following features defined in the System Interfaces volume of IEEE Std 1003.1-2001 shall apply unless the utility or function description states otherwise:

- 1. The user ID of the file shall be set to the effective user ID of the calling process.
- 2. The group ID of the file shall be set to the effective group ID of the calling process or the group ID of the directory in which the file is being created.
- 3. If the file is a regular file, the permission bits of the file shall be set to:

```
S_IROTH | S_IWOTH | S_IRGRP | S_IWGRP | S_IRUSR | S_IWUSR
```

(see the description of *File Modes* in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 13, Headers, <**sys/stat.h>**) except that the bits specified by the file mode creation mask of the process shall be cleared. If the file is a directory, the permission bits shall be set to:

```
S_IRWXU | S_IRWXG | S_IRWXO
```

except that the bits specified by the file mode creation mask of the process shall be cleared.

- 4. The *st_atime*, *st_ctime*, and *st_mtime* fields of the file shall be updated as specified in the System Interfaces volume of IEEE Std 1003.1-2001, Section 2.5, Standard I/O Streams.
- 5. If the file is a directory, it shall be an empty directory; otherwise, the file shall have length zero.
- 6. If the file is a symbolic link, the effect shall be undefined unless the {POSIX2_SYMLINKS} variable is in effect for the directory in which the symbolic link would be created.
- 7. Unless otherwise specified, the file created shall be a regular file.

When an attempt is made to create a file that already exists, the utility shall take the action indicated in Table 1-1 (on page 5) corresponding to the type of the file the utility is trying to create and the type of the existing file, unless the utility description states otherwise.

Table 1-1 Actions when Creating a File that Already Exists

					N	lew	Ty	pe				Function
Existing Type	В	C	D	F	L	M	P	Q	R	S	T	Creating New
A fattach()-ed STREAM	F	F	F	F	F	_	_	_	OF	_	_	N/A
B Block Special	F	F	F	F	F	_	_	_	OF	_	_	mknod()**
C Character Special	F	F	F	F	F	_	_	_	OF	_	_	mknod()**
D Directory	F	F	F	F	F	_	_	_	F	_	_	mkdir()
F FIFO Special File	F	F	F	F	F	_	_	_	Ο	_	_	mkfifo()
L Symbolic Link	F	F	F	F	F	_	_	_	FL	_	_	symlink()
M Shared Memory	F	F	F	F	F	_	_	_	—	_	_	shm_open()
P Semaphore	F	F	F	F	F	_	_	_		_	_	sem_open()
Q Message Queue	F	F	F	F	F	_	_	_		_	_	mq_open()
R Regular File	F	F	F	F	F	_	_	_	RF	_		open()
S Socket	F	F	F	F	F	_	_	_	_	_	_	bind()
T Typed Memory	F	F	F	F	F	_	_		_	_		*

The following codes are used in Table 1-1:

- Fail. The attempt to create the new file shall fail and the utility shall either continue with its operation or exit immediately with a non-zero exit status, depending on the description of the utility.
- FL Follow link. Unless otherwise specified, the symbolic link shall be followed as specified for pathname resolution, and the operation performed shall be as if the target of the symbolic link (after all resolution) had been named. If the target of the symbolic link does not exist, it shall be as if that nonexistent target had been named directly.
- O Open FIFO. When attempting to create a regular file, and the existing file is a FIFO special file:
 - 1. If the FIFO is not already open for reading, the attempt shall block until the FIFO is opened for reading.
 - 2. Once the FIFO is open for reading, the utility shall open the FIFO for writing and continue with its operation.
- OF The named file shall be opened with the consequences defined for that file type.
- RF Regular file. When attempting to create a regular file, and the existing file is a regular file:
 - 1. The user ID, group ID, and permission bits of the file shall not be changed.
 - 2. The file shall be truncated to zero length.
 - 3. The *st_ctime* and *st_mtime* fields shall be marked for update.
- The effect is implementation-defined unless specified by the utility description.
- * There is no portable way to create a file of this type.
- ** Not portable.
- When a file is to be appended, the file shall be opened in a manner equivalent to using the O_APPEND flag, without the O_TRUNC flag, in the *open()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001.
- When a file is to be read or written, the file shall be opened with an access mode corresponding to the operation to be performed. If file access permissions deny access, the requested operation shall fail.

179 1.7.1.5 File Removal

180

181

182

183

184

185

186

187

188

190 191

192

193

194

195 196

197

198 199

201

202

204

205206

207208

209

210

211212

213

214

215

216217

218

When a directory that is the root directory or current working directory of any process is removed, the effect is implementation-defined. If file access permissions deny access, the requested operation shall fail. Otherwise, when a file is removed:

- 1. Its directory entry shall be removed from the file system.
- 2. The link count of the file shall be decremented.
- 3. If the file is an empty directory (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.143, Empty Directory):
 - a. If no process has the directory open, the space occupied by the directory shall be freed and the directory shall no longer be accessible.
 - b. If one or more processes have the directory open, the directory contents shall be preserved until all references to the file have been closed.
- 4. If the file is a directory that is not empty, the *st_ctime* field shall be marked for update.
- 5. If the file is not a directory:
 - a. If the link count becomes zero:
 - i. If no process has the file open, the space occupied by the file shall be freed and the file shall no longer be accessible.
 - ii. If one or more processes have the file open, the file contents shall be preserved until all references to the file have been closed.
 - b. If the link count is not reduced to zero, the *st_ctime* field shall be marked for update.
- 6. The *st_ctime* and *st_mtime* fields of the containing directory shall be marked for update.

200 1.7.1.6 File Time Values

All files shall have the three time values described by the Base Definitions volume of IEEE Std 1003.1-2001, Section 4.7, File Times Update.

203 1.7.1.7 File Contents

When a reference is made to the contents of a file, *pathname*, this means the equivalent of all of the data placed in the space pointed to by *buf* when performing the *read()* function calls in the following operations defined in the System Interfaces volume of IEEE Std 1003.1-2001:

```
while (read (fildes, buf, nbytes) > 0)
:
```

If the file is indicated by a pathname, the file descriptor shall be determined by the equivalent of the following operation defined in the System Interfaces volume of IEEE Std 1003.1-2001:

```
fildes = open (pathname, O_RDONLY);
```

The value of *nbytes* in the above sequence is unspecified; if the file is of a type where the data returned by *read()* would vary with different values, the value shall be one that results in the most data being returned.

If the *read()* function calls would return an error, it is unspecified whether the contents of the file are considered to include any data from offsets in the file beyond where the error would be returned.

219 1.7.1.8 Pathname Resolution

The pathname resolution algorithm, described by the Base Definitions volume of IEEE Std 1003.1-2001, Section 4.11, Pathname Resolution, shall be used by implementations conforming to this volume of IEEE Std 1003.1-2001; see also the Base Definitions volume of IEEE Std 1003.1-2001, Section 4.5, File Hierarchy.

224 1.7.1.9 Changing the Current Working Directory

When the current working directory (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.436, Working Directory) is to be changed, unless the utility or function description states otherwise, the operation shall succeed unless a call to the *chdir()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001 would fail when invoked with the new working directory pathname as its argument.

230 1.7.1.10 Establish the Locale

The functionality of the *setlocale()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001 shall be available on all systems conforming to this volume of IEEE Std 1003.1-2001; that is, utilities that require the capability of establishing an international operating environment shall be permitted to set the specified category of the international environment.

236 1.7.1.11 Actions Equivalent to Functions

Some utility descriptions specify that a utility performs actions equivalent to a function defined in the System Interfaces volume of IEEE Std 1003.1-2001. Such specifications require only that the external effects be equivalent, not that any effect within the utility and visible only to the utility be equivalent.

1.7.2 Concepts Derived from the ISO C Standard

Some of the standard utilities perform complex data manipulation using their own procedure and arithmetic languages, as defined in their EXTENDED DESCRIPTION or OPERANDS sections. Unless otherwise noted, the arithmetic and semantic concepts (precision, type conversion, control flow, and so on) shall be equivalent to those defined in the ISO C standard, as described in the following sections. Note that there is no requirement that the standard utilities be implemented in any particular programming language.

248 1.7.2.1 Arithmetic Precision and Operations

Integer variables and constants, including the values of operands and option-arguments, used by the standard utilities listed in this volume of IEEE Std 1003.1-2001 shall be implemented as equivalent to the ISO C standard **signed long** data type; floating point shall be implemented as equivalent to the ISO C standard **double** type. Conversions between types shall be as described in the ISO C standard. All variables shall be initialized to zero if they are not otherwise assigned by the input to the application.

Arithmetic operators and control flow keywords shall be implemented as equivalent to those in the cited ISO C standard section, as listed in Table 1-2 (on page 8).

Table 1-2 Selected ISO C Standard Operators and Control Flow Keywords

258		
259	Operation	ISO C Standard Equivalent Reference
260	()	Section 6.5.1, Primary Expressions
261	postfix ++	Section 6.5.2, Postfix Operators
262	postfix	_
263	unary +	Section 6.5.3, Unary Operators
264	unary -	
265	prefix ++	
266	prefix	
267	~	
268	!	
269	sizeof()	
270	*	Section 6.5.5, Multiplicative Operators
271	/	
272	%	
273	+	Section 6.5.6, Additive Operators
274	-	
275	<<	Section 6.5.7, Bitwise Shift Operators
276	>>	
277	<, <=	Section 6.5.8, Relational Operators
278	>, >=	
279	==	Section 6.5.9, Equality Operators
280	!=	
281	&	Section 6.5.10, Bitwise AND Operator
282	^	Section 6.5.11, Bitwise Exclusive OR Operator
283		Section 6.5.12, Bitwise Inclusive OR Operator
284	&&	Section 6.5.13, Logical AND Operator
285		Section 6.5.14, Logical OR Operator
286	expr?expr:expr	Section 6.5.15, Conditional Operator
287	=, *=, /=, %=, +=, -=	Section 6.5.16, Assignment Operators
288	<<=, >>=, &=, ^=, =	
289	if ()	Section 6.8.4, Selection Statements
290	if () else	
291	switch ()	
292	while ()	Section 6.8.5, Iteration Statements
293	do while ()	
294	for ()	
295	goto	Section 6.8.6, Jump Statements
296	continue	
297	break	
298	return	

The evaluation of arithmetic expressions shall be equivalent to that described in Section 6.5, Expressions, of the ISO C standard.

299

300

303 304

305

308

309

310

311

312

313

314

315

316

317

318

320

321

322 323

324

328

329

330

301 1.7.2.2 Mathematical Functions

Any mathematical functions with the same names as those in the following sections of the ISO C standard:

- Section 7.12, Mathematics, <math.h>
- Section 7.20.2, Pseudo-Random Sequence Generation Functions

shall be implemented to return the results equivalent to those returned from a call to the corresponding function described in the ISO C standard.

1.8 Portability

Some of the utilities in the Shell and Utilities volume of IEEE Std 1003.1-2001 and functions in the System Interfaces volume of IEEE Std 1003.1-2001 describe functionality that might not be fully portable to systems meeting the requirements for POSIX conformance (see the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 2, Conformance).

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

Unless the primary task of a utility is to produce textual material on its standard output, application developers should not rely on the format or content of any such material that may be produced. Where the primary task *is* to provide such material, but the output format is incompletely specified, the description is marked with the OF margin code and shading. Application developers are warned not to expect that the output of such an interface on one system is any guide to its behavior on another system.

1.8.1 Codes

Codes and their meanings are listed in the Base Definitions volume of IEEE Std 1003.1-2001, but are repeated here for convenience:

325 ADV Advisory Information

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

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

331 AIO Asynchronous Input and Output

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

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

337 BAR Barriers

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

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

Portability Introduction

343 BE **Batch Environment Services and Utilities** The functionality described is optional. 344 345 Where applicable, utilities are marked with the BE margin legend in the SYNOPSIS section. Where additional semantics apply to a utility, the material is identified by use of the BE margin 346 347 legend. C-Language Development Utilities 348 CD The functionality described is optional. 349 Where applicable, utilities are marked with the CD margin legend in the SYNOPSIS section. 350 Where additional semantics apply to a utility, the material is identified by use of the CD margin 351 legend. 352 CPT Process CPU-Time Clocks 353 The functionality described is optional. The functionality described is also an extension to the ISO C standard. 355 Where applicable, functions are marked with the CPT margin legend in the SYNOPSIS section. 356 Where additional semantics apply to a function, the material is identified by use of the CPT margin legend. 358 Clock Selection 359 CS The functionality described is optional. The functionality described is also an extension to the 360 ISO C standard. 361 362 Where applicable, functions are marked with the CS margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the CS 363 margin legend. 364 Extension to the ISO C standard 365 CXThe functionality described is an extension to the ISO C standard. Application writers may make 366 use of an extension as it is supported on all IEEE Std 1003.1-2001-conforming systems. 367 With each function or header from the ISO C standard, a statement to the effect that "any 368 conflict is unintentional" is included. That is intended to refer to a direct conflict. 369 IEEE Std 1003.1-2001 acts in part as a profile of the ISO C standard, and it may choose to further constrain behaviors allowed to vary by the ISO C standard. Such limitations are not considered 371 conflicts. 372 373 Where additional semantics apply to a function or header, the material is identified by use of the CX margin legend. 374 375 FD **FORTRAN Development Utilities** The functionality described is optional. 376 Where applicable, utilities are marked with the FD margin legend in the SYNOPSIS section. 377 Where additional semantics apply to a utility, the material is identified by use of the FD margin 378 legend. 379 **FORTRAN Runtime Utilities** 380 FR The functionality described is optional. 381 Where applicable, utilities are marked with the FR margin legend in the SYNOPSIS section. Where additional semantics apply to a utility, the material is identified by use of the FR margin 383 legend. 384 File Synchronization FSC 385 The functionality described is optional. The functionality described is also an extension to the 386

387

ISO C standard.

Introduction Portability

388 Where applicable, functions are marked with the FSC margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the FSC 389 margin legend. 390 IPV6 391 IP6 The functionality described is optional. The functionality described is also an extension to the 392 ISO C standard. 393 Where applicable, functions are marked with the IP6 margin legend in the SYNOPSIS section. 394 Where additional semantics apply to a function, the material is identified by use of the IP6 395 396 margin legend. Advisory Information and either Memory Mapped Files or Shared Memory Objects 397 MC1 The functionality described is optional. The functionality described is also an extension to the 398 ISO C standard. 399 This is a shorthand notation for combinations of multiple option codes. 400 401 Where applicable, functions are marked with the MC1 margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the MC1 margin legend. 403 Refer to the Base Definitions volume of IEEE Std 1003.1-2001, Section 1.5.2, Margin Code 404 Notation. Memory Mapped Files, Shared Memory Objects, or Memory Protection 406 MC2 407 The functionality described is optional. The functionality described is also an extension to the ISO C standard. 408 This is a shorthand notation for combinations of multiple option codes. 409 410 Where applicable, functions are marked with the MC2 margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the MC2 margin legend. 412 413 Refer to the Base Definitions volume of IEEE Std 1003.1-2001, Section 1.5.2, Margin Code Notation. Memory Mapped Files, Shared Memory Objects, or Typed Memory Objects 415 MC3 The functionality described is optional. The functionality described is also an extension to the 416 ISO C standard. 417 This is a shorthand notation for combinations of multiple option codes. 418 Where applicable, functions are marked with the MC3 margin legend in the SYNOPSIS section. 419 Where additional semantics apply to a function, the material is identified by use of the MC3 margin legend. 421 Refer to the Base Definitions volume of IEEE Std 1003.1-2001, Section 1.5.2, Margin Code 422 Notation. 423 Memory Mapped Files 424 MF The functionality described is optional. The functionality described is also an extension to the 425 ISO C standard. 426 Where applicable, functions are marked with the MF margin legend in the SYNOPSIS section. 427 Where additional semantics apply to a function, the material is identified by use of the MF 428 margin legend. 429 **Process Memory Locking** 430 MI. The functionality described is optional. The functionality described is also an extension to the 431

Portability Introduction

432		ISO C standard.
433 434 435		Where applicable, functions are marked with the ML margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the MI margin legend.
436 437 438	MLR	Range Memory Locking The functionality described is optional. The functionality described is also an extension to the ISO C standard.
439 440 441		Where applicable, functions are marked with the MLR margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the MLF margin legend.
442 443 444	MON	Monotonic Clock The functionality described is optional. The functionality described is also an extension to the ISO C standard.
445 446 447		Where applicable, functions are marked with the MON margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the MON margin legend.
448 449 450	MPR	Memory Protection The functionality described is optional. The functionality described is also an extension to the ISO C standard.
451 452 453		Where applicable, functions are marked with the MPR margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the MPR margin legend.
454 455 456	MSG	Message Passing The functionality described is optional. The functionality described is also an extension to the ISO C standard.
457 458 459		Where applicable, functions are marked with the MSG margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the MSC margin legend.
460 461 462	MX	IEC 60559 Floating-Point Option The functionality described is optional. The functionality described is also an extension to the ISO C standard.
463 464 465		Where applicable, functions are marked with the MX margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the MX margin legend.
466 467 468 469	OB	Obsolescent The functionality described may be withdrawn in a future version of this volume of IEEE Std 1003.1-2001. Strictly Conforming POSIX Applications and Strictly Conforming XS Applications shall not use obsolescent features.
470		Where applicable, the material is identified by use of the OB margin legend.
471 472 473 474	OF	Output Format Incompletely Specified The functionality described is an XSI extension. The format of the output produced by the utility is not fully specified. It is therefore not possible to post-process this output in a consistent fashion. Typical problems include unknown length of strings and unspecified field delimiters.
475		Where applicable, the material is identified by use of the OF margin legend.

Introduction Portability

476 477	ОН	Optional Header In the SYNOPSIS section of some interfaces in the System Interfaces volume of
478		IEEE Std 1003.1-2001 an included header is marked as in the following example:
479 480 481	ОН	<pre>#include <sys types.h=""> #include <grp.h> struct group *getgrnam(const char *name);</grp.h></sys></pre>
482 483		The OH margin legend indicates that the marked header is not required on XSI-conformant systems.
484 485 486	PIO	Prioritized Input and Output The functionality described is optional. The functionality described is also an extension to the ISO C standard.
487 488 489		Where applicable, functions are marked with the PIO margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the PIO margin legend.
490 491 492	PS	Process Scheduling The functionality described is optional. The functionality described is also an extension to the ISO C standard.
493 494 495		Where applicable, functions are marked with the PS margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the PS margin legend.
496 497 498	RS	Raw Sockets The functionality described is optional. The functionality described is also an extension to the ISO C standard.
499 500 501		Where applicable, functions are marked with the RS margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the RS margin legend.
502 503 504	RTS	Realtime Signals Extension The functionality described is optional. The functionality described is also an extension to the ISO C standard.
505 506 507		Where applicable, functions are marked with the RTS margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the RTS margin legend.
508 509	SD	Software Development Utilities The functionality described is optional.
510 511 512		Where applicable, utilities are marked with the SD margin legend in the SYNOPSIS section. Where additional semantics apply to a utility, the material is identified by use of the SD margin legend.
513 514 515	SEM	Semaphores The functionality described is optional. The functionality described is also an extension to the ISO C standard.
516 517 518		Where applicable, functions are marked with the SEM margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the SEM margin legend.
519 520	SHM	Shared Memory Objects The functionality described is optional. The functionality described is also an extension to the

Portability Introduction

521		ISO C standard.
522 523 524		Where applicable, functions are marked with the SHM margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the SHM margin legend.
525 526 527	SIO	Synchronized Input and Output The functionality described is optional. The functionality described is also an extension to the ISO C standard.
528 529 530		Where applicable, functions are marked with the SIO margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the SIO margin legend.
531 532 533	SPI	Spin Locks The functionality described is optional. The functionality described is also an extension to the ISO C standard.
534 535 536		Where applicable, functions are marked with the SPI margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the SPI margin legend.
537 538 539	SPN	Spawn The functionality described is optional. The functionality described is also an extension to the ISO C standard.
540 541 542		Where applicable, functions are marked with the SPN margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the SPN margin legend.
543 544 545	SS	Process Sporadic Server The functionality described is optional. The functionality described is also an extension to the ISO C standard.
546 547 548		Where applicable, functions are marked with the SS margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the SS margin legend.
549 550 551	TCT	Thread CPU-Time Clocks The functionality described is optional. The functionality described is also an extension to the ISO C standard.
552 553 554		Where applicable, functions are marked with the TCT margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TCT margin legend.
555 556 557	TEF	Trace Event Filter The functionality described is optional. The functionality described is also an extension to the ISO C standard.
558 559 560		Where applicable, functions are marked with the TEF margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TEF margin legend.
561 562 563	THR	Threads The functionality described is optional. The functionality described is also an extension to the ISO C standard.
564 565		Where applicable, functions are marked with the THR margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the THR

Introduction Portability

566		margin legend.
567 568 569	TMO	Timeouts The functionality described is optional. The functionality described is also an extension to the ISO C standard.
570 571 572		Where applicable, functions are marked with the TMO margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TMO margin legend.
573 574 575	TMR	Timers The functionality described is optional. The functionality described is also an extension to the ISO C standard.
576 577 578		Where applicable, functions are marked with the TMR margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TMR margin legend.
579 580 581	TPI	Thread Priority Inheritance The functionality described is optional. The functionality described is also an extension to the ISO C standard.
582 583 584		Where applicable, functions are marked with the TPI margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TPI margin legend.
585 586 587	TPP	Thread Priority Protection The functionality described is optional. The functionality described is also an extension to the ISO C standard.
588 589 590		Where applicable, functions are marked with the TPP margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TPP margin legend.
591 592 593	TPS	Thread Execution Scheduling The functionality described is optional. The functionality described is also an extension to the ISO C standard.
594 595 596		Where applicable, functions are marked with the TPS margin legend for the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TPS margin legend.
597 598 599	TRC	Trace The functionality described is optional. The functionality described is also an extension to the ISO C standard.
600 601 602		Where applicable, functions are marked with the TRC margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TRC margin legend.
603 604 605	TRI	Trace Inherit The functionality described is optional. The functionality described is also an extension to the ISO C standard.
606 607 608		Where applicable, functions are marked with the TRI margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TRI margin legend.
609 610	TRL	Trace Log The functionality described is optional. The functionality described is also an extension to the

Portability Introduction

611		ISO C standard.
612 613 614		Where applicable, functions are marked with the TRL margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TRL margin legend.
615 616 617	TSA	Thread Stack Address Attribute The functionality described is optional. The functionality described is also an extension to the ISO C standard.
618 619 620		Where applicable, functions are marked with the TSA margin legend for the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TSA margin legend.
621 622 623	TSF	Thread-Safe Functions The functionality described is optional. The functionality described is also an extension to the ISO C standard.
624 625 626		Where applicable, functions are marked with the TSF margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TSF margin legend.
627 628 629	TSH	Thread Process-Shared Synchronization The functionality described is optional. The functionality described is also an extension to the ISO C standard.
630 631 632		Where applicable, functions are marked with the TSH margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TSH margin legend.
633 634 635	TSP	Thread Sporadic Server The functionality described is optional. The functionality described is also an extension to the ISO C standard.
636 637 638		Where applicable, functions are marked with the TSP margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TSP margin legend.
639 640 641	TSS	Thread Stack Size Attribute The functionality described is optional. The functionality described is also an extension to the ISO C standard.
642 643 644		Where applicable, functions are marked with the TSS margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TSS margin legend.
645 646 647	TYM	Typed Memory Objects The functionality described is optional. The functionality described is also an extension to the ISO C standard.
648 649 650		Where applicable, functions are marked with the TYM margin legend in the SYNOPSIS section. Where additional semantics apply to a function, the material is identified by use of the TYM margin legend.
651 652	UP	User Portability Utilities The functionality described is optional.
653 654 655		Where applicable, utilities are marked with the UP margin legend in the SYNOPSIS section. Where additional semantics apply to a utility, the material is identified by use of the UP margin legend.

Introduction Portability

656 XSI Extension 657 The functi 658 the ISO C 659 systems su

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

If an entire SYNOPSIS section is shaded and marked XSI, all the functionality described in that reference page is an extension. See the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.439, XSI.

663 XSR XSI STREAMS

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

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

1.9 Utility Limits

This section lists magnitude limitations imposed by a specific implementation. The braces notation, {LIMIT}, is used in this volume of IEEE Std 1003.1-2001 to indicate these values, but the braces are not part of the name.

Table 1-3 Utility Limit Minimum Values

Name	Description	Value
{POSIX2_BC_BASE_MAX}	The maximum <i>obase</i> value allowed by the <i>bc</i> utility.	99
{POSIX2_BC_DIM_MAX}	The maximum number of elements permitted in an array by the <i>bc</i> utility.	2 048
{POSIX2_BC_SCALE_MAX}	The maximum <i>scale</i> value allowed by the <i>bc</i> utility.	99
{POSIX2_BC_STRING_MAX}	The maximum length of a string constant accepted by the <i>bc</i> utility.	1 000
{POSIX2_COLL_WEIGHTS_MAX}	The maximum number of weights that can be assigned to an entry of the <i>LC_COLLATE</i> order keyword in the locale definition file; see the border_start keyword in the Base Definitions volume of IEEE Std 1003.1-2001, Section 7.3.2, LC_COLLATE.	2
{POSIX2_EXPR_NEST_MAX}	The maximum number of expressions that can be nested within parentheses by the <i>expr</i> utility.	32
{POSIX2_LINE_MAX}	Unless otherwise noted, the maximum length, in bytes, of the input line of a utility (either standard input or another file), when the utility is described as processing text files. The length includes room for the trailing <newline>.</newline>	2 048

Utility Limits Introduction

Name	Description	Value
{POSIX2_RE_DUP_MAX}	The maximum number of repeated occurrences of a BRE permitted when using the interval notation $\{m,n'\}$; see the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.3.6, BREs Matching Multiple Characters.	255

The values specified in Table 1-3 (on page 17) represent the lowest values conforming implementations shall provide and, consequently, the largest values on which an application can rely without further enquiries, as described below. These values shall be accessible to applications via the *getconf* utility (see *getconf* (on page 484)).

Implementations may provide more liberal, or less restrictive, values than shown in Table 1-3 (on page 17). These possibly more liberal values are accessible using the symbols in Table 1-4.

The <code>sysconf()</code> function defined in the System Interfaces volume of IEEE Std 1003.1-2001 or the <code>getconf</code> utility return the value of each symbol on each specific implementation. The value so retrieved is the largest, or most liberal, value that is available throughout the session lifetime, as determined at session creation. The literal names shown in the table apply only to the <code>getconf</code> utility; the high-level language binding describes the exact form of each name to be used by the interfaces in that binding.

All numeric limits defined by the System Interfaces volume of IEEE Std 1003.1-2001, such as {PATH_MAX}, shall also apply to this volume of IEEE Std 1003.1-2001. All the utilities defined by this volume of IEEE Std 1003.1-2001 are implicitly limited by these values, unless otherwise noted in the utility descriptions.

It is not guaranteed that the application can actually reach the specified limit of an implementation in any given case, or at all, as a lack of virtual memory or other resources may prevent this. The limit value indicates only that the implementation does not specifically impose any arbitrary, more restrictive limit.

Table 1-4 Symbolic Utility Limits

Name	Description	Minimum Value
{BC_BASE_MAX}	The maximum <i>obase</i> value allowed by the <i>bc</i> utility.	{POSIX2_BC_BASE_MAX}
{BC_DIM_MAX}	The maximum number of elements permitted in an array by the <i>bc</i> utility.	{POSIX2_BC_DIM_MAX}
{BC_SCALE_MAX}	The maximum <i>scale</i> value allowed by the <i>bc</i> utility.	{POSIX2_BC_SCALE_MAX}
{BC_STRING_MAX}	The maximum length of a string constant accepted by the <i>bc</i> utility.	{POSIX2_BC_STRING_MAX}
{COLL_WEIGHTS_MAX}	The maximum number of weights that can be assigned to an entry of the <i>LC_COLLATE</i> order keyword in the locale definition file; see the	{POSIX2_COLL_WEIGHTS_MAX}

Introduction Utility Limits

743			
744	Name	Description	Minimum Value
745		order_start keyword in the	
746		Base Definitions volume of	
747		IEEE Std 1003.1-2001,	
748		Section 7.3.2, LC_COLLATE.	
749	{EXPR_NEST_MAX}	The maximum number of	{POSIX2_EXPR_NEST_MAX}
750		expressions that can be	
751		nested within parentheses	
752		by the <i>expr</i> utility.	
753	{LINE_MAX}	Unless otherwise noted, the	{POSIX2_LINE_MAX}
754		maximum length, in bytes,	
755		of the input line of a utility	
756		(either standard input or	
757		another file), when the	
758		utility is described as	
759		processing text files. The	
760		length includes room for the	
761		trailing <newline>.</newline>	
762	{RE_DUP_MAX}	The maximum number of	{POSIX2_RE_DUP_MAX}
763		repeated occurrences of a	
764		BRE permitted when using	
765		the interval notation	
766		$\{m,n\}$; see the Base	
767		Definitions volume of	
768		IEEE Std 1003.1-2001,	
769		Section 9.3.6, BREs	
770		Matching Multiple	
771		Characters.	

The following value may be a constant within an implementation or may vary from one pathname to another.

{POSIX2_SYMLINKS}

When referring to a directory, the system supports the creation of symbolic links within that directory; for non-directory files, the meaning of {POSIX2_SYMLINKS} is undefined.

1.10 Grammar Conventions

Portions of this volume of IEEE Std 1003.1-2001 are expressed in terms of a special grammar notation. It is used to portray the complex syntax of certain program input. The grammar is based on the syntax used by the *yacc* utility. However, it does not represent fully functional *yacc* input, suitable for program use; the lexical processing and all semantic requirements are described only in textual form. The grammar is not based on source used in any traditional implementation and has not been tested with the semantic code that would normally be required to accompany it. Furthermore, there is no implication that the partial *yacc* code presented represents the most efficient, or only, means of supporting the complex syntax within the utility. Implementations may use other programming languages or algorithms, as long as the syntax supported is the same as that represented by the grammar.

The following typographical conventions are used in the grammar; they have no significance except to aid in reading.

Grammar Conventions Introduction

- The identifiers for the reserved words of the language are shown with a leading capital letter. (These are terminals in the grammar; for example, **While**, **Case**.)
- The identifiers for terminals in the grammar are all named with uppercase letters and underscores; for example, NEWLINE, ASSIGN_OP, NAME.
- The identifiers for non-terminals are all lowercase.

1.11 Utility Description Defaults

This section describes all of the subsections used within the utility descriptions, including:

- Intended usage of the section
- · Global defaults that affect all the standard utilities
- The meanings of notations used in this volume of IEEE Std 1003.1-2001 that are specific to individual utility sections

NAME

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

SYNOPSIS

The SYNOPSIS section summarizes the syntax of the calling sequence for the utility, including options, option-arguments, and operands. Standards for utility naming are described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines; for describing the utility's arguments in the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.1, Utility Argument Syntax.

DESCRIPTION

The DESCRIPTION section describes the actions of the utility. If the utility has a very complex set of subcommands or its own procedural language, an EXTENDED DESCRIPTION section is also provided. Most explanations of optional functionality are omitted here, as they are usually explained in the OPTIONS section.

As stated in Section 1.7.1.11 (on page 7), some functions are described in terms of equivalent functionality. When specific functions are cited, the implementation shall provide equivalent functionality including side effects associated with successful execution of the function. The treatment of errors and intermediate results from the individual functions cited is generally not specified by this volume of IEEE Std 1003.1-2001. See the utility's EXIT STATUS and CONSEQUENCES OF ERRORS sections for all actions associated with errors encountered by the utility.

OPTIONS

The OPTIONS section describes the utility options and option-arguments, and how they modify the actions of the utility. Standard utilities that have options either fully comply with the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines or describe all deviations. Apparent disagreements between functionality descriptions in the OPTIONS and DESCRIPTION (or EXTENDED DESCRIPTION) sections are always resolved in favor of the OPTIONS section.

Each OPTIONS section that uses the phrase "The ... utility shall conform to the Utility Syntax Guidelines ..." refers only to the use of the utility as specified by this volume of IEEE Std 1003.1-2001; implementation extensions should also conform to the guidelines, but may allow exceptions for historical practice.

Unless otherwise stated in the utility description, when given an option unrecognized by the implementation, or when a required option-argument is not provided, standard utilities shall issue a diagnostic message to standard error and exit with a non-zero exit status.

All utilities in this volume of IEEE Std 1003.1-2001 shall be capable of processing arguments using eight-bit transparency.

Default Behavior: When this section is listed as "None.", it means that the implementation need not support any options. Standard utilities that do not accept options, but that do accept operands, shall recognize "--" as a first argument to be discarded.

The requirement for recognizing "--" is because conforming applications need a way to shield their operands from any arbitrary options that the implementation may provide as an extension. For example, if the standard utility *foo* is listed as taking no options, and the application needed to give it a pathname with a leading hyphen, it could safely do it as:

```
foo -- -myfile
```

and avoid any problems with -m used as an extension.

OPERANDS

The OPERANDS section describes the utility operands, and how they affect the actions of the utility. Apparent disagreements between functionality descriptions in the OPERANDS and DESCRIPTION (or EXTENDED DESCRIPTION) sections shall be resolved in favor of the OPERANDS section.

If an operand naming a file can be specified as '-', which means to use the standard input instead of a named file, this is explicitly stated in this section. Unless otherwise stated, the use of multiple instances of '-' to mean standard input in a single command produces unspecified results.

Unless otherwise stated, the standard utilities that accept operands shall process those operands in the order specified in the command line.

Default Behavior: When this section is listed as "None.", it means that the implementation need not support any operands.

STDIN

The STDIN section describes the standard input of the utility. This section is frequently merely a reference to the following section, as many utilities treat standard input and input files in the same manner. Unless otherwise stated, all restrictions described in the INPUT FILES section shall apply to this section as well.

Use of a terminal for standard input can cause any of the standard utilities that read standard input to stop when used in the background. For this reason, applications should not use interactive features in scripts to be placed in the background.

The specified standard input format of the standard utilities shall not depend on the existence or value of the environment variables defined in this volume of IEEE Std 1003.1-2001, except as provided by this volume of IEEE Std 1003.1-2001.

Default Behavior: When this section is listed as "Not used.", it means that the standard input shall not be read when the utility is used as described by this volume of IEEE Std 1003.1-2001.

INPUT FILES

The INPUT FILES section describes the files, other than the standard input, used as input by the utility. It includes files named as operands and option-arguments as well as other files that are referred to, such as start-up and initialization files, databases, and so on. Commonly-used files are generally described in one place and cross-referenced by other utilities.

All utilities in this volume of IEEE Std 1003.1-2001 shall be capable of processing input files using eight-bit transparency.

When a standard utility reads a seekable input file and terminates without an error before it reaches end-of-file, the utility shall ensure that the file offset in the open file description is properly positioned just past the last byte processed by the utility. For files that are not seekable, the state of the file offset in the open file description for that file is unspecified. A conforming application shall not assume that the following three commands are equivalent:

```
tail -n +2 file
(sed -n 1q; cat) < file
cat file | (sed -n 1q; cat)</pre>
```

The second command is equivalent to the first only when the file is seekable. The third command leaves the file offset in the open file description in an unspecified state. Other utilities, such as *head*, *read*, and *sh*, have similar properties.

Some of the standard utilities, such as filters, process input files a line or a block at a time and have no restrictions on the maximum input file size. Some utilities may have size limitations that are not as obvious as file space or memory limitations. Such limitations should reflect resource limitations of some sort, not arbitrary limits set by implementors. Implementations shall document those utilities that are limited by constraints other than file system space, available memory, and other limits specifically cited by this volume of IEEE Std 1003.1-2001, and identify what the constraint is and indicate a way of estimating when the constraint would be reached. Similarly, some utilities descend the directory tree (recursively). Implementations shall also document any limits that they may have in descending the directory tree that are beyond limits cited by this volume of IEEE Std 1003.1-2001.

When an input file is described as a "text file", the utility produces undefined results if given input that is not from a text file, unless otherwise stated. Some utilities (for example, *make*, *read*, *sh*) allow for continued input lines using an escaped <newline> convention; unless otherwise stated, the utility need not be able to accumulate more than {LINE_MAX} bytes from a set of multiple, continued input lines. Thus, for a conforming application the total of all the continued lines in a set cannot exceed {LINE_MAX}. If a utility using the escaped <newline> convention detects an end-of-file condition immediately after an escaped <newline>, the results are unspecified.

Record formats are described in a notation similar to that used by the C-language function, <code>printf()</code>. See the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5, File Format Notation for a description of this notation. The format description is intended to be sufficiently rigorous to allow other applications to generate these input files. However, since <code><blank>s</code> can legitimately be included in some of the fields described by the standard utilities, particularly in locales other than the POSIX locale, this intent is not always realized.

Default Behavior: When this section is listed as "None.", it means that no input files are required to be supplied when the utility is used as described by this volume of

XSI

IEEE Std 1003.1-2001.

ENVIRONMENT VARIABLES

The ENVIRONMENT VARIABLES section lists what variables affect the utility's execution.

The entire manner in which environment variables described in this volume of IEEE Std 1003.1-2001 affect the behavior of each utility is described in the ENVIRONMENT VARIABLES section for that utility, in conjunction with the global effects of the *LANG*, *LC_ALL*, and *NLSPATH* environment variables described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables. The existence or value of environment variables described in this volume of IEEE Std 1003.1-2001 shall not otherwise affect the specified behavior of the standard utilities. Any effects of the existence or value of environment variables not described by this volume of IEEE Std 1003.1-2001 upon the standard utilities are unspecified.

For those standard utilities that use environment variables as a means for selecting a utility to execute (such as *CC* in *make*), the string provided to the utility is subjected to the path search described for *PATH* in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables.

All utilities in this volume of IEEE Std 1003.1-2001 shall be capable of processing environment variable names and values using eight-bit transparency.

Default Behavior: When this section is listed as "None.", it means that the behavior of the utility is not directly affected by environment variables described by this volume of IEEE Std 1003.1-2001 when the utility is used as described by this volume of IEEE Std 1003.1-2001.

ASYNCHRONOUS EVENTS

The ASYNCHRONOUS EVENTS section lists how the utility reacts to such events as signals and what signals are caught.

Default Behavior: When this section is listed as "Default.", or it refers to "the standard action for all other signals; see Section 1.11 (on page 20)" it means that the action taken as a result of the signal shall be one of the following:

- 1. The action shall be that inherited from the parent according to the rules of inheritance of signal actions defined in the System Interfaces volume of IEEE Std 1003.1-2001.
- 2. When no action has been taken to change the default, the default action shall be that specified by the System Interfaces volume of IEEE Std 1003.1-2001.
- 3. The result of the utility's execution is as if default actions had been taken.

A utility is permitted to catch a signal, perform some additional processing (such as deleting temporary files), restore the default signal action (or action inherited from the parent process), and resignal itself.

STDOUT

The STDOUT section completely describes the standard output of the utility. This section is frequently merely a reference to the following section, OUTPUT FILES, because many utilities treat standard output and output files in the same manner.

Use of a terminal for standard output may cause any of the standard utilities that write standard output to stop when used in the background. For this reason, applications should not use interactive features in scripts to be placed in the background.

Record formats are described in a notation similar to that used by the C-language function, *printf()*. See the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5, File Format Notation for a description of this notation.

The specified standard output of the standard utilities shall not depend on the existence or value of the environment variables defined in this volume of IEEE Std 1003.1-2001, except as provided by this volume of IEEE Std 1003.1-2001.

Some of the standard utilities describe their output using the verb display, defined in the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.132, Display. Output described in the STDOUT sections of such utilities may be produced using means other than standard output. When standard output is directed to a terminal, the output described shall be written directly to the terminal. Otherwise, the results are undefined.

Default Behavior: When this section is listed as "Not used.", it means that the standard output shall not be written when the utility is used as described by this volume of IEEE Std 1003.1-2001.

STDERR

The STDERR section describes the standard error output of the utility. Only those messages that are purposely sent by the utility are described.

Use of a terminal for standard error may cause any of the standard utilities that write standard error output to stop when used in the background. For this reason, applications should not use interactive features in scripts to be placed in the background.

The format of diagnostic messages for most utilities is unspecified, but the language and cultural conventions of diagnostic and informative messages whose format is unspecified by this volume of IEEE Std 1003.1-2001 should be affected by the setting of LC_MESSAGES and NLSPATH.

The specified standard error output of standard utilities shall not depend on the existence or value of the environment variables defined in this volume of IEEE Std 1003.1-2001, except as provided by this volume of IEEE Std 1003.1-2001.

Default Behavior: When this section is listed as "The standard error shall be used only for diagnostic messages.", it means that, unless otherwise stated, the diagnostic messages shall be sent to the standard error only when the exit status is non-zero and the utility is used as described by this volume of IEEE Std 1003.1-2001.

When this section is listed as "Not used.", it means that the standard error shall not be used when the utility is used as described in this volume of IEEE Std 1003.1-2001.

OUTPUT FILES

The OUTPUT FILES section completely describes the files created or modified by the utility. Temporary or system files that are created for internal usage by this utility or other parts of the implementation (for example, spool, log, and audit files) are not described in this, or any, section. The utilities creating such files and the names of such files are unspecified. If applications are written to use temporary or intermediate files, they should use the TMPDIR environment variable, if it is set and represents an accessible directory, to select the location of temporary files.

Implementations shall ensure that temporary files, when used by the standard utilities, are named so that different utilities or multiple instances of the same utility can operate simultaneously without regard to their working directories, or any other process characteristic other than process ID. There are two exceptions to this rule:

991

969

970

972

973

974

975

976

977 978

979

980

981

982

983

984

985

986

988

989

992

995

996

993 994

XSI

997 998 999

1000 1001 1002

1004 1005 1006

1007

1008

1003

1009 1010 1011

1012

1013 1014

24

- 10.15 Resources for temporary files other than the name space (for example, disk space, available directory entries, or number of processes allowed) are not guaranteed.
 - Certain standard utilities generate output files that are intended as input for other utilities (for example, *lex* generates *lex.yy.c*), and these cannot have unique names. These cases are explicitly identified in the descriptions of the respective utilities.

Any temporary file created by the implementation shall be removed by the implementation upon a utility's successful exit, exit because of errors, or before termination by any of the SIGHUP, SIGINT, or SIGTERM signals, unless specified otherwise by the utility description.

Receipt of the SIGQUIT signal should generally cause termination (unless in some debugging mode) that would bypass any attempted recovery actions.

Record formats are described in a notation similar to that used by the C-language function, *printf*(); see the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5, File Format Notation for a description of this notation.

Default Behavior: When this section is listed as "None.", it means that no files are created or modified as a consequence of direct action on the part of the utility when the utility is used as described by this volume of IEEE Std 1003.1-2001. However, the utility may create or modify system files, such as log files, that are outside the utility's normal execution environment.

EXTENDED DESCRIPTION

The EXTENDED DESCRIPTION section provides a place for describing the actions of very complicated utilities, such as text editors or language processors, which typically have elaborate command languages.

Default Behavior: When this section is listed as "None.", no further description is necessary.

EXIT STATUS

The EXIT STATUS section describes the values the utility shall return to the calling program, or shell, and the conditions that cause these values to be returned. Usually, utilities return zero for successful completion and values greater than zero for various error conditions. If specific numeric values are listed in this section, the system shall use those values for the errors described. In some cases, status values are listed more loosely, such as >0. A strictly conforming application shall not rely on any specific value in the range shown and shall be prepared to receive any value in the range.

For example, a utility may list zero as a successful return, 1 as a failure for a specific reason, and >1 as "an error occurred". In this case, unspecified conditions may cause a 2 or 3, or other value, to be returned. A conforming application should be written so that it tests for successful exit status values (zero in this case), rather than relying upon the single specific error value listed in this volume of IEEE Std 1003.1-2001. In that way, it has maximum portability, even on implementations with extensions.

Unspecified error conditions may be represented by specific values not listed in this volume of IEEE Std 1003.1-2001.

CONSEQUENCES OF ERRORS

The CONSEQUENCES OF ERRORS section describes the effects on the environment, file systems, process state, and so on, when error conditions occur. It does not describe error messages produced or exit status values used.

 The many reasons for failure of a utility are generally not specified by the utility descriptions. Utilities may terminate prematurely if they encounter: invalid usage of options, arguments, or environment variables; invalid usage of the complex syntaxes expressed in EXTENDED DESCRIPTION sections; difficulties accessing, creating, reading, or writing files; or difficulties associated with the privileges of the process.

The following shall apply to each utility, unless otherwise stated:

• If the requested action cannot be performed on an operand representing a file, directory, user, process, and so on, the utility shall issue a diagnostic message to standard error and continue processing the next operand in sequence, but the final exit status shall be returned as non-zero.

For a utility that recursively traverses a file hierarchy (such as find or chown $-\mathbf{R}$), if the requested action cannot be performed on a file or directory encountered in the hierarchy, the utility shall issue a diagnostic message to standard error and continue processing the remaining files in the hierarchy, but the final exit status shall be returned as non-zero.

- If the requested action characterized by an option or option-argument cannot be performed, the utility shall issue a diagnostic message to standard error and the exit status returned shall be non-zero.
- When an unrecoverable error condition is encountered, the utility shall exit with a non-zero exit status.
- A diagnostic message shall be written to standard error whenever an error condition occurs.

When a utility encounters an error condition several actions are possible, depending on the severity of the error and the state of the utility. Included in the possible actions of various utilities are: deletion of temporary or intermediate work files; deletion of incomplete files; validity checking of the file system or directory.

Default Behavior: When this section is listed as "Default.", it means that any changes to the environment are unspecified.

APPLICATION USAGE

This section is informative.

The APPLICATION USAGE section gives advice to the application programmer or user about the way the utility should be used.

EXAMPLES

This section is informative.

The EXAMPLES section gives one or more examples of usage, where appropriate. In the event of conflict between an example and a normative part of the specification, the normative material is to be taken as correct.

In all examples, quoting has been used, showing how sample commands (utility names combined with arguments) could be passed correctly to a shell (see *sh*) or as a string to the *system*() function defined in the System Interfaces volume of IEEE Std 1003.1-2001. Such quoting would not be used if the utility is invoked using one of the *exec* functions defined in the System Interfaces volume of IEEE Std 1003.1-2001.

RATIONALE

This section is informative.

1109

1110

1111 1112

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123 1124

1126

1127

1128

1129

This section contains historical information concerning the contents of this volume of IEEE Std 1003.1-2001 and why features were included or discarded by the standard developers.

FUTURE DIRECTIONS

This section is informative.

The FUTURE DIRECTIONS section should be used as a guide to current thinking; there is not necessarily a commitment to implement all of these future directions in their entirety.

SEE ALSO

This section is informative.

The SEE ALSO section lists related entries.

CHANGE HISTORY

This section is informative.

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

Certain of the standard utilities describe how they can invoke other utilities or applications, such as by passing a command string to the command interpreter. The external influences (STDIN, ENVIRONMENT VARIABLES, and so on) and external effects (STDOUT, CONSEQUENCES OF ERRORS, and so on) of such invoked utilities are not described in the section concerning the standard utility that invokes them.

1125 1.12 Considerations for Utilities in Support of Files of Arbitrary Size

The following utilities support files of any size up to the maximum that can be created by the implementation. This support includes correct writing of file size-related values (such as file sizes and offsets, line numbers, and block counts) and correct interpretation of command line arguments that contain such values.

	8	
1130	basename	Return non-directory portion of pathname.
1131	cat	Concatenate and print files.
1132	cd	Change working directory.
1133	chgrp	Change file group ownership.
1134	chmod	Change file modes.
1135	chown	Change file ownership.
1136	cksum	Write file checksums and sizes.
1137	cmp	Compare two files.
1138	ср	Copy files.
1139	dd	Convert and copy a file.
1140	df	Report free disk space.
1141	dirname	Return directory portion of pathname.
1142	du	Estimate file space usage.

1143	find	Find files.
1144	ln	Link files.
1145	ls	List directory contents.
1146	mkdir	Make directories.
1147	mv	Move files.
1148	pathchk	Check pathnames.
1149	pwd	Return working directory name.
1150	rm	Remove directory entries.
1151	rmdir	Remove directories.
1152	sh	Shell, the standard command language interpreter.
1153	sum	Print checksum and block or byte count of a file.
1154	test	Evaluate expression.
1155	touch	Change file access and modification times.
1156	ulimit	Set or report file size limit.
1157 1158	Exceptions to follows:	to the requirement that utilities support files of any size up to the maximum are as
1159	1. Uses o	f files as command scripts, or for configuration or control, are exempt. For example,

- 1. Uses of files as command scripts, or for configuration or control, are exempt. For example, it is not required that *sh* be able to read an arbitrarily large **.profile**.
- 2. Shell input and output redirection are exempt. For example, it is not required that the redirections *sum* < *file* or *echo foo* > *file* succeed for an arbitrarily large existing file.

1.13 **Built-In Utilities** 1163

1160

1161

1162

1164

1165

1166 1167

1168

1169

1170

1175

1176 1177

1178

1179

Any of the standard utilities may be implemented as regular built-in utilities within the command language interpreter. This is usually done to increase the performance of frequently used utilities or to achieve functionality that would be more difficult in a separate environment. The utilities named in Table 1-5 are frequently provided in built-in form. All of the utilities named in the table have special properties in terms of command search order within the shell, as described in Section 2.9.1.1 (on page 48).

Table 1-5 Regular Built-In Utilities

1171	alias	false	jobs	read	wait
1172	bg	fc	kill	true	
1173	$c\bar{d}$	fg	newgrp	umask	
1174	command	getopts	pwd	unalias	

However, all of the standard utilities, including the regular built-ins in the table, but not the special built-ins described in Section 2.14 (on page 64), shall be implemented in a manner so that they can be accessed via the *exec* family of functions as defined in the System Interfaces volume of IEEE Std 1003.1-2001 and can be invoked directly by those standard utilities that require it (env, find, nice, nohup, time, xargs).

 This chapter contains the definition of the Shell Command Language.

1182 2.1 Shell Introduction

The shell is a command language interpreter. This chapter describes the syntax of that command language as it is used by the *sh* utility and the *system()* and *popen()* functions defined in the System Interfaces volume of IEEE Std 1003.1-2001.

The shell operates according to the following general overview of operations. The specific details are included in the cited sections of this chapter.

- 1. The shell reads its input from a file (see *sh*), from the -c option or from the *system*() and *popen*() functions defined in the System Interfaces volume of IEEE Std 1003.1-2001. If the first line of a file of shell commands starts with the characters "#!", the results are unspecified.
- 2. The shell breaks the input into tokens: words and operators; see Section 2.3 (on page 31).
- 3. The shell parses the input into simple commands (see Section 2.9.1 (on page 47)) and compound commands (see Section 2.9.4 (on page 52)).
- 4. The shell performs various expansions (separately) on different parts of each command, resulting in a list of pathnames and fields to be treated as a command and arguments; see Section 2.6 (on page 36).
- 5. The shell performs redirection (see Section 2.7 (on page 43)) and removes redirection operators and their operands from the parameter list.
- 6. The shell executes a function (see Section 2.9.5 (on page 54)), built-in (see Section 2.14 (on page 64)), executable file, or script, giving the names of the arguments as positional parameters numbered 1 to *n*, and the name of the command (or in the case of a function within a script, the name of the script) as the positional parameter numbered 0 (see Section 2.9.1.1 (on page 48)).
- 7. The shell optionally waits for the command to complete and collects the exit status (see Section 2.8.2 (on page 46)).

2.2 Quoting

Quoting is used to remove the special meaning of certain characters or words to the shell. Quoting can be used to preserve the literal meaning of the special characters in the next paragraph, prevent reserved words from being recognized as such, and prevent parameter expansion and command substitution within here-document processing (see Section 2.7.4 (on page 44)).

The application shall quote the following characters if they are to represent themselves:

```
| & ; < > ( ) $ ` \ " ' <space> <tab> <newline>
```

and the following may need to be quoted under certain circumstances. That is, these characters may be special depending on conditions described elsewhere in this volume of IEEE Std 1003.1-2001:

```
* ? [ # ~ = %
```

The various quoting mechanisms are the escape character, single-quotes, and double-quotes. The here-document represents another form of quoting; see Section 2.7.4 (on page 44).

2.2.1 Escape Character (Backslash)

A backslash that is not quoted shall preserve the literal value of the following character, with the exception of a <newline>. If a <newline> follows the backslash, the shell shall interpret this as line continuation. The backslash and <newline>s shall be removed before splitting the input into tokens. Since the escaped <newline> is removed entirely from the input and is not replaced by any white space, it cannot serve as a token separator.

1227 2.2.2 Single-Quotes

Enclosing characters in single-quotes (' ') shall preserve the literal value of each character within the single-quotes. A single-quote cannot occur within single-quotes.

1230 2.2.3 Double-Quotes

Enclosing characters in double-quotes ("") shall preserve the literal value of all characters within the double-quotes, with the exception of the characters dollar sign, backquote, and backslash, as follows:

\$ The dollar sign shall retain its special meaning introducing parameter expansion (see Section 2.6.2 (on page 37)), a form of command substitution (see Section 2.6.3 (on page 40)), and arithmetic expansion (see Section 2.6.4 (on page 41)).

The input characters within the quoted string that are also enclosed between "\$ (" and the matching ')' shall not be affected by the double-quotes, but rather shall define that command whose output replaces the "\$ (...)" when the word is expanded. The tokenizing rules in Section 2.3 (on page 31), not including the alias substitutions in Section 2.3.1 (on page 32), shall be applied recursively to find the matching ')'.

Within the string of characters from an enclosed " $$\{$ " to the matching ' $\}$ ', an even number of unescaped double-quotes or single-quotes, if any, shall occur. A preceding backslash character shall be used to escape a literal ' $\{$ ' or ' $\}$ '. The rule in Section 2.6.2 (on page 37) shall be used to determine the matching ' $\{$ '.

The backquote shall retain its special meaning introducing the other form of command substitution (see Section 2.6.3 (on page 40)). The portion of the quoted string from the initial backquote and the characters up to the next backquote that is not preceded by a backslash,

having escape characters removed, defines that command whose output replaces "`...`" when the word is expanded. Either of the following cases produces undefined results:

- A single-quoted or double-quoted string that begins, but does not end, within the "`...` "sequence
- A "'...'" sequence that begins, but does not end, within the same double-quoted string
- The backslash shall retain its special meaning as an escape character (see Section 2.2.1 (on page 30)) only when followed by one of the following characters when considered special:

```
$ ' " \ <newline>
```

The application shall ensure that a double-quote is preceded by a backslash to be included within double-quotes. The parameter '@' has special meaning inside double-quotes and is described in Section 2.5.2 (on page 34).

2.3 Token Recognition

The shell shall read its input in terms of lines from a file, from a terminal in the case of an interactive shell, or from a string in the case of sh –c or system(). The input lines can be of unlimited length. These lines shall be parsed using two major modes: ordinary token recognition and processing of here-documents.

When an **io_here** token has been recognized by the grammar (see Section 2.10 (on page 55)), one or more of the subsequent lines immediately following the next **NEWLINE** token form the body of one or more here-documents and shall be parsed according to the rules of Section 2.7.4 (on page 44).

When it is not processing an **io_here**, the shell shall break its input into tokens by applying the first applicable rule below to the next character in its input. The token shall be from the current position in the input until a token is delimited according to one of the rules below; the characters forming the token are exactly those in the input, including any quoting characters. If it is indicated that a token is delimited, and no characters have been included in a token, processing shall continue until an actual token is delimited.

- 1. If the end of input is recognized, the current token shall be delimited. If there is no current token, the end-of-input indicator shall be returned as the token.
- 2. If the previous character was used as part of an operator and the current character is not quoted and can be used with the current characters to form an operator, it shall be used as part of that (operator) token.
- 3. If the previous character was used as part of an operator and the current character cannot be used with the current characters to form an operator, the operator containing the previous character shall be delimited.
- 4. If the current character is backslash, single-quote, or double-quote ('\',''', or '"') and it is not quoted, it shall affect quoting for subsequent characters up to the end of the quoted text. The rules for quoting are as described in Section 2.2 (on page 30). During token recognition no substitutions shall be actually performed, and the result token shall contain exactly the characters that appear in the input (except for <newline> joining), unmodified, including any embedded or enclosing quotes or substitution operators, between the quote mark and the end of the quoted text. The token shall not be delimited by the end of the quoted field.

- 5. If the current character is an unquoted '\$' or '\', the shell shall identify the start of any candidates for parameter expansion (Section 2.6.2 (on page 37)), command substitution (Section 2.6.3 (on page 40)), or arithmetic expansion (Section 2.6.4 (on page 41)) from their introductory unquoted character sequences: '\$' or "\${ ", "\$ (" or '\', and "\$ ((", respectively. The shell shall read sufficient input to determine the end of the unit to be expanded (as explained in the cited sections). While processing the characters, if instances of expansions or quoting are found nested within the substitution, the shell shall recursively process them in the manner specified for the construct that is found. The characters found from the beginning of the substitution to its end, allowing for any recursion necessary to recognize embedded constructs, shall be included unmodified in the result token, including any embedded or enclosing substitution operators or quotes. The token shall not be delimited by the end of the substitution.
- 6. If the current character is not quoted and can be used as the first character of a new operator, the current token (if any) shall be delimited. The current character shall be used as the beginning of the next (operator) token.
- 7. If the current character is an unquoted <newline>, the current token shall be delimited.
- 8. If the current character is an unquoted <blank>, any token containing the previous character is delimited and the current character shall be discarded.
- 9. If the previous character was part of a word, the current character shall be appended to that word.
- 10. If the current character is a '#', it and all subsequent characters up to, but excluding, the next <newline> shall be discarded as a comment. The <newline> that ends the line is not considered part of the comment.
- 11. The current character is used as the start of a new word.
- Once a token is delimited, it is categorized as required by the grammar in Section 2.10 (on page 55).

1318 2.3.1 Alias Substitution

The processing of aliases shall be supported on all XSI-conformant systems or if the system supports the User Portability Utilities option (and the rest of this section is not further shaded for these options).

After a token has been delimited, but before applying the grammatical rules in Section 2.10 (on page 55), a resulting word that is identified to be the command name word of a simple command shall be examined to determine whether it is an unquoted, valid alias name. However, reserved words in correct grammatical context shall not be candidates for alias substitution. A valid alias name (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.10, Alias Name) shall be one that has been defined by the *alias* utility and not subsequently undefined using *unalias*. Implementations also may provide predefined valid aliases that are in effect when the shell is invoked. To prevent infinite loops in recursive aliasing, if the shell is not currently processing an alias of the same name, the word shall be replaced by the value of the alias; otherwise, it shall not be replaced.

If the value of the alias replacing the word ends in a <black>, the shell shall check the next command word for alias substitution; this process shall continue until a word is found that is not a valid alias or an alias value does not end in a <black>.

When used as specified by this volume of IEEE Std 1003.1-2001, alias definitions shall not be inherited by separate invocations of the shell or by the utility execution environments invoked by the shell; see Section 2.12 (on page 61).

2.4 Reserved Words

1338

1339 1340

1345

1346

1347

1348 1349

1350

1355

1356

1358 1359

1360 1361

1363

1364

1365

1366

1367

1368 1369 Reserved words are words that have special meaning to the shell; see Section 2.9 (on page 47). The following words shall be recognized as reserved words:

1341	!	do	esac	in
1342	{	done	fi	then
1343	}	elif	for	until
1344	case	else	if	while

This recognition shall only occur when none of the characters is quoted and when the word is used as:

- The first word of a command
- The first word following one of the reserved words other than case, for, or in
- The third word in a **case** command (only **in** is valid in this case)
- The third word in a **for** command (only **in** and **do** are valid in this case)
- See the grammar in Section 2.10 (on page 55).

The following words may be recognized as reserved words on some implementations (when none of the characters are quoted), causing unspecified results:

1354 [[]] function select

Words that are the concatenation of a name and a colon (':') are reserved; their use produces unspecified results.

1357 **2.5** Parameters and Variables

A parameter can be denoted by a name, a number, or one of the special characters listed in Section 2.5.2 (on page 34). A variable is a parameter denoted by a name.

A parameter is set if it has an assigned value (null is a valid value). Once a variable is set, it can only be unset by using the *unset* special built-in command.

1362 2.5.1 Positional Parameters

A positional parameter is a parameter denoted by the decimal value represented by one or more digits, other than the single digit 0. The digits denoting the positional parameters shall always be interpreted as a decimal value, even if there is a leading zero. When a positional parameter with more than one digit is specified, the application shall enclose the digits in braces (see Section 2.6.2 (on page 37)). Positional parameters are initially assigned when the shell is invoked (see *sh*), temporarily replaced when a shell function is invoked (see Section 2.9.5 (on page 54)), and can be reassigned with the *set* special built-in command.

2.5.2 Special Parameters

 Listed below are the special parameters and the values to which they shall expand. Only the values of the special parameters are listed; see Section 2.6 (on page 36) for a detailed summary of all the stages involved in expanding words.

- Expands to the positional parameters, starting from one. When the expansion occurs within double-quotes, and where field splitting (see Section 2.6.5 (on page 42)) is performed, each positional parameter shall expand as a separate field, with the provision that the expansion of the first parameter shall still be joined with the beginning part of the original word (assuming that the expanded parameter was embedded within a word), and the expansion of the last parameter shall still be joined with the last part of the original word. If there are no positional parameters, the expansion of '@' shall generate zero fields, even when '@' is double-quoted.
- * Expands to the positional parameters, starting from one. When the expansion occurs within a double-quoted string (see Section 2.2.3 (on page 30)), it shall expand to a single field with the value of each parameter separated by the first character of the *IFS* variable, or by a <space> if *IFS* is unset. If *IFS* is set to a null string, this is not equivalent to unsetting it; its first character does not exist, so the parameter values are concatenated.
- # Expands to the decimal number of positional parameters. The command name (parameter 0) shall not be counted in the number given by '#' because it is a special parameter, not a positional parameter.
- ? Expands to the decimal exit status of the most recent pipeline (see Section 2.9.2 (on page 49)).
- (Hyphen.) Expands to the current option flags (the single-letter option names concatenated into a string) as specified on invocation, by the *set* special built-in command, or implicitly by the shell.
- \$ Expands to the decimal process ID of the invoked shell. In a subshell (see Section 2.12 (on page 61)), '\$' shall expand to the same value as that of the current shell.
- ! Expands to the decimal process ID of the most recent background command (see Section 2.9.3 (on page 50)) executed from the current shell. (For example, background commands executed from subshells do not affect the value of "\$!" in the current shell environment.) For a pipeline, the process ID is that of the last command in the pipeline.
- 0 (Zero.) Expands to the name of the shell or shell script. See *sh* (on page 850) for a detailed description of how this name is derived.

See the description of the *IFS* variable in Section 2.5.3.

1404 2.5.3 Shell Variables

Variables shall be initialized from the environment (as defined by the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables and the *exec* function in the System Interfaces volume of IEEE Std 1003.1-2001) and can be given new values with variable assignment commands. If a variable is initialized from the environment, it shall be marked for export immediately; see the *export* special built-in. New variables can be defined and initialized with variable assignments, with the *read* or *getopts* utilities, with the *name* parameter in a **for** loop, with the \${name=word}\$ expansion, or with other mechanisms provided as implementation extensions.

The following variables shall affect the execution of the shell:

1414 UP XSI 1415 1416	ENV	The processing of the <i>ENV</i> shell variable shall be supported on all XSI-conformant systems or if the system supports the User Portability Utilities option.
1417 1418 1419 1420 1421 1422 1423		This variable, when and only when an interactive shell is invoked, shall be subjected to parameter expansion (see Section 2.6.2 (on page 37)) by the shell and the resulting value shall be used as a pathname of a file containing shell commands to execute in the current environment. The file need not be executable. If the expanded value of <i>ENV</i> is not an absolute pathname, the results are unspecified. <i>ENV</i> shall be ignored if the user's real and effective user IDs or real and effective group IDs are different.
1424 1425	HOME	The pathname of the user's home directory. The contents of <i>HOME</i> are used in tilde expansion (see Section 2.6.1 (on page 37)).
1426 1427 1428 1429 1430 1431	IFS	(Input Field Separators.) A string treated as a list of characters that is used for field splitting and to split lines into fields with the <i>read</i> command. If <i>IFS</i> is not set, the shell shall behave as if the value of <i>IFS</i> is <space>, <tab>, and <newline>; see Section 2.6.5 (on page 42). Implementations may ignore the value of <i>IFS</i> in the environment at the time the shell is invoked, treating <i>IFS</i> as if it were not set.</newline></tab></space>
1432 1433 1434 1435	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
1436 1437 1438	LC_ALL	The value of this variable overrides the LC^* variables and $LANG$, as described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables.
1439 1440	LC_COLLATE	Determine the behavior of range expressions, equivalence classes, and multi- character collating elements within pattern matching.
1441 1442 1443 1444 1445 1446 1447 1448	LC_CTYPE	Determine the interpretation of sequences of bytes of text data as characters (for example, single-byte as opposed to multi-byte characters), which characters are defined as letters (character class alpha) and <blank>s (character class blank), and the behavior of character classes within pattern matching. Changing the value of LC_CTYPE after the shell has started shall not affect the lexical processing of shell commands in the current shell execution environment or its subshells. Invoking a shell script or performing exec sh subjects the new shell to the changes in LC_CTYPE.</blank>
1449	LC_MESSAGES	Determine the language in which messages should be written.
1450 1451 1452 1453 1454 1455 1456	LINENO	Set by the shell to a decimal number representing the current sequential line number (numbered starting with 1) within a script or function before it executes each command. If the user unsets or resets <i>LINENO</i> , the variable may lose its special meaning for the life of the shell. If the shell is not currently executing a script or function, the value of <i>LINENO</i> is unspecified. This volume of IEEE Std 1003.1-2001 specifies the effects of the variable only for systems supporting the User Portability Utilities option.
1457 XSI 1458	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
1459 1460	PATH	A string formatted as described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables, used to effect

command interpretation; see Section 2.9	01111 (011 page 10)1
1462 PPID Set by the shell to the decimal process II	
In a subshell (see Section 2.12 (on page 1991)	
value as that of the parent of the curre	
<u></u>	e same value. This volume of
	•
1468 <i>PS1</i> Each time an interactive shell is ready	
variable shall be subjected to paramet	
error. The default value shall be "\$ ".	
implementation-defined privileges,	
implementation-defined value. The sh	
character '!' in <i>PS1</i> with the history f	
typed. Escaping the '!' with another '	
character '!' in the prompt. This vol	
the effects of the variable only for sys	stems supporting the User Portability
1477 Utilities option.	
1478 <i>PS2</i> Each time the user enters a <newline></newline>	prior to completing a command line in
an interactive shell, the value of this value	ariable shall be subjected to parameter
expansion and written to standard er	rror. The default value is "> ". This
volume of IEEE Std 1003.1-2001 specifi	ies the effects of the variable only for
systems supporting the User Portability	Utilities option.
1483 $PS4$ When an execution trace (set $-x$) is be	eing performed in an interactive shell,
before each line in the execution trace	ce, the value of this variable shall be
subjected to parameter expansion and	
value is "+ ". This volume of IEEE Sto	
variable only for systems supporting the	e User Portability Utilities option.
1488 <i>PWD</i> Set by the shell to be an absolute pathr	name of the current working directory,
containing no components of type sym	abolic link, no components that are dot,
and no components that are dot-dot	t when the shell is initialized. If an
application sets or unsets the value of	<i>PWD</i> , the behaviors of the <i>cd</i> and <i>pwd</i>
1492 utilities are unspecified.	•

2.6 Word Expansions

This section describes the various expansions that are performed on words. Not all expansions are performed on every word, as explained in the following sections.

Tilde expansions, parameter expansions, command substitutions, arithmetic expansions, and quote removals that occur within a single word expand to a single field. It is only field splitting or pathname expansion that can create multiple fields from a single word. The single exception to this rule is the expansion of the special parameter '@' within double-quotes, as described in Section 2.5.2 (on page 34).

The order of word expansion shall be as follows:

1. Tilde expansion (see Section 2.6.1 (on page 37)), parameter expansion (see Section 2.6.2 (on page 37)), command substitution (see Section 2.6.3 (on page 40)), and arithmetic expansion (see Section 2.6.4 (on page 41)) shall be performed, beginning to end. See item 5 in Section 2.3 (on page 31).

1509

1510

1511

1512 1513

1514

1515

1516

1517

1518

1519

1520 1521

1522

1523

1524 1525

1526

1527

1529

1530

1531 1532

15331534

1535

1536

1537

15381539

1540 1541

1542

1543

1544

1546

1547

1548

1549

- 2. Field splitting (see Section 2.6.5 (on page 42)) shall be performed on the portions of the fields generated by step 1, unless *IFS* is null.
 - 3. Pathname expansion (see Section 2.6.6 (on page 42)) shall be performed, unless *set* –**f** is in effect.
 - 4. Quote removal (see Section 2.6.7 (on page 42)) shall always be performed last.

The expansions described in this section shall occur in the same shell environment as that in which the command is executed.

If the complete expansion appropriate for a word results in an empty field, that empty field shall be deleted from the list of fields that form the completely expanded command, unless the original word contained single-quote or double-quote characters.

The '\$' character is used to introduce parameter expansion, command substitution, or arithmetic evaluation. If an unquoted '\$' is followed by a character that is either not numeric, the name of one of the special parameters (see Section 2.5.2 (on page 34)), a valid first character of a variable name, a left curly brace (' $\{'\}$) or a left parenthesis, the result is unspecified.

2.6.1 Tilde Expansion

A "tilde-prefix" consists of an unquoted tilde character at the beginning of a word, followed by all of the characters preceding the first unquoted slash in the word, or all the characters in the word if there is no slash. In an assignment (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 4.21, Variable Assignment), multiple tilde-prefixes can be used: at the beginning of the word (that is, following the equal sign of the assignment), following any unquoted colon, or both. A tilde-prefix in an assignment is terminated by the first unquoted colon or slash. If none of the characters in the tilde-prefix are quoted, the characters in the tildeprefix following the tilde are treated as a possible login name from the user database. A portable login name cannot contain characters outside the set given in the description of the LOGNAME environment variable in the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.3, Other Environment Variables. If the login name is null (that is, the tilde-prefix contains only the tilde), the tilde-prefix is replaced by the value of the variable HOME. If HOME is unset, the results are unspecified. Otherwise, the tilde-prefix shall be replaced by a pathname of the initial working directory associated with the login name obtained using the getpwnam() function as defined in the System Interfaces volume of IEEE Std 1003.1-2001. If the system does not recognize the login name, the results are undefined.

2.6.2 Parameter Expansion

The format for parameter expansion is as follows:

```
${expression}
```

where *expression* consists of all characters until the matching '}'. Any '}' escaped by a backslash or within a quoted string, and characters in embedded arithmetic expansions, command substitutions, and variable expansions, shall not be examined in determining the matching '}'.

The simplest form for parameter expansion is:

```
${parameter}
```

The value, if any, of *parameter* shall be substituted.

The parameter name or symbol can be enclosed in braces, which are optional except for positional parameters with more than one digit or when *parameter* is followed by a character that could be interpreted as part of the name. The matching closing brace shall be determined by

counting brace levels, skipping over enclosed quoted strings, and command substitutions.

If the parameter name or symbol is not enclosed in braces, the expansion shall use the longest valid name (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.230, Name), whether or not the symbol represented by that name exists.

If a parameter expansion occurs inside double-quotes:

- Pathname expansion shall not be performed on the results of the expansion.
- Field splitting shall not be performed on the results of the expansion, with the exception of '@'; see Section 2.5.2 (on page 34).

In addition, a parameter expansion can be modified by using one of the following formats. In each case that a value of *word* is needed (based on the state of *parameter*, as described below), *word* shall be subjected to tilde expansion, parameter expansion, command substitution, and arithmetic expansion. If *word* is not needed, it shall not be expanded. The '}' character that delimits the following parameter expansion modifications shall be determined as described previously in this section and in Section 2.2.3 (on page 30). (For example, \${foo-bar}xyz} would result in the expansion of foo followed by the string xyz} if foo is set, else the string "barxyz}").

\${parameter:-word} Use Default Values. If parameter is unset or null, the expansion of word shall be substituted; otherwise, the value of parameter shall be substituted.

S{parameter:=word} Assign Default Values. If parameter is unset or null, the expansion of word shall be assigned to parameter. In all cases, the final value of parameter shall be substituted. Only variables, not positional parameters or special parameters, can be assigned in this way.

\$\{\text{parameter:?[word]}\}\ Indicate Error if Null or Unset. If parameter is unset or null, the expansion of word (or a message indicating it is unset if word is omitted) shall be written to standard error and the shell exits with a non-zero exit status. Otherwise, the value of parameter shall be substituted. An interactive shell need not exit.

\${parameter:+word} **Use Alternative Value**. If parameter is unset or null, null shall be substituted; otherwise, the expansion of word shall be substituted.

In the parameter expansions shown previously, use of the colon in the format shall result in a test for a parameter that is unset or null; omission of the colon shall result in a test for a parameter that is only unset. The following table summarizes the effect of the colon:

	parameter Set and Not Null	parameter Set But Null	parameter Unset
\${parameter:-word}	substitute parameter	substitute word	substitute word
\${parameter-word}	substitute parameter	substitute null	substitute word
\${parameter:=word}	substitute parameter	assign <i>word</i>	assign word
\${parameter=word}	substitute parameter	substitute null	assign word
\${parameter:?word}	substitute parameter	error, exit	error, exit
\${parameter?word}	substitute parameter	substitute null	error, exit
\${parameter:+word}	substitute word	substitute null	substitute null
\${parameter+word}	substitute word	substitute word	substitute null

In all cases shown with "substitute", the expression is replaced with the value shown. In all cases shown with "assign", *parameter* is assigned that value, which also replaces the expression.

```
1594
              ${#parameter}
                                     String Length. The length in characters of the value of parameter shall be
                                    substituted. If parameter is '*' or '@', the result of the expansion is
1595
                                     unspecified.
1596
              The following four varieties of parameter expansion provide for substring processing. In each
1597
1598
              case, pattern matching notation (see Section 2.13 (on page 62)), rather than regular expression
              notation, shall be used to evaluate the patterns. If parameter is '*' or '@', the result of the
1599
              expansion is unspecified. Enclosing the full parameter expansion string in double-quotes shall
1600
              not cause the following four varieties of pattern characters to be quoted, whereas quoting
1601
              characters within the braces shall have this effect.
1602
1603
              ${parameter%word}
                                     Remove Smallest Suffix Pattern. The word shall be expanded to produce
1604
                                     a pattern. The parameter expansion shall then result in parameter, with the
                                    smallest portion of the suffix matched by the pattern deleted.
1605
                                    Remove Largest Suffix Pattern. The word shall be expanded to produce a
              ${parameter%%word}
1606
1607
                                     pattern. The parameter expansion shall then result in parameter, with the
1608
                                     largest portion of the suffix matched by the pattern deleted.
                                     Remove Smallest Prefix Pattern. The word shall be expanded to produce
1609
              ${parameter#word}
                                     a pattern. The parameter expansion shall then result in parameter, with the
1610
1611
                                     smallest portion of the prefix matched by the pattern deleted.
              ${parameter##word}
                                     Remove Largest Prefix Pattern. The word shall be expanded to produce a
1612
1613
                                     pattern. The parameter expansion shall then result in parameter, with the
1614
                                     largest portion of the prefix matched by the pattern deleted.
1615
              Examples
1616
              ${parameter:-word}
1617
                  In this example, ls is executed only if x is null or unset. (The \$(ls) command substitution
                  notation is explained in Section 2.6.3 (on page 40).)
1618
1619
                      \{x:-\$(ls)\}
1620
              ${parameter:=word}
                  unset X
1621
                  echo ${X:=abc}
1622
                  abc
1623
              ${parameter:?word}
1624
                  unset posix
1625
                  echo ${posix:?}
1626
1627
                  sh: posix: parameter null or not set
              ${parameter:+word}
1628
                  set a b c
1629
                  echo \{3:+posix\}
1630
1631
                  posix
1632
              ${#parameter}
                  HOME=/usr/posix
1633
                  echo ${#HOME}
1634
1635
              ${parameter%word}
1636
                  x=file.c
1637
1638
                  echo \{x\%.c\}.o
```

```
1639
                  file.o
             ${parameter%%word}
1640
1641
                  x=posix/src/std
                  echo \{x\%\%/*\}
1642
1643
                  posix
             ${parameter#word}
1644
                  x=$HOME/src/cmd
1645
                  echo ${x#$HOME}
1646
1647
                  /src/cmd
             ${parameter##word}
1648
                  x=/one/two/three
1649
                  echo ${x##*/}
1650
                  three
1651
             The double-quoting of patterns is different depending on where the double-quotes are placed:
1652
              "${x#*}"
                           The asterisk is a pattern character.
1653
             ${x#"*"}
                           The literal asterisk is quoted and not special.
1654
```

1655 **2.6.3 Command Substitution**

Command substitution allows the output of a command to be substituted in place of the command name itself. Command substitution shall occur when the command is enclosed as follows:

1659 \$ (command)

1656

1658

1660

16611662

1663

1664

1665 1666

1667 1668

1669

1670

1671

1672 1673

1674

1675

1676

16771678

1679

1680

1681 1682 or (backquoted version):

'command'

The shell shall expand the command substitution by executing *command* in a subshell environment (see Section 2.12 (on page 61)) and replacing the command substitution (the text of *command* plus the enclosing "\$()" or backquotes) with the standard output of the command, removing sequences of one or more <newline>s at the end of the substitution. Embedded <newline>s before the end of the output shall not be removed; however, they may be treated as field delimiters and eliminated during field splitting, depending on the value of *IFS* and quoting that is in effect.

With the \$(command) form, all characters following the open parenthesis to the matching closing parenthesis constitute the command. Any valid shell script can be used for command, except a script consisting solely of redirections which produces unspecified results.

The results of command substitution shall not be processed for further tilde expansion, parameter expansion, command substitution, or arithmetic expansion. If a command substitution occurs inside double-quotes, field splitting and pathname expansion shall not be performed on the results of the substitution.

1683 Command substitution can be nested. To specify nesting within the backquoted version, the application shall precede the inner backquotes with backslashes, for example:

```
1685 \ 'command\'
```

1686

1688

1689

1691

1692

1693

1694

1695

1696

1698

1699

1700 1701

1702

1703

1704 1705

1706

1707

1708 1709

1710

1711

If the command substitution consists of a single subshell, such as:

```
1687 $ ( (command) )
```

a conforming application shall separate the "\$(" and '(' into two tokens (that is, separate them with white space). This is required to avoid any ambiguities with arithmetic expansion.

1690 2.6.4 Arithmetic Expansion

Arithmetic expansion provides a mechanism for evaluating an arithmetic expression and substituting its value. The format for arithmetic expansion shall be as follows:

```
$((expression))
```

The expression shall be treated as if it were in double-quotes, except that a double-quote inside the expression is not treated specially. The shell shall expand all tokens in the expression for parameter expansion, command substitution, and quote removal.

Next, the shell shall treat this as an arithmetic expression and substitute the value of the expression. The arithmetic expression shall be processed according to the rules given in Section 1.7.2.1 (on page 7), with the following exceptions:

- · Only signed long integer arithmetic is required.
- Only the decimal-constant, octal-constant, and hexadecimal-constant constants specified in the ISO C standard, Section 6.4.4.1 are required to be recognized as constants.
- The sizeof() operator and the prefix and postfix "++" and "--" operators are not required.
- Selection, iteration, and jump statements are not supported.

As an extension, the shell may recognize arithmetic expressions beyond those listed. The shell may use a signed integer type with a rank larger than the rank of **signed long**. The shell may use a real-floating type instead of **signed long** as long as it does not affect the results in cases where there is no overflow. If the expression is invalid, the expansion fails and the shell shall write a message to standard error indicating the failure.

Examples

A simple example using arithmetic expansion:

1722 1723

1724

1725

1726

17281729

1731

1732 1733

1734

1735

1736

1737

17381739

1740

1719 2.6.5 Field Splitting

After parameter expansion (Section 2.6.2 (on page 37)), command substitution (Section 2.6.3 (on page 40)), and arithmetic expansion (Section 2.6.4 (on page 41)), the shell shall scan the results of expansions and substitutions that did not occur in double-quotes for field splitting and multiple fields can result.

The shell shall treat each character of the *IFS* as a delimiter and use the delimiters to split the results of parameter expansion and command substitution into fields.

1. If the value of *IFS* is a <space>, <tab>, and <newline>, or if it is unset, any sequence of <space>s, <tab>s, or <newline>s at the beginning or end of the input shall be ignored and any sequence of those characters within the input shall delimit a field. For example, the input:

<newline><space><tab>foo<tab><tab>bar<space>

yields two fields, foo and bar.

- 2. If the value of *IFS* is null, no field splitting shall be performed.
- 3. Otherwise, the following rules shall be applied in sequence. The term "*IFS* white space" is used to mean any sequence (zero or more instances) of white space characters that are in the *IFS* value (for example, if *IFS* contains <space>/<comma>/<tab>, any sequence of <space>s and <tab>s is considered *IFS* white space).
 - a. *IFS* white space shall be ignored at the beginning and end of the input.
 - b. Each occurrence in the input of an *IFS* character that is not *IFS* white space, along with any adjacent *IFS* white space, shall delimit a field, as described previously.
 - c. Non-zero-length *IFS* white space shall delimit a field.

1741 **2.6.6 Pathname Expansion**

After field splitting, if *set* –**f** is not in effect, each field in the resulting command line shall be expanded using the algorithm described in Section 2.13 (on page 62), qualified by the rules in Section 2.13.3 (on page 63).

1745 **2.6.7 Quote Removal**

The quote characters: '\', ''', and '"' (backslash, single-quote, double-quote) that were present in the original word shall be removed unless they have themselves been quoted.

1748 2.7 Redirection

Redirection is used to open and close files for the current shell execution environment (see Section 2.12 (on page 61)) or for any command. Redirection operators can be used with numbers representing file descriptors (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.165, File Descriptor) as described below.

The overall format used for redirection is:

```
[n]redir-op word
```

The number n is an optional decimal number designating the file descriptor number; the application shall ensure it is delimited from any preceding text and immediately precede the redirection operator redir-op. If n is quoted, the number shall not be recognized as part of the redirection expression. For example:

```
echo \2>a
```

writes the character 2 into file **a**. If any part of *redir-op* is quoted, no redirection expression is recognized. For example:

```
echo 2\>a
```

writes the characters 2>a to standard output. The optional number, redirection operator, and word shall not appear in the arguments provided to the command to be executed (if any).

Open files are represented by decimal numbers starting with zero. The largest possible value is implementation-defined; however, all implementations shall support at least 0 to 9, inclusive, for use by the application. These numbers are called "file descriptors". The values 0, 1, and 2 have special meaning and conventional uses and are implied by certain redirection operations; they are referred to as *standard input*, *standard output*, and *standard error*, respectively. Programs usually take their input from standard input, and write output on standard output. Error messages are usually written on standard error. The redirection operators can be preceded by one or more digits (with no intervening

sllowed) to designate the file descriptor number.

If the redirection operator is "<<" or "<<-", the word that follows the redirection operator shall be subjected to quote removal; it is unspecified whether any of the other expansions occur. For the other redirection operators, the word that follows the redirection operator shall be subjected to tilde expansion, parameter expansion, command substitution, arithmetic expansion, and quote removal. Pathname expansion shall not be performed on the word by a non-interactive shell; an interactive shell may perform it, but shall do so only when the expansion would result in one word.

If more than one redirection operator is specified with a command, the order of evaluation is from beginning to end.

A failure to open or create a file shall cause a redirection to fail.

1784 2.7.1 Redirecting Input

Input redirection shall cause the file whose name results from the expansion of *word* to be opened for reading on the designated file descriptor, or standard input if the file descriptor is not specified.

1788 The general format for redirecting input is:

1789 [n] < word

1790 1791

1796

1797

1798

1799

1800

1801

1802 1803

1805

1806

1807

1808

1809

1811

1812

1816

1817

1818

1822 1823 where the optional *n* represents the file descriptor number. If the number is omitted, the redirection shall refer to standard input (file descriptor 0).

1792 2.7.2 Redirecting Output

1793 The two general formats for redirecting output are:

1794 [n] > word 1795 [n] > | word

where the optional n represents the file descriptor number. If the number is omitted, the redirection shall refer to standard output (file descriptor 1).

Output redirection using the '>' format shall fail if the *noclobber* option is set (see the description of *set* –**C**) and the file named by the expansion of *word* exists and is a regular file. Otherwise, redirection using the '>' or ">| " formats shall cause the file whose name results from the expansion of *word* to be created and opened for output on the designated file descriptor, or standard output if none is specified. If the file does not exist, it shall be created; otherwise, it shall be truncated to be an empty file after being opened.

1804 2.7.3 Appending Redirected Output

Appended output redirection shall cause the file whose name results from the expansion of word to be opened for output on the designated file descriptor. The file is opened as if the <code>open()</code> function as defined in the System Interfaces volume of IEEE Std 1003.1-2001 was called with the O_APPEND flag. If the file does not exist, it shall be created.

The general format for appending redirected output is as follows:

1810 [n] >> word

where the optional *n* represents the file descriptor number. If the number is omitted, the redirection refers to standard output (file descriptor 1).

1813 2.7.4 Here-Document

The redirection operators "<<" and "<<-" both allow redirection of lines contained in a shell input file, known as a "here-document", to the input of a command.

The here-document shall be treated as a single word that begins after the next <newline> and continues until there is a line containing only the delimiter and a <newline>, with no <blank>s in between. Then the next here-document starts, if there is one. The format is as follows:

 1819
 [n] << word</td>

 1820
 here-document

 1821
 delimiter

where the optional *n* represents the file descriptor number. If the number is omitted, the here-document refers to standard input (file descriptor 0).

If any character in *word* is quoted, the delimiter shall be formed by performing quote removal on *word*, and the here-document lines shall not be expanded. Otherwise, the delimiter shall be the *word* itself.

If no characters in *word* are quoted, all lines of the here-document shall be expanded for parameter expansion, command substitution, and arithmetic expansion. In this case, the backslash in the input behaves as the backslash inside double-quotes (see Section 2.2.3 (on page 30)). However, the double-quote character ('"') shall not be treated specially within a here-document, except when the double-quote appears within "\$()", ""\", or "\${}".

If the redirection symbol is "<<-", all leading <tab>s shall be stripped from input lines and the line containing the trailing delimiter. If more than one "<<" or "<<-" operator is specified on a line, the here-document associated with the first operator shall be supplied first by the application and shall be read first by the shell.

Examples

An example of a here-document follows:

```
1838 cat <<eof1; cat <<eof2
1839 Hi,
1840 eof1
1841 Helene.
1842 eof2
```

1843 2.7.5 Duplicating an Input File Descriptor

The redirection operator:

```
1845 [n] < \&word
```

shall duplicate one input file descriptor from another, or shall close one. If *word* evaluates to one or more digits, the file descriptor denoted by n, or standard input if n is not specified, shall be made to be a copy of the file descriptor denoted by *word*; if the digits in *word* do not represent a file descriptor already open for input, a redirection error shall result; see Section 2.8.1 (on page 46). If *word* evaluates to '-', file descriptor n, or standard input if n is not specified, shall be closed. Attempts to close a file descriptor that is not open shall not constitute an error. If *word* evaluates to something else, the behavior is unspecified.

1853 2.7.6 Duplicating an Output File Descriptor

The redirection operator:

```
1855 [n] >&word
```

shall duplicate one output file descriptor from another, or shall close one. If *word* evaluates to one or more digits, the file descriptor denoted by n, or standard output if n is not specified, shall be made to be a copy of the file descriptor denoted by *word*; if the digits in *word* do not represent a file descriptor already open for output, a redirection error shall result; see Section 2.8.1 (on page 46). If *word* evaluates to '-', file descriptor n, or standard output if n is not specified, is closed. Attempts to close a file descriptor that is not open shall not constitute an error. If *word* evaluates to something else, the behavior is unspecified.

1863 2.7.7 Open File Descriptors for Reading and Writing

The redirection operator:

[n] <> word

shall cause the file whose name is the expansion of *word* to be opened for both reading and writing on the file descriptor denoted by *n*, or standard input if *n* is not specified. If the file does not exist, it shall be created.

2.8 Exit Status and Errors

2.8.1 Consequences of Shell Errors

For a non-interactive shell, an error condition encountered by a special built-in (see Section 2.14 (on page 64)) or other type of utility shall cause the shell to write a diagnostic message to standard error and exit as shown in the following table:

Error	Special Built-In	Other Utilities
Shell language syntax error	Shall exit	Shall exit
Utility syntax error (option or operand error)	Shall exit	Shall not exit
Redirection error	Shall exit	Shall not exit
Variable assignment error	Shall exit	Shall not exit
Expansion error	Shall exit	Shall exit
Command not found	N/A	May exit
Dot script not found	Shall exit	N/A

An expansion error is one that occurs when the shell expansions defined in Section 2.6 (on page 36) are carried out (for example, " $\{x!y\}$ ", because '!' is not a valid operator); an implementation may treat these as syntax errors if it is able to detect them during tokenization, rather than during expansion.

If any of the errors shown as "shall exit" or "(may) exit" occur in a subshell, the subshell shall (respectively may) exit with a non-zero status, but the script containing the subshell shall not exit because of the error.

In all of the cases shown in the table, an interactive shell shall write a diagnostic message to standard error without exiting.

2.8.2 Exit Status for Commands

Each command has an exit status that can influence the behavior of other shell commands. The exit status of commands that are not utilities is documented in this section. The exit status of the standard utilities is documented in their respective sections.

If a command is not found, the exit status shall be 127. If the command name is found, but it is not an executable utility, the exit status shall be 126. Applications that invoke utilities without using the shell should use these exit status values to report similar errors.

If a command fails during word expansion or redirection, its exit status shall be greater than zero.

Internally, for purposes of deciding whether a command exits with a non-zero exit status, the shell shall recognize the entire status value retrieved for the command by the equivalent of the *wait()* function WEXITSTATUS macro (as defined in the System Interfaces volume of IEEE Std 1003.1-2001). When reporting the exit status with the special parameter '?', the shell

 shall report the full eight bits of exit status available. The exit status of a command that terminated because it received a signal shall be reported as greater than 128.

1906 2.9 Shell Commands

This section describes the basic structure of shell commands. The following command descriptions each describe a format of the command that is only used to aid the reader in recognizing the command type, and does not formally represent the syntax. Each description discusses the semantics of the command; for a formal definition of the command language, consult Section 2.10 (on page 55).

A *command* is one of the following:

- Simple command (see Section 2.9.1)
- Pipeline (see Section 2.9.2 (on page 49))
- List compound-list (see Section 2.9.3 (on page 50))
- Compound command (see Section 2.9.4 (on page 52))
- Function definition (see Section 2.9.5 (on page 54))

Unless otherwise stated, the exit status of a command shall be that of the last simple command executed by the command. There shall be no limit on the size of any shell command other than that imposed by the underlying system (memory constraints, {ARG_MAX}, and so on).

2.9.1 Simple Commands

A "simple command" is a sequence of optional variable assignments and redirections, in any sequence, optionally followed by words and redirections, terminated by a control operator.

When a given simple command is required to be executed (that is, when any conditional construct such as an AND-OR list or a **case** statement has not bypassed the simple command), the following expansions, assignments, and redirections shall all be performed from the beginning of the command text to the end:

- 1. The words that are recognized as variable assignments or redirections according to Section 2.10.2 (on page 56) are saved for processing in steps 3 and 4.
- 2. The words that are not variable assignments or redirections shall be expanded. If any fields remain following their expansion, the first field shall be considered the command name and remaining fields are the arguments for the command.
- 3. Redirections shall be performed as described in Section 2.7 (on page 43).
- 4. Each variable assignment shall be expanded for tilde expansion, parameter expansion, command substitution, arithmetic expansion, and quote removal prior to assigning the value.

In the preceding list, the order of steps 3 and 4 may be reversed for the processing of special built-in utilities; see Section 2.14 (on page 64).

If no command name results, variable assignments shall affect the current execution environment. Otherwise, the variable assignments shall be exported for the execution environment of the command and shall not affect the current execution environment (except for special built-ins). If any of the variable assignments attempt to assign a value to a read-only variable, a variable assignment error shall occur. See Section 2.8.1 (on page 46) for the consequences of these errors.

If there is no command name, any redirections shall be performed in a subshell environment; it is unspecified whether this subshell environment is the same one as that used for a command substitution within the command. (To affect the current execution environment, see the *exec* special built-in.) If any of the redirections performed in the current shell execution environment fail, the command shall immediately fail with an exit status greater than zero, and the shell shall write an error message indicating the failure. See Section 2.8.1 (on page 46) for the consequences of these failures on interactive and non-interactive shells.

If there is a command name, execution shall continue as described in Section 2.9.1.1. If there is no command name, but the command contained a command substitution, the command shall complete with the exit status of the last command substitution performed. Otherwise, the command shall complete with a zero exit status.

1956 2.9.1.1 Command Search and Execution

If a simple command results in a command name and an optional list of arguments, the following actions shall be performed:

- 1. If the command name does not contain any slashes, the first successful step in the following sequence shall occur:
 - a. If the command name matches the name of a special built-in utility, that special built-in utility shall be invoked.
 - b. If the command name matches the name of a function known to this shell, the function shall be invoked as described in Section 2.9.5 (on page 54). If the implementation has provided a standard utility in the form of a function, it shall not be recognized at this point. It shall be invoked in conjunction with the path search in step 1d.
 - c. If the command name matches the name of a utility listed in the following table, that utility shall be invoked.

alias	false	jobs	read	wait
bg	fc	kill	true	
cd	fg	newgrp	umask	
command	getopts	pwd	unalias	

- d. Otherwise, the command shall be searched for using the *PATH* environment variable as described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables:
 - i. If the search is successful:
 - a. If the system has implemented the utility as a regular built-in or as a shell function, it shall be invoked at this point in the path search.
 - b. Otherwise, the shell executes the utility in a separate utility environment (see Section 2.12 (on page 61)) with actions equivalent to calling the *execve*() function as defined in the System Interfaces volume of IEEE Std 1003.1-2001 with the *path* argument set to the pathname resulting from the search, *arg*0 set to the command name, and the remaining arguments set to the operands, if any.

If the *execve()* function fails due to an error equivalent to the [ENOEXEC] error defined in the System Interfaces volume of IEEE Std 1003.1-2001, the shell shall execute a command equivalent to having a shell invoked with the command name as its first operand, with any remaining arguments

passed to the new shell. If the executable file is not a text file, the shell may bypass this command execution. In this case, it shall write an error message, and shall return an exit status of 126.

Once a utility has been searched for and found (either as a result of this specific search or as part of an unspecified shell start-up activity), an implementation may remember its location and need not search for the utility again unless the *PATH* variable has been the subject of an assignment. If the remembered location fails for a subsequent invocation, the shell shall repeat the search to find the new location for the utility, if any.

- ii. If the search is unsuccessful, the command shall fail with an exit status of 127 and the shell shall write an error message.
- 2. If the command name contains at least one slash, the shell shall execute the utility in a separate utility environment with actions equivalent to calling the *execve()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001 with the *path* and *arg0* arguments set to the command name, and the remaining arguments set to the operands, if any.

If the <code>execve()</code> function fails due to an error equivalent to the <code>[ENOEXEC]</code> error, the shell shall execute a command equivalent to having a shell invoked with the command name as its first operand, with any remaining arguments passed to the new shell. If the executable file is not a text file, the shell may bypass this command execution. In this case, it shall write an error message and shall return an exit status of 126.

2.9.2 Pipelines

A *pipeline* is a sequence of one or more commands separated by the control operator $' \mid '$. The standard output of all but the last command shall be connected to the standard input of the next command.

The format for a pipeline is:

```
[!] command1 [ | command2 ...]
```

The standard output of *command1* shall be connected to the standard input of *command2*. The standard input, standard output, or both of a command shall be considered to be assigned by the pipeline before any redirection specified by redirection operators that are part of the command (see Section 2.7 (on page 43)).

If the pipeline is not in the background (see Section 2.9.3.1 (on page 50)), the shell shall wait for the last command specified in the pipeline to complete, and may also wait for all commands to complete.

Exit Status

If the reserved word! does not precede the pipeline, the exit status shall be the exit status of the last command specified in the pipeline. Otherwise, the exit status shall be the logical NOT of the exit status of the last command. That is, if the last command returns zero, the exit status shall be 1; if the last command returns greater than zero, the exit status shall be zero.

2029 2.9.3 Lists

2032

2033

2034

20352036

2037

2038

2039

2040

2041

2042

2043

2044

2058

2059

20602061

2062

2063 2064

2065

2066 2067

2068

2069 2070

2030 An AND-OR list is a sequence of one or more pipelines separated by the operators "&&" and "| | ".

A *list* is a sequence of one or more AND-OR lists separated by the operators '; ' and '&' and optionally terminated by '; ', '&', or <newline>.

The operators "&&" and "||" shall have equal precedence and shall be evaluated with left associativity. For example, both of the following commands write solely **bar** to standard output:

```
false && echo foo || echo bar true || echo foo && echo bar
```

A ';' or <newline> terminator shall cause the preceding AND-OR list to be executed sequentially; an '&' shall cause asynchronous execution of the preceding AND-OR list.

The term "compound-list" is derived from the grammar in Section 2.10 (on page 55); it is equivalent to a sequence of *lists*, separated by <newline>s, that can be preceded or followed by an arbitrary number of <newline>s.

Examples

The following is an example that illustrates <newline>s in compound-lists:

```
while
2045
2046
                    # a couple of <newline>s
2047
                    # a list
                    date && who || ls; cat file
2048
                    # a couple of <newline>s
2049
2050
                    # another list
                    wc file > output & true
2051
2052
               do
                    # 2 lists
2053
2054
                    ls
2055
                    cat file
2056
               done
```

2057 2.9.3.1 Asynchronous Lists

If a command is terminated by the control operator ampersand ('&'), the shell shall execute the command asynchronously in a subshell. This means that the shell shall not wait for the command to finish before executing the next command.

The format for running a command in the background is:

```
command1 & [command2 & ... ]
```

The standard input for an asynchronous list, before any explicit redirections are performed, shall be considered to be assigned to a file that has the same properties as /dev/null. If it is an interactive shell, this need not happen. In all cases, explicit redirection of standard input shall override this activity.

When an element of an asynchronous list (the portion of the list ended by an ampersand, such as *command1*, above) is started by the shell, the process ID of the last command in the asynchronous list element shall become known in the current shell execution environment; see Section 2.12 (on page 61). This process ID shall remain known until:

- 2071 1. The command terminates and the application waits for the process ID.
- 2072 2. Another asynchronous list invoked before "\$!" (corresponding to the previous asynchronous list) is expanded in the current execution environment.

The implementation need not retain more than the {CHILD_MAX} most recent entries in its list of known process IDs in the current shell execution environment.

Exit Status

The exit status of an asynchronous list shall be zero.

2078 2.9.3.2 Sequential Lists

2076

2079 Commands that are separated by a semicolon (';') shall be executed sequentially.

2080 The format for executing commands sequentially shall be:

```
2081 command1 [; command2] ...
```

Each command shall be expanded and executed in the order specified.

2083 Exit Status

The exit status of a sequential list shall be the exit status of the last command in the list.

2085 2.9.3.3 AND Lists

2086 The control operator "&&" denotes an AND list. The format shall be:

```
2087 command1 [ && command2] ...
```

First *command1* shall be executed. If its exit status is zero, *command2* shall be executed, and so on, until a command has a non-zero exit status or there are no more commands left to execute. The commands are expanded only if they are executed.

2091 Exit Status

The exit status of an AND list shall be the exit status of the last command that is executed in the list.

2094 2.9.3.4 OR Lists

2095

2099

The control operator " | | " denotes an OR List. The format shall be:

```
2096 command1 [ | command2] ...
```

First, *command1* shall be executed. If its exit status is non-zero, *command2* shall be executed, and so on, until a command has a zero exit status or there are no more commands left to execute.

Exit Status

The exit status of an OR list shall be the exit status of the last command that is executed in the list.

2105

2106

2107 2108

2116

2117

2121

2122

2123

2129

2130

2131

2135

2102 2.9.4 Compound Commands

The shell has several programming constructs that are "compound commands", which provide control flow for commands. Each of these compound commands has a reserved word or control operator at the beginning, and a corresponding terminator reserved word or operator at the end. In addition, each can be followed by redirections on the same line as the terminator. Each redirection shall apply to all the commands within the compound command that do not explicitly override that redirection.

2109 2.9.4.1 Grouping Commands

2110 The format for grouping commands is as follows:

2111	(compound-list)	Execute compound-list in a subshell environment; see Section 2.12 (on page		
2112		61). Variable assignments and built-in commands that affect the		
2113		environment shall not remain in effect after the list finishes.		
2114	{ compound-list;}	Execute <i>compound-list</i> in the current process environment. The semicolon		
2115		shown here is an example of a control operator delimiting the } reserved		

55); a <newline> is frequently used.

word. Other delimiters are possible, as shown in Section 2.10 (on page

2118 Exit Status

The exit status of a grouping command shall be the exit status of *compound-list*.

2120 2.9.4.2 The for Loop

The **for** loop shall execute a sequence of commands for each member in a list of *items*. The **for** loop requires that the reserved words **do** and **done** be used to delimit the sequence of commands.

The format for the **for** loop is as follows:

```
2125 for name [ in [word ... ]]
2126 do
2127 compound-list
2128 done
```

First, the list of words following **in** shall be expanded to generate a list of items. Then, the variable *name* shall be set to each item, in turn, and the *compound-list* executed each time. If no items result from the expansion, the *compound-list* shall not be executed. Omitting:

```
2132 in word...
```

shall be equivalent to:

```
2134 in "$@"
```

Exit Status

The exit status of a **for** command shall be the exit status of the last command that executes. If there are no items, the exit status shall be zero.

2138 2.9.4.3 Case Conditional Construct

The conditional construct **case** shall execute the *compound-list* corresponding to the first one of several *patterns* (see Section 2.13 (on page 62)) that is matched by the string resulting from the tilde expansion, parameter expansion, command substitution, arithmetic expansion, and quote removal of the given word. The reserved word **in** shall denote the beginning of the patterns to be matched. Multiple patterns with the same *compound-list* shall be delimited by the '|' symbol. The control operator ')' terminates a list of patterns corresponding to a given action. The *compound-list* for each list of patterns, with the possible exception of the last, shall be terminated with ";;". The **case** construct terminates with the reserved word **esac** (**case** reversed).

The format for the **case** construct is as follows:

The ";; " is optional for the last *compound-list*.

In order from the beginning to the end of the **case** statement, each *pattern* that labels a *compound-list* shall be subjected to tilde expansion, parameter expansion, command substitution, and arithmetic expansion, and the result of these expansions shall be compared against the expansion of *word*, according to the rules described in Section 2.13 (on page 62) (which also describes the effect of quoting parts of the pattern). After the first match, no more patterns shall be expanded, and the *compound-list* shall be executed. The order of expansion and comparison of multiple *patterns* that label a *compound-list* statement is unspecified.

Exit Status

The exit status of **case** shall be zero if no patterns are matched. Otherwise, the exit status shall be the exit status of the last command executed in the *compound-list*.

2164 2.9.4.4 The if Conditional Construct

The **if** command shall execute a *compound-list* and use its exit status to determine whether to execute another *compound-list*.

The format for the **if** construct is as follows:

```
2168
                if compound-list
2169
                then
2170
                     compound-list
                [elif compound-list
2171
2172
2173
                     compound-list] ...
2174
                [else
                     compound-list]
2175
2176
```

The **if** *compound-list* shall be executed; if its exit status is zero, the **then** *compound-list* shall be executed and the command shall complete. Otherwise, each **elif** *compound-list* shall be executed, in turn, and if its exit status is zero, the **then** *compound-list* shall be executed and the command shall complete. Otherwise, the **else** *compound-list* shall be executed.

2181 Exit Status

The exit status of the **if** command shall be the exit status of the **then** or **else** *compound-list* that was executed, or zero, if none was executed.

2184 2.9.4.5 The while Loop

2187

2205

2206

2207

2214

2216

2217

2218

2219

The **while** loop shall continuously execute one *compound-list* as long as another *compound-list* has a zero exit status.

The format of the **while** loop is as follows:

```
      2188
      while compound-list-1

      2189
      do

      2190
      compound-list-2

      2191
      done
```

The *compound-list-1* shall be executed, and if it has a non-zero exit status, the **while** command shall complete. Otherwise, the *compound-list-2* shall be executed, and the process shall repeat.

2194 Exit Status

The exit status of the **while** loop shall be the exit status of the last *compound-list-2* executed, or zero if none was executed.

2197 2.9.4.6 The until Loop

The **until** loop shall continuously execute one *compound-list* as long as another *compound-list* has a non-zero exit status.

2200 The format of the **until** loop is as follows:

```
      2201
      until compound-list-1

      2202
      do

      2203
      compound-list-2

      2204
      done
```

The *compound-list-1* shall be executed, and if it has a zero exit status, the **until** command completes. Otherwise, the *compound-list-2* shall be executed, and the process repeats.

Exit Status

The exit status of the **until** loop shall be the exit status of the last *compound-list-2* executed, or zero if none was executed.

2210 2.9.5 Function Definition Command

A function is a user-defined name that is used as a simple command to call a compound command with new positional parameters. A function is defined with a "function definition command".

The format of a function definition command is as follows:

```
fname() compound-command[io-redirect ...]
```

The function is named *fname*; the application shall ensure that it is a name (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.230, Name). An implementation may allow other characters in a function name as an extension. The implementation shall maintain separate name spaces for functions and variables.

The argument *compound-command* represents a compound command, as described in Section 2.9.4 (on page 52).

When the function is declared, none of the expansions in Section 2.6 (on page 36) shall be performed on the text in *compound-command* or *io-redirect*; all expansions shall be performed as normal each time the function is called. Similarly, the optional *io-redirect* redirections and any variable assignments within *compound-command* shall be performed during the execution of the function itself, not the function definition. See Section 2.8.1 (on page 46) for the consequences of failures of these operations on interactive and non-interactive shells.

When a function is executed, it shall have the syntax-error and variable-assignment properties described for special built-in utilities in the enumerated list at the beginning of Section 2.14 (on page 64).

The *compound-command* shall be executed whenever the function name is specified as the name of a simple command (see Section 2.9.1.1 (on page 48)). The operands to the command temporarily shall become the positional parameters during the execution of the *compound-command*; the special parameter '#' also shall be changed to reflect the number of operands. The special parameter 0 shall be unchanged. When the function completes, the values of the positional parameters and the special parameter '#' shall be restored to the values they had before the function was executed. If the special built-in *return* is executed in the *compound-command*, the function completes and execution shall resume with the next command after the function call.

Exit Status

The exit status of a function definition shall be zero if the function was declared successfully; otherwise, it shall be greater than zero. The exit status of a function invocation shall be the exit status of the last command executed by the function.

2.10 Shell Grammar

The following grammar defines the Shell Command Language. This formal syntax shall take precedence over the preceding text syntax description.

2247 2.10.1 Shell Grammar Lexical Conventions

The input language to the shell must be first recognized at the character level. The resulting tokens shall be classified by their immediate context according to the following rules (applied in order). These rules shall be used to determine what a "token" is that is subject to parsing at the token level. The rules for token recognition in Section 2.3 (on page 31) shall apply.

- 1. A <newline> shall be returned as the token identifier **NEWLINE**.
- 2. If the token is an operator, the token identifier for that operator shall result.
- 3. If the string consists solely of digits and the delimiter character is one of '<' or '>', the token identifier **IO NUMBER** shall be returned.
- 4. Otherwise, the token identifier **TOKEN** results.

Further distinction on **TOKEN** is context-dependent. It may be that the same **TOKEN** yields **WORD**, a **NAME**, an **ASSIGNMENT**, or one of the reserved words below, dependent upon the context. Some of the productions in the grammar below are annotated with a rule number from the following list. When a **TOKEN** is seen where one of those annotated productions could be used to reduce the symbol, the applicable rule shall be applied to convert the token identifier

2267

2268

2269

2270

2271

2272

2273 2274

2275

2276

2277 2278

2279

2280

2281 2282

2283

2284 2285

2286

2287

2288 2289

2290

2291

2292

2293

2294

2295

2296

2297

2298

2299

2300

2301

2302

type of the **TOKEN** to a token identifier acceptable at that point in the grammar. The reduction shall then proceed based upon the token identifier type yielded by the rule applied. When more than one rule applies, the highest numbered rule shall apply (which in turn may refer to another rule). (Note that except in rule 7, the presence of an '=' in the token has no effect.)

The **WORD** tokens shall have the word expansion rules applied to them immediately before the associated command is executed, not at the time the command is parsed.

2.10.2 Shell Grammar Rules

1. [Command Name]

When the **TOKEN** is exactly a reserved word, the token identifier for that reserved word shall result. Otherwise, the token **WORD** shall be returned. Also, if the parser is in any state where only a reserved word could be the next correct token, proceed as above.

Note:

Because at this point quote marks are retained in the token, quoted strings cannot be recognized as reserved words. This rule also implies that reserved words are not recognized except in certain positions in the input, such as after a <newline> or semicolon; the grammar presumes that if the reserved word is intended, it is properly delimited by the user, and does not attempt to reflect that requirement directly. Also note that line joining is done before tokenization, as described in Section 2.2.1 (on page 30), so escaped <newline>s are already removed at this point.

Rule 1 is not directly referenced in the grammar, but is referred to by other rules, or applies globally.

[Redirection to or from filename]

The expansions specified in Section 2.7 (on page 43) shall occur. As specified there, exactly one field can result (or the result is unspecified), and there are additional requirements on pathname expansion.

[Redirection from here-document]

Quote removal shall be applied to the word to determine the delimiter that is used to find the end of the here-document that begins after the next <newline>.

4. [Case statement termination]

When the **TOKEN** is exactly the reserved word **esac**, the token identifier for **esac** shall result. Otherwise, the token **WORD** shall be returned.

5. [NAME in for]

When the **TOKEN** meets the requirements for a name (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.230, Name), the token identifier **NAME** shall result. Otherwise, the token **WORD** shall be returned.

6. [Third word of **for** and **case**]

a. [case only]

When the **TOKEN** is exactly the reserved word **in**, the token identifier for **in** shall result. Otherwise, the token **WORD** shall be returned.

b. **[for** only]

When the **TOKEN** is exactly the reserved word **in** or **do**, the token identifier for **in** or **do** shall result, respectively. Otherwise, the token **WORD** shall be returned.

2304

2305

2306

2307

2308

23092310

2311

2312

2313

2314 2315

2316

23172318

2319

2320

2321

2322

2323

2324

2325

2326

2327 2328 (For a. and b.: As indicated in the grammar, a *linebreak* precedes the tokens **in** and **do**. If <newline>s are present at the indicated location, it is the token after them that is treated in this fashion.)

- [Assignment preceding command name]
 - a. [When the first word]

If the **TOKEN** does not contain the character '=', rule 1 is applied. Otherwise, 7b shall be applied.

b. [Not the first word]

If the **TOKEN** contains the equal sign character:

- If it begins with ' = ', the token **WORD** shall be returned.
- If all the characters preceding '=' form a valid name (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.230, Name), the token ASSIGNMENT_WORD shall be returned. (Quoted characters cannot participate in forming a valid name.)
- Otherwise, it is unspecified whether it is ASSIGNMENT_WORD or WORD that is returned.

Assignment to the **NAME** shall occur as specified in Section 2.9.1 (on page 47).

8. [NAME in function]

When the **TOKEN** is exactly a reserved word, the token identifier for that reserved word shall result. Otherwise, when the **TOKEN** meets the requirements for a name, the token identifier **NAME** shall result. Otherwise, rule 7 applies.

9. [Body of function]

Word expansion and assignment shall never occur, even when required by the rules above, when this rule is being parsed. Each **TOKEN** that might either be expanded or have assignment applied to it shall instead be returned as a single **WORD** consisting only of characters that are exactly the token described in Section 2.3 (on page 31).

```
-----
2329
2330
            The grammar symbols
            */
2331
2332
         %token
                WORD
2333
         %token ASSIGNMENT WORD
2334
         %token NAME
         %token NEWLINE
2335
         %token IO NUMBER
2336
         /* The following are the operators mentioned above. */
2337
         %token
                AND IF
                                 DSEMI
                         OR IF
2338
                '&&'
                         ' | | '
                                        */
2339
                                 ';;'
         %token
                DLESS
                      DGREAT LESSAND GREATAND
                                              LESSGREAT DLESSDASH
2340
                ' << '
                      ' >> '
                             '<&'
                                     '>&'
                                              ' <> '
                                                        ′ <<--′
2341
2342
         %token CLOBBER
                ' > | '
2343
         /*
```

```
2344
           /* The following are the reserved words. */
           %token If
                                         Elif
2345
                         Then
                                 Else
                                                  Γi
                                                        Do
                                                              Done
2346
                   'if'
                         'then'
                                 'else' 'elif' 'fi' 'do'
                                                              'done'
2347
           %token Case
                           Esac
                                   While
                                            Until
                                            'until'
                                                      'for'
2348
           /*
                   'case'
                           'esac'
                                   'while'
                                                              * /
2349
           /* These are reserved words, not operator tokens, and are
2350
              recognized when reserved words are recognized. */
2351
           %token Lbrace
                             Rbrace
                                       Bang
2352
                   ′ { ′
                             ' } '
                                       111
2353
           %token In
                   'in'
2354
           /*
           /* -----
2355
2356
              The Grammar
              */
2357
2358
           %start complete_command
2359
2360
           complete command : list separator
2361
                              list
2362
2363
           list
                              list separator op and or
2364
                                                 and or
2365
                                                       pipeline
2366
           and or
2367
                              and_or AND_IF linebreak pipeline
2368
                              and or OR IF linebreak pipeline
2369
2370
           pipeline
                                   pipe sequence
2371
                              Bang pipe_sequence
2372
2373
           pipe sequence
                                                           command
2374
                              pipe_sequence '|' linebreak command
2375
                              simple command
2376
           command
2377
                              compound command
2378
                              compound_command redirect_list
                              function definition
2379
2380
2381
           compound command : brace group
                              subshell
2382
2383
                              for clause
                              case clause
2384
                              if clause
2385
                              while clause
2386
                              until clause
2387
2388
                              '(' compound list ')'
2389
           subshell
2390
2391
           compound_list
                                           term
2392
                            newline list term
```

```
2393
                                              term separator
2394
                              | newline_list term separator
2395
2396
                              : term separator and or
           term
2397
2398
            for clause
                               For name linebreak
2399
                                                                                  do group
2400
                                For name linebreak in
                                                                  sequential sep do group
                              | For name linebreak in wordlist sequential_sep do_group
2401
2402
2403
           name
                              : NAME
                                                           /* Apply rule 5 */
2404
                                                            /* Apply rule 6 */
2405
                              : In
2406
                              : wordlist WORD
2407
           wordlist
2408
                                          WORD
2409
                              : Case WORD linebreak in linebreak case_list
2410
           case clause
2411
                                Case WORD linebreak in linebreak case list ns Esac
                                Case WORD linebreak in linebreak
2412
                                                                                  Esac
2413
           case list ns
2414
                              : case_list case_item_ns
2415
                                           case item ns
2416
                               case list case item
2417
           case list
2418
                                           case item
2419
                                    pattern ')'
2420
                                                                 linebreak
           case item ns
2421
                                    pattern ')' compound list linebreak
                                '(' pattern ')'
2422
2423
                                '(' pattern ')' compound list linebreak
2424
2425
           case item
                                    pattern ')' linebreak
                                                                 DSEMI linebreak
                                    pattern ')' compound list DSEMI linebreak
2426
                                '(' pattern ')' linebreak
2427
                                                               DSEMI linebreak
2428
                                '(' pattern ')' compound list DSEMI linebreak
2429
2430
                                             WORD
                                                           /* Apply rule 4 */
           pattern
                                pattern '|' WORD
2431
                                                           /* Do not apply rule 4 */
2432
2433
           if clause
                              : If compound_list Then compound_list else_part Fi
                              If compound list Then compound list
2434
2435
2436
           else part
                              : Elif compound list Then else part
2437
                                Else compound list
2438
2439
           while clause
                              : While compound list do group
2440
2441
           until clause
                              : Until compound list do group
2442
2443
           function_definition : fname '(' ')' linebreak function_body
2444
```

```
2445
            function body
                               : compound command
                                                                      /* Apply rule 9 */
2446
                                 compound_command redirect_list
                                                                     /* Apply rule 9 */
2447
2448
                                                                      /* Apply rule 8 */
            fname
                                 NAME
2449
2450
            brace group
                                 Lbrace compound list Rbrace
2451
                                 Do compound list Done
                                                                      /* Apply rule 6 */
2452
            do group
2453
2454
            simple command
                                  cmd prefix cmd word cmd suffix
2455
                                  cmd prefix cmd word
2456
                                  cmd prefix
2457
                                  cmd_name cmd_suffix
2458
                                  cmd name
2459
2460
            cmd name
                                 WORD
                                                            /* Apply rule 7a */
2461
                                                            /* Apply rule 7b */
            cmd word
                                 WORD
2462
2463
2464
            cmd prefix
                                              io redirect
                                  cmd prefix io redirect
2465
                                              ASSIGNMENT WORD
2466
                                  cmd prefix ASSIGNMENT WORD
2467
2468
2469
            cmd suffix
                                              io redirect
2470
                                  cmd suffix io redirect
2471
                                              WORD
2472
                                  cmd suffix WORD
2473
            redirect list
2474
                                                  io redirect
                                 redirect_list io_redirect
2475
2476
2477
            io redirect
                                             io_file
                                  IO NUMBER io file
2478
2479
                                             io here
2480
                                  IO NUMBER io here
2481
                                 ′<′
2482
            io file
                                             filename
                                 LESSAND
                                             filename
2483
2484
                                             filename
2485
                                 GREATAND
                                            filename
                                             filename
2486
                                 DGREAT
                                 LESSGREAT filename
2487
                                             filename
2488
                                 CLOBBER
2489
                                                               /* Apply rule 2 */
2490
            filename
                                 WORD
2491
2492
            io here
                                 DLESS
                                             here end
                               :
2493
                                 DLESSDASH here end
2494
2495
            here end
                                 WORD
                                                               /* Apply rule 3 */
2496
```

```
2497
            newline list
                                                  NEWLINE
2498
                                   newline list NEWLINE
2499
            linebreak
                                 : newline list
2500
2501
                                   /* empty */
2502
2503
            separator op
2504
2505
2506
            separator
                                   separator op linebreak
2507
                                   newline list
2508
                                  ';' linebreak
2509
            sequential sep
                                  newline list
2510
2511
```

2.11 Signals and Error Handling

2512

2513

2514 2515

2516

2517

2518

2519 2520

2521

25222523

25252526

2527

2530

25312532

2533

When a command is in an asynchronous list, the shell shall prevent SIGQUIT and SIGINT signals from the keyboard from interrupting the command. Otherwise, signals shall have the values inherited by the shell from its parent (see also the *trap* special built-in).

When a signal for which a trap has been set is received while the shell is waiting for the completion of a utility executing a foreground command, the trap associated with that signal shall not be executed until after the foreground command has completed. When the shell is waiting, by means of the *wait* utility, for asynchronous commands to complete, the reception of a signal for which a trap has been set shall cause the *wait* utility to return immediately with an exit status >128, immediately after which the trap associated with that signal shall be taken.

If multiple signals are pending for the shell for which there are associated trap actions, the order of execution of trap actions is unspecified.

2.12 Shell Execution Environment

A shell execution environment consists of the following:

- Open files inherited upon invocation of the shell, plus open files controlled by exec
- Working directory as set by cd
- File creation mask set by *umask*
- Current traps set by *trap*
 - Shell parameters that are set by variable assignment (see the *set* special built-in) or from the System Interfaces volume of IEEE Std 1003.1-2001 environment inherited by the shell when it begins (see the *export* special built-in)
 - Shell functions; see Section 2.9.5 (on page 54)
- Options turned on at invocation or by *set*
- Process IDs of the last commands in asynchronous lists known to this shell environment; see Section 2.9.3.1 (on page 50)

• Shell aliases; see Section 2.3.1 (on page 32)

Utilities other than the special built-ins (see Section 2.14 (on page 64)) shall be invoked in a separate environment that consists of the following. The initial value of these objects shall be the same as that for the parent shell, except as noted below.

- Open files inherited on invocation of the shell, open files controlled by the *exec* special builtin plus any modifications, and additions specified by any redirections to the utility
- Current working directory
- File creation mask
- If the utility is a shell script, traps caught by the shell shall be set to the default values and traps ignored by the shell shall be set to be ignored by the utility; if the utility is not a shell script, the trap actions (default or ignore) shall be mapped into the appropriate signal handling actions for the utility
- Variables with the *export* attribute, along with those explicitly exported for the duration of the command, shall be passed to the utility environment variables

The environment of the shell process shall not be changed by the utility unless explicitly specified by the utility description (for example, *cd* and *umask*).

A subshell environment shall be created as a duplicate of the shell environment, except that signal traps set by that shell environment shall be set to the default values. Changes made to the subshell environment shall not affect the shell environment. Command substitution, commands that are grouped with parentheses, and asynchronous lists shall be executed in a subshell environment. Additionally, each command of a multi-command pipeline is in a subshell environment; as an extension, however, any or all commands in a pipeline may be executed in the current environment. All other commands shall be executed in the current shell environment.

2.13 Pattern Matching Notation

The pattern matching notation described in this section is used to specify patterns for matching strings in the shell. Historically, pattern matching notation is related to, but slightly different from, the regular expression notation described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 9, Regular Expressions. For this reason, the description of the rules for this pattern matching notation are based on the description of regular expression notation, modified to include backslash escape processing.

2568 2.13.1 Patterns Matching a Single Character

The following patterns matching a single character shall match a single character: ordinary characters, special pattern characters, and pattern bracket expressions. The pattern bracket expression also shall match a single collating element. A backslash character shall escape the following character. The escaping backslash shall be discarded.

An ordinary character is a pattern that shall match itself. It can be any character in the supported character set except for NUL, those special shell characters in Section 2.2 (on page 30) that require quoting, and the following three special pattern characters. Matching shall be based on the bit pattern used for encoding the character, not on the graphic representation of the character. If any character (ordinary, shell special, or pattern special) is quoted, that pattern shall match the character itself. The shell special characters always require quoting.

When unquoted and outside a bracket expression, the following three characters shall have special meaning in the specification of patterns:

- ? A question-mark is a pattern that shall match any character.
- * An asterisk is a pattern that shall match multiple characters, as described in Section 2.13.2.
 - [The open bracket shall introduce a pattern bracket expression.

The description of basic regular expression bracket expressions in the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.3.5, RE Bracket Expression shall also apply to the pattern bracket expression, except that the exclamation mark character ('!') shall replace the circumflex character ('^') in its role in a "non-matching list" in the regular expression notation. A bracket expression starting with an unquoted circumflex character produces unspecified results.

When pattern matching is used where shell quote removal is not performed (such as in the argument to the *find* – *name* primary when *find* is being called using one of the *exec* functions as defined in the System Interfaces volume of IEEE Std 1003.1-2001, or in the *pattern* argument to the *finmatch()* function), special characters can be escaped to remove their special meaning by preceding them with a backslash character. This escaping backslash is discarded. The sequence "\\" represents one literal backslash. All of the requirements and effects of quoting on ordinary, shell special, and special pattern characters shall apply to escaping in this context.

2.13.2 Patterns Matching Multiple Characters

The following rules are used to construct patterns matching multiple characters from patterns matching a single character:

- 1. The asterisk (' * ') is a pattern that shall match any string, including the null string.
- 2. The concatenation of patterns matching a single character is a valid pattern that shall match the concatenation of the single characters or collating elements matched by each of the concatenated patterns.
- 3. The concatenation of one or more patterns matching a single character with one or more asterisks is a valid pattern. In such patterns, each asterisk shall match a string of zero or more characters, matching the greatest possible number of characters that still allows the remainder of the pattern to match the string.

2.13.3 Patterns Used for Filename Expansion

The rules described so far in Section 2.13.1 (on page 62) and Section 2.13.2 are qualified by the following rules that apply when pattern matching notation is used for filename expansion:

- 1. The slash character in a pathname shall be explicitly matched by using one or more slashes in the pattern; it shall neither be matched by the asterisk or question-mark special characters nor by a bracket expression. Slashes in the pattern shall be identified before bracket expressions; thus, a slash cannot be included in a pattern bracket expression used for filename expansion. If a slash character is found following an unescaped open square bracket character before a corresponding closing square bracket is found, the open bracket shall be treated as an ordinary character. For example, the pattern "a[b/c]d" does not match such pathnames as **abd** or **a/d**. It only matches a pathname of literally **a[b/c]d**.
- 2. If a filename begins with a period (' . '), the period shall be explicitly matched by using a period as the first character of the pattern or immediately following a slash character. The leading period shall not be matched by:

- The asterisk or question-mark special characters
- A bracket expression containing a non-matching list, such as "[!a]", a range expression, such as "[%-0]", or a character class expression, such as "[[:punct:]]"

It is unspecified whether an explicit period in a bracket expression matching list, such as "[.abc]", can match a leading period in a filename.

3. Specified patterns shall be matched against existing filenames and pathnames, as appropriate. Each component that contains a pattern character shall require read permission in the directory containing that component. Any component, except the last, that does not contain a pattern character shall require search permission. For example, given the pattern:

```
/foo/bar/x*/bam
```

search permission is needed for directories / and **foo**, search and read permissions are needed for directory **bar**, and search permission is needed for each \mathbf{x}^* directory. If the pattern matches any existing filenames or pathnames, the pattern shall be replaced with those filenames and pathnames, sorted according to the collating sequence in effect in the current locale. If the pattern contains an invalid bracket expression or does not match any existing filenames or pathnames, the pattern string shall be left unchanged.

2.14 Special Built-In Utilities

The following "special built-in" utilities shall be supported in the shell command language. The output of each command, if any, shall be written to standard output, subject to the normal redirection and piping possible with all commands.

The term "built-in" implies that the shell can execute the utility directly and does not need to search for it. An implementation may choose to make any utility a built-in; however, the special built-in utilities described here differ from regular built-in utilities in two respects:

- 1. A syntax error in a special built-in utility may cause a shell executing that utility to abort, while a syntax error in a regular built-in utility shall not cause a shell executing that utility to abort. (See Section 2.8.1 (on page 46) for the consequences of errors on interactive and non-interactive shells.) If a special built-in utility encountering a syntax error does not abort the shell, its exit value shall be non-zero.
- 2. Variable assignments specified with special built-in utilities remain in effect after the built-in completes; this shall not be the case with a regular built-in or other utility.

The special built-in utilities in this section need not be provided in a manner accessible via the *exec* family of functions defined in the System Interfaces volume of IEEE Std 1003.1-2001.

Some of the special built-ins are described as conforming to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines. For those that are not, the requirement in Section 1.11 (on page 20) that "--" be recognized as a first argument to be discarded does not apply and a conforming application shall not use that argument.

Shell Command Language break

2659 **NAME** break — exit from for, while, or until loop 2660 2661 **SYNOPSIS** break [n] 2662 2663 DESCRIPTION The break utility shall exit from the smallest enclosing for, while, or until loop, if any; or from the 2664 *n*th enclosing loop if *n* is specified. The value of *n* is an unsigned decimal integer greater than or 2665 equal to 1. The default shall be equivalent to n=1. If n is greater than the number of enclosing 2666 loops, the outermost enclosing loop shall be exited. Execution shall continue with the command 2667 2668 immediately following the loop. **OPTIONS** 2669 None. 2670 **OPERANDS** 2671 See the DESCRIPTION. 2672 **STDIN** 2673 Not used. 2674 **INPUT FILES** 2675 None. 2676 **ENVIRONMENT VARIABLES** 2677 2678 None. **ASYNCHRONOUS EVENTS** Default. 2680 **STDOUT** 2681 2682 Not used. **STDERR** 2683 2684 The standard error shall be used only for diagnostic messages. **OUTPUT FILES** 2685 None. 2686 EXTENDED DESCRIPTION 2687 2688 None. **EXIT STATUS** 2689 Successful completion. 2690 >0 The *n* value was not an unsigned decimal integer greater than or equal to 1. 2691 **CONSEQUENCES OF ERRORS** 2692 Default. 2693

break Shell Command Language

```
APPLICATION USAGE
2694
2695
             None.
     EXAMPLES
2696
             for i in * do
2697
                  if test -d "$i" then break fi done
2698
     RATIONALE
2699
             In early proposals, consideration was given to expanding the syntax of break and continue to refer
2700
             to a label associated with the appropriate loop as a preferable alternative to the n method.
2701
             However, this volume of IEEE Std 1003.1-2001 does reserve the name space of command names
2702
2703
             ending with a colon. It is anticipated that a future implementation could take advantage of this
             and provide something like:
2704
2705
             outofloop: for i in a b c d e
2706
             do
                  for j in 0 1 2 3 4 5 6 7 8 9
2707
2708
                  do
                        if test -r "${i}${j}"
2709
2710
                        then break outofloop
2711
                        fi
2712
                  done
2713
             done
             and that this might be standardized after implementation experience is achieved.
2714
    FUTURE DIRECTIONS
2715
             None.
2716
     SEE ALSO
2717
             Section 2.14 (on page 64)
2718
    CHANGE HISTORY
2719
2720
    Issue 6
             IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page
2721
             sections use terms as described in the Utility Description Defaults (Section 1.11). No change in
2722
             behavior is intended.
2723
```

colon

```
2724
    NAME
             colon — null utility
2725
2726
     SYNOPSIS
             : [argument ...]
2727
     DESCRIPTION
2728
             This utility shall only expand command arguments. It is used when a command is needed, as in
             the then condition of an if command, but nothing is to be done by the command.
2730
     OPTIONS
2731
             None.
     OPERANDS
2733
             See the DESCRIPTION.
2734
     STDIN
2735
             Not used.
2736
    INPUT FILES
2737
2738
     ENVIRONMENT VARIABLES
2739
             None.
2740
     ASYNCHRONOUS EVENTS
2741
             Default.
2742
    STDOUT
2743
             Not used.
2744
     STDERR
             The standard error shall be used only for diagnostic messages.
2746
     OUTPUT FILES
2747
2748
             None.
     EXTENDED DESCRIPTION
2749
2750
             None.
     EXIT STATUS
2751
             Zero.
2752
     CONSEQUENCES OF ERRORS
2753
             Default.
2754
     APPLICATION USAGE
2755
2756
             None.
     EXAMPLES
2757
             : ${X=abc}
2758
             if
2759
                      false
2760
             then
2761
             else
                      echo $X
             fi
2762
2763
             abc
             As with any of the special built-ins, the null utility can also have variable assignments and
2764
             redirections associated with it, such as:
2765
```

colon

2766	X=Y : > Z
2767 2768	which sets variable x to the value y (so that it persists after the null utility completes) and creates or truncates file z .
2769	RATIONALE
2770	None.
2771	FUTURE DIRECTIONS
2772	None.
2773	SEE ALSO
2774	Section 2.14 (on page 64)
2775	CHANGE HISTORY
2776	Issue 6
2777	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page
2778	sections use terms as described in the Utility Description Defaults (Section 1.11). No change in
2779	behavior is intended.

Shell Command Language continue

2780 **NAME** 2781 continue — continue for, while, or until loop 2782 **SYNOPSIS** continue [n] 2783 **DESCRIPTION** 2784 The *continue* utility shall return to the top of the smallest enclosing **for**, **while**, or **until** loop, or to 2785 the top of the nth enclosing loop, if n is specified. This involves repeating the condition list of a 2786 while or until loop or performing the next assignment of a for loop, and re-executing the loop if 2787 appropriate. 2788 The value of *n* is a decimal integer greater than or equal to 1. The default shall be equivalent to 2789 n=1. If n is greater than the number of enclosing loops, the outermost enclosing loop shall be 2790 2791 used. **OPTIONS** 2792 2793 None. **OPERANDS** 2794 See the DESCRIPTION. 2795 **STDIN** 2796 Not used. 2797 **INPUT FILES** 2798 2799 None. **ENVIRONMENT VARIABLES** 2800 None. 2801 ASYNCHRONOUS EVENTS 2802 2803 Default. **STDOUT** 2804 2805 Not used. **STDERR** 2806 The standard error shall be used only for diagnostic messages. 2807 **OUTPUT FILES** 2808 2809 None. **EXTENDED DESCRIPTION** 2810 2811 None. **EXIT STATUS** 2812 Successful completion. 2813 >0 The *n* value was not an unsigned decimal integer greater than or equal to 1. 2814 **CONSEQUENCES OF ERRORS** 2815 Default. 2816

continue

```
APPLICATION USAGE
2817
2818
             None.
    EXAMPLES
2819
             for i in *
2820
             do
2821
                  if test -d "$i"
2822
                  then continue
2823
2824
                  echo "\"$i\"" is not a directory.
2825
2826
    RATIONALE
2827
             None.
2828
    FUTURE DIRECTIONS
2829
             None.
2830
    SEE ALSO
2831
             Section 2.14 (on page 64)
2832
    CHANGE HISTORY
2833
    Issue 6
2834
2835
             IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page
             sections use terms as described in the Utility Description Defaults (Section 1.11). No change in
2836
2837
             behavior is intended.
```

dot

2838 2839	NAME dot — execute commands in the current environment
2840 2841	SYNOPSIS . file
2842 2843	DESCRIPTION The shell shall execute commands from the <i>file</i> in the current environment.
2844 2845 2846 2847 2848	If <i>file</i> does not contain a slash, the shell shall use the search path specified by <i>PATH</i> to find the directory containing <i>file</i> . Unlike normal command search, however, the file searched for by the <i>dot</i> utility need not be executable. If no readable file is found, a non-interactive shell shall abort; an interactive shell shall write a diagnostic message to standard error, but this condition shall not be considered a syntax error.
2849 2850	OPTIONS None.
2851 2852	OPERANDS See the DESCRIPTION.
2853 2854	STDIN Not used.
2855 2856	INPUT FILES See the DESCRIPTION.
2857 2858	ENVIRONMENT VARIABLES See the DESCRIPTION.
2859 2860	ASYNCHRONOUS EVENTS Default.
2861 2862	STDOUT Not used.
2863 2864	STDERR The standard error shall be used only for diagnostic messages.
2865 2866	OUTPUT FILES None.
2867 2868	EXTENDED DESCRIPTION None.
2869 2870	EXIT STATUS Returns the value of the last command executed, or a zero exit status if no command is executed.
2871 2872	CONSEQUENCES OF ERRORS Default.

dot

```
APPLICATION USAGE
2873
2874
             None.
     EXAMPLES
2875
             cat foobar
2876
             foo=hello bar=world
2877
              . foobar
2878
             echo $foo $bar
2879
2880
             hello world
     RATIONALE
2881
2882
             Some older implementations searched the current directory for the file, even if the value of PATH
             disallowed it. This behavior was omitted from this volume of IEEE Std 1003.1-2001 due to
2883
             concerns about introducing the susceptibility to trojan horses that the user might be trying to
2884
             avoid by leaving dot out of PATH.
2885
             The KornShell version of dot takes optional arguments that are set to the positional parameters.
2886
             This is a valid extension that allows a dot script to behave identically to a function.
2887
     FUTURE DIRECTIONS
2888
             None.
2889
    SEE ALSO
2890
2891
             Section 2.14 (on page 64)
     CHANGE HISTORY
2892
     Issue 6
2893
             IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page
2894
             sections use terms as described in the Utility Description Defaults (Section 1.11). No change in
2895
```

2896

behavior is intended.

```
2897
    NAME
2898
             eval — construct command by concatenating arguments
2899
    SYNOPSIS
             eval [argument ...]
2900
2901
    DESCRIPTION
             The eval utility shall construct a command by concatenating arguments together, separating each
2902
             with a <space>. The constructed command shall be read and executed by the shell.
2903
     OPTIONS
2904
             None.
2905
     OPERANDS
2906
             See the DESCRIPTION.
2907
    STDIN
2908
             Not used.
2909
    INPUT FILES
2910
2911
    ENVIRONMENT VARIABLES
2912
             None.
2913
    ASYNCHRONOUS EVENTS
2914
             Default.
2915
    STDOUT
2916
             Not used.
2917
    STDERR
2918
             The standard error shall be used only for diagnostic messages.
2919
     OUTPUT FILES
2920
2921
             None.
    EXTENDED DESCRIPTION
2922
2923
             None.
    EXIT STATUS
2924
             If there are no arguments, or only null arguments, eval shall return a zero exit status; otherwise, it
2925
             shall return the exit status of the command defined by the string of concatenated arguments
2926
2927
             separated by <space>s.
     CONSEQUENCES OF ERRORS
2928
             Default.
2929
    APPLICATION USAGE
2930
             None.
2931
     EXAMPLES
2932
             foo=10 x=foo
2933
2934
             y='$'$x
2935
             echo $y
             $foo
2936
2937
             eval y='$'$x
             echo $y
2938
2939
             10
```

2940 2941	RATIONALE None.
2942 2943	FUTURE DIRECTIONS None.
2944 2945	SEE ALSO Section 2.14 (on page 64)
2946	CHANGE HISTORY
2947	Issue 6
2948	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page
2949	sections use terms as described in the Utility Description Defaults (Section 1.11). No change in
2950	behavior is intended.

2990

2991

2951 **NAME** exec — execute commands and open, close, or copy file descriptors 2952 2953 **SYNOPSIS** 2954 exec [command [argument ...]] DESCRIPTION 2955 The exec utility shall open, close, and/or copy file descriptors as specified by any redirections as 2956 part of the command. 2957 If exec is specified without command or arguments, and any file descriptors with numbers greater 2958 than 2 are opened with associated redirection statements, it is unspecified whether those file descriptors remain open when the shell invokes another utility. Scripts concerned that child 2960 shells could misuse open file descriptors can always close them explicitly, as shown in one of the 2961 following examples. 2962 If exec is specified with command, it shall replace the shell with command without creating a new 2963 process. If *arguments* are specified, they shall be arguments to *command*. Redirection affects the 2964 current shell execution environment. 2965 **OPTIONS** 2966 None. 2967 **OPERANDS** 2968 See the DESCRIPTION. 2969 **STDIN** 2970 Not used. 2971 **INPUT FILES** 2972 None. 2973 2974 **ENVIRONMENT VARIABLES** None. 2975 2976 ASYNCHRONOUS EVENTS Default. 2977 2978 STDOUT Not used. 2979 **STDERR** 2980 The standard error shall be used only for diagnostic messages. 2981 2982 **OUTPUT FILES** None. 2983 EXTENDED DESCRIPTION 2984 None. 2985 **EXIT STATUS** 2986 If command is specified, exec shall not return to the shell; rather, the exit status of the process shall 2987 2988

If *command* is specified, *exec* shall not return to the shell; rather, the exit status of the process shall be the exit status of the program implementing *command*, which overlaid the shell. If *command* is not found, the exit status shall be 127. If *command* is found, but it is not an executable utility, the exit status shall be 126. If a redirection error occurs (see Section 2.8.1 (on page 46)), the shell shall exit with a value in the range 1–125. Otherwise, *exec* shall return a zero exit status.

```
CONSEQUENCES OF ERRORS
2992
2993
              Default.
     APPLICATION USAGE
2994
              None.
2995
     EXAMPLES
2996
              Open readfile as file descriptor 3 for reading:
2997
2998
              exec 3< readfile
              Open writefile as file descriptor 4 for writing:
2999
              exec 4> writefile
3000
              Make file descriptor 5 a copy of file descriptor 0:
3001
              exec 5<&0
3002
3003
              Close file descriptor 3:
              exec 3<&-
3004
              Cat the file maggie by replacing the current shell with the cat utility:
3005
              exec cat maggie
3006
     RATIONALE
3007
              Most historical implementations were not conformant in that:
3008
3009
              foo=bar exec cmd
3010
              did not pass foo to cmd.
     FUTURE DIRECTIONS
3011
3012
              None.
     SEE ALSO
3013
3014
              Section 2.14 (on page 64)
     CHANGE HISTORY
     Issue 6
3016
              IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page
3017
3018
              sections use terms as described in the Utility Description Defaults (Section 1.11). No change in
3019
              behavior is intended.
```

3020 3021	NAME exit — cause the shell to exit			
3022	SYNOPSIS			
3023	exit [n]			
3024 3025 3026 3027	DESCRIPTION The <i>exit</i> utility shall cause the shell to exit with the exit status specified by the unsigned decimal integer <i>n</i> . If <i>n</i> is specified, but its value is not between 0 and 255 inclusively, the exit status is undefined.			
3028 3029	A <i>trap</i> on EXIT shall be executed before the shell terminates, except when the <i>exit</i> utility is invoked in that <i>trap</i> itself, in which case the shell shall exit immediately.			
3030 3031	OPTIONS None.			
3032 3033	OPERANDS See the DESCRIPTION.			
3034 3035	STDIN Not used.			
3036 3037	INPUT FILES None.			
3038 3039	ENVIRONMENT VARIABLES None.			
3040 3041	ASYNCHRONOUS EVENTS Default.			
3042 3043	STDOUT Not used.			
3044 3045	STDERR The standard error shall be used only for diagnostic messages.			
3046 3047	OUTPUT FILES None.			
3048 3049	EXTENDED DESCRIPTION None.			
3050 3051 3052 3053 3054	EXIT STATUS The exit status shall be <i>n</i> , if specified. Otherwise, the value shall be the exit value of the last command executed, or zero if no command was executed. When <i>exit</i> is executed in a <i>trap</i> action, the last command is considered to be the command that executed immediately preceding the <i>trap</i> action.			
3055 3056	CONSEQUENCES OF ERRORS Default.			

```
APPLICATION USAGE
3057
3058
              None.
     EXAMPLES
3059
              Exit with a true value:
3060
3061
              exit 0
              Exit with a false value:
3062
3063
     RATIONALE
3064
              As explained in other sections, certain exit status values have been reserved for special uses and
3065
3066
              should be used by applications only for those purposes:
                       A file to be executed was found, but it was not an executable utility.
3067
               127
                       A utility to be executed was not found.
3068
3069
              >128
                       A command was interrupted by a signal.
     FUTURE DIRECTIONS
3070
              None.
3071
     SEE ALSO
3072
              Section 2.14 (on page 64)
3073
     CHANGE HISTORY
     Issue 6
3075
              IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page
3076
              sections use terms as described in the Utility Description Defaults (Section 1.11). No change in
3077
3078
              behavior is intended.
```

```
3079
    NAME
              export — set the export attribute for variables
3080
3081
              export name[=word]...
3082
              export -p
3083
     DESCRIPTION
3084
3085
              The shell shall give the export attribute to the variables corresponding to the specified names,
              which shall cause them to be in the environment of subsequently executed commands. If the
3086
              name of a variable is followed by =word, then the value of that variable shall be set to word.
3087
              The export special built-in shall support the Base Definitions volume of IEEE Std 1003.1-2001,
3088
              Section 12.2, Utility Syntax Guidelines.
3089
              When -p is specified, export shall write to the standard output the names and values of all
3090
              exported variables, in the following format:
3091
              "export %s=%s\n", <name>, <value>
3092
              if name is set, and:
3093
              "export %s\n", <name>
3094
              if name is unset.
3095
              The shell shall format the output, including the proper use of quoting, so that it is suitable for
3096
              reinput to the shell as commands that achieve the same exporting results, except:
3097
               1. Read-only variables with values cannot be reset.
3098
3099
                   Variables that were unset at the time they were output need not be reset to the unset state
                   if a value is assigned to the variable between the time the state was saved and the time at
3100
                   which the saved output is reinput to the shell.
3101
3102
              When no arguments are given, the results are unspecified.
     OPTIONS
3103
              See the DESCRIPTION.
3104
     OPERANDS
3105
              See the DESCRIPTION.
3106
     STDIN
3107
              Not used.
3108
     INPUT FILES
3109
3110
     ENVIRONMENT VARIABLES
3111
              None.
3112
     ASYNCHRONOUS EVENTS
3113
              Default.
3114
    STDOUT
              See the DESCRIPTION.
3116
```

```
3117
     STDERR
              The standard error shall be used only for diagnostic messages.
3118
3119
     OUTPUT FILES
              None.
3120
3121
     EXTENDED DESCRIPTION
              None.
     EXIT STATUS
3123
              Zero.
3124
3125
     CONSEQUENCES OF ERRORS
              Default.
3126
     APPLICATION USAGE
3127
3128
              None.
     EXAMPLES
3129
              Export PWD and HOME variables:
3130
3131
              export PWD HOME
3132
              Set and export the PATH variable:
              export PATH=/local/bin:$PATH
3133
3134
              Save and restore all exported variables:
3135
              export -p > temp-file
3136
              unset a lot of variables
3137
              ... processing
3138
              . temp-file
     RATIONALE
3139
3140
              Some historical shells use the no-argument case as the functional equivalent of what is required
3141
              here with -p. This feature was left unspecified because it is not historical practice in all shells,
              and some scripts may rely on the now-unspecified results on their implementations. Attempts to
3142
3143
              specify the -\mathbf{p} output as the default case were unsuccessful in achieving consensus. The -\mathbf{p}
3144
              option was added to allow portable access to the values that can be saved and then later restored
              using; for example, a dot script.
3145
     FUTURE DIRECTIONS
3146
3147
              None.
     SEE ALSO
3148
              Section 2.14 (on page 64)
3149
     CHANGE HISTORY
     Issue 6
3151
3152
              IEEE PASC Interpretation 1003.2 #203 is applied, clarifying the format when a variable is unset.
              IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page
3153
              sections use terms as described in the Utility Description Defaults (Section 1.11). No change in
3154
3155
              behavior is intended.
              IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/6 is applied, adding the following text to
3156
              the end of the first paragraph of the DESCRIPTION: "If the name of a variable is followed by
3157
3158
              =word, then the value of that variable shall be set to word.". The reason for this change is that the
              SYNOPSIS for export includes:
3159
```

export name [=word]...but the meaning of the optional "=word" is never explained in the text.

readonly Shell Command Language

```
3162
     NAME
              readonly — set the readonly attribute for variables
3163
3164
3165
              readonly name[=word]...
3166
              readonly -p
     DESCRIPTION
3167
3168
              The variables whose names are specified shall be given the readonly attribute. The values of
              variables with the readonly attribute cannot be changed by subsequent assignment, nor can those
3169
              variables be unset by the unset utility. If the name of a variable is followed by =word, then the
3170
              value of that variable shall be set to word.
3171
3172
              The readonly special built-in shall support the Base Definitions volume of IEEE Std 1003.1-2001,
3173
              Section 12.2, Utility Syntax Guidelines.
              When -\mathbf{p} is specified, readonly writes to the standard output the names and values of all read-
3174
3175
              only variables, in the following format:
              "readonly %s=%s\n", <name>, <value>
3176
              if name is set, and
3177
              "readonly %s\n", <name>
3178
              if name is unset.
3179
              The shell shall format the output, including the proper use of quoting, so that it is suitable for
3180
3181
              reinput to the shell as commands that achieve the same value and readonly attribute-setting
3182
              results in a shell execution environment in which:
3183
                   Variables with values at the time they were output do not have the readonly attribute set.
                2. Variables that were unset at the time they were output do not have a value at the time at
3184
3185
                    which the saved output is reinput to the shell.
              When no arguments are given, the results are unspecified.
3186
     OPTIONS
3187
              See the DESCRIPTION.
3188
     OPERANDS
3189
              See the DESCRIPTION.
3190
     STDIN
3191
              Not used.
3192
     INPUT FILES
3193
              None.
3194
     ENVIRONMENT VARIABLES
3195
3196
     ASYNCHRONOUS EVENTS
3197
              Default.
3198
     STDOUT
3199
              See the DESCRIPTION.
3200
```

readonly Shell Command Language

3201 3202	STDERR The standard error shall be used only for diagnostic messages.	ı
3203 3204	OUTPUT FILES None.	
3205 3206	EXTENDED DESCRIPTION None.	
3207 3208	EXIT STATUS Zero.	
3209 3210	CONSEQUENCES OF ERRORS Default.	1
3211 3212	APPLICATION USAGE None.	
3213 3214	EXAMPLES readonly HOME PWD	
3215 3216 3217	RATIONALE Some historical shells preserve the <i>readonly</i> attribute across separate invocations. This volume of IEEE Std 1003.1-2001 allows this behavior, but does not require it.	
3218 3219 3220	The $-\mathbf{p}$ option allows portable access to the values that can be saved and then later restored using, for example, a <i>dot</i> script. Also see the RATIONALE for <i>export</i> (on page 79) for a description of the no-argument and $-\mathbf{p}$ output cases and a related example.	
3221 3222 3223 3224 3225	Read-only functions were considered, but they were omitted as not being historical practice or particularly useful. Furthermore, functions must not be read-only across invocations to preclude "spoofing" (spoofing is the term for the practice of creating a program that acts like a well-known utility with the intent of subverting the real intent of the user) of administrative or security-relevant (or security-conscious) shell scripts.	
3226 3227	FUTURE DIRECTIONS None.	
3228 3229	SEE ALSO Section 2.14 (on page 64)	
3230	CHANGE HISTORY	
3231 3232	Issue 6 IEEE PASC Interpretation 1003.2 #203 is applied, clarifying the format when a variable is unset.	
3233 3234 3235	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page sections use terms as described in the Utility Description Defaults (Section 1.11). No change in behavior is intended.	
3236 3237 3238 3239	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/7 is applied, adding the following text to the end of the first paragraph of the DESCRIPTION: "If the name of a variable is followed by =word, then the value of that variable shall be set to word.". The reason for this change is that the SYNOPSIS for readonly includes:	
3240	readonly name[=word]	
3241	but the meaning of the optional "=word" is never explained in the text.	

```
3242
    NAME
             return — return from a function
3243
3244
     SYNOPSIS
3245
             return [n]
     DESCRIPTION
3246
             The return utility shall cause the shell to stop executing the current function or dot script. If the
             shell is not currently executing a function or dot script, the results are unspecified.
3248
     OPTIONS
3249
             None.
     OPERANDS
3251
             See the DESCRIPTION.
3252
     STDIN
3253
             Not used.
3254
     INPUT FILES
3255
3256
     ENVIRONMENT VARIABLES
3257
             None.
3258
     ASYNCHRONOUS EVENTS
3259
             Default.
3260
     STDOUT
3261
             Not used.
3262
     STDERR
3263
             The standard error shall be used only for diagnostic messages.
3264
     OUTPUT FILES
3265
             None.
3266
     EXTENDED DESCRIPTION
3267
3268
             None.
     EXIT STATUS
3269
             The value of the special parameter '?' shall be set to n, an unsigned decimal integer, or to the
3270
             exit status of the last command executed if n is not specified. If the value of n is greater than 255,
3271
             the results are undefined. When return is executed in a trap action, the last command is
3272
3273
             considered to be the command that executed immediately preceding the trap action.
     CONSEQUENCES OF ERRORS
3274
             Default.
3275
     APPLICATION USAGE
3276
             None.
     EXAMPLES
3278
3279
             None.
     RATIONALE
3280
             The behavior of return when not in a function or dot script differs between the System V shell
3281
             and the KornShell. In the System V shell this is an error, whereas in the KornShell, the effect is
3282
             the same as exit.
3283
```

3284	The results of returning a number greater than 255 are undefined because of differing practices				
3285	in the various historical implementations. Some shells AND out all but the low-order 8 bits;				
3286	others allow larger values, but not of unlimited size.				
0200	others and windger variety, but not or unminited size.				
3287	See the discussion of appropriate exit status values under exit (on page 77).				
3288	FUTURE DIRECTIONS				
3289	None.				
3290	SEE ALSO				
3291	Section 2.14 (on page 64)				
3292	CHANGE HISTORY				
3293	Issue 6				
3294	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page				
3295	sections use terms as described in the Utility Description Defaults (Section 1.11). No change in				
	behavior is intended.				
3296	benavior is intended.				

```
3297
    NAME
            set — set or unset options and positional parameters
3298
3299
    SYNOPSIS
             set [-abCefmnuvx] [-h] [-o option] [argument...]
3300
    XSI
            set [+abCefmnuvx] [+h] [+o option] [argument...]
3301
    XSI
3302
            set -- [argument...]
3303
            set -o
            set +o
    DESCRIPTION
3305
```

If no *options* or *arguments* are specified, *set* shall write the names and values of all shell variables in the collation sequence of the current locale. Each *name* shall start on a separate line, using the format:

```
"%s=%sn", <name>, <value>
```

The *value* string shall be written with appropriate quoting; see the description of shell quoting in Section 2.2 (on page 30). The output shall be suitable for reinput to the shell, setting or resetting, as far as possible, the variables that are currently set; read-only variables cannot be reset.

When options are specified, they shall set or unset attributes of the shell, as described below. When *arguments* are specified, they cause positional parameters to be set or unset, as described below. Setting or unsetting attributes and positional parameters are not necessarily related actions, but they can be combined in a single invocation of *set*.

The *set* special built-in shall support the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines except that options can be specified with either a leading hyphen (meaning enable the option) or plus sign (meaning disable it) unless otherwise specified.

Implementations shall support the options in the following list in both their hyphen and plussign forms. These options can also be specified as options to *sh*.

- -a When this option is on, the *export* attribute shall be set for each variable to which an assignment is performed; see the Base Definitions volume of IEEE Std 1003.1-2001, Section 4.21, Variable Assignment. If the assignment precedes a utility name in a command, the *export* attribute shall not persist in the current execution environment after the utility completes, with the exception that preceding one of the special built-in utilities causes the *export* attribute to persist after the built-in has completed. If the assignment does not precede a utility name in the command, or if the assignment is a result of the operation of the *getopts* or *read* utilities, the *export* attribute shall persist until the variable is unset.
- -b This option shall be supported if the implementation supports the User Portability Utilities option. It shall cause the shell to notify the user asynchronously of background job completions. The following message is written to standard error:

```
"[%d]%c %s%s\n", <job-number>, <current>, <status>, <job-name> where the fields shall be as follows:
```

<current>

The character '+' identifies the job that would be used as a default for the *fg* or *bg* utilities; this job can also be specified using the *job_id* "%+" or "%%". The character '-' identifies the job that would become the default if the current default job were to exit; this job can also be specified using the *job_id* "%-". For other jobs, this field is a <space>. At most one job can be identified with '+' and at most one job can be identified with '-'.

3341 3342 3343			If there is any suspended job, then the current job shall be a suspended job. If there are at least two suspended jobs, then the previous job also shall be a suspended job.		
3344 3345 3346		<job-number< td=""><td colspan="2">A number that can be used to identify the process group to the <i>wait</i>, <i>fg</i>, and <i>kill</i> utilities. Using these utilities, the job can be identified prefixing the job number with '%'.</td></job-number<>	A number that can be used to identify the process group to the <i>wait</i> , <i>fg</i> , and <i>kill</i> utilities. Using these utilities, the job can be identified prefixing the job number with '%'.		
3347		<status></status>	Unspecified.		
3348		<job-name></job-name>	Unspecified.		
3349 3350 3351		ID from the	ell notifies the user a job has been completed, it may remove the job's process list of those known in the current shell execution environment; see Section age 50). Asynchronous notification shall not be enabled by default.		
3352 3353 3354	-С	operator (se	C.) Prevent existing files from being overwritten by the shell's $'>'$ redirection e Section 2.7.2 (on page 44)); the $"> $ "redirection operator shall override this ion for an individual file.		
3355 3356 3357 3358	-е	2.8.1 (on page following a v	When this option is on, if a simple command fails for any of the reasons listed in Section 2.8.1 (on page 46) or returns an exit status value >0, and is not part of the compound list following a while , until , or if keyword, and is not a part of an AND or OR list, and is not a pipeline preceded by the ! reserved word, then the shell shall immediately exit.		
3359	− f	The shell sha	ll disable pathname expansion.		
3360 XSI 3361	-h		Locate and remember utilities invoked by functions as those functions are defined (the utilities are normally located when the function is executed).		
3362 3363 3364 3365 3366 3367 3368 3369 3370	-m	This option shall be supported if the implementation supports the User Portability Utilities option. All jobs shall be run in their own process groups. Immediately before the shell issues a prompt after completion of the background job, a message reporting the exit status of the background job shall be written to standard error. If a foreground job stops, the shell shall write a message to standard error to that effect, formatted as described by the <i>jobs</i> utility. In addition, if a job changes status other than exiting (for example, if it stops for input or output or is stopped by a SIGSTOP signal), the shell shall write a similar message immediately prior to writing the next prompt. This option is enabled by default for interactive shells.			
3371 3372	-n		The shell shall read commands but does not execute them; this can be used to check for shell script syntax errors. An interactive shell may ignore this option.		
3373	-о	Write the cur	Write the current settings of the options to standard output in an unspecified format.		
3374 3375	+0	Write the current option settings to standard output in a format that is suitable for reinput to the shell as commands that achieve the same options settings.			
3376 3377 3378 3379	−o (This option is supported if the system supports the User Portability Utilities option. It shall set various options, many of which shall be equivalent to the single option letters. The following values of <i>option</i> shall be supported:			
3380		allexport	Equivalent to $-\mathbf{a}$.		
3381		errexit	Equivalent to $-\mathbf{e}$.		
3382 3383 3384		<i>ignoreeof</i> Prevent an interactive shell from exiting on end-of-file. This setting prevents accidental logouts when <control>-D is entered. A user shall explicitly <i>exit</i> to leave the interactive shell.</control>			

3385 3386	mo		Equivalent to -m. This option is supported if the system supports the User Portability Utilities option.		
3387	noc	clobber	Equivalent to -C (uppercase C).		
3388	nog	glob	Equivalent to $-\mathbf{f}$.		
3389	noe	exec	Equivalent to $-\mathbf{n}$.		
3390 3391	noi		Prevent the entry of function definitions into the command history; see Command History List (on page 854).		
3392	not	tify	Equivalent to $-\mathbf{b}$.		
3393	not	unset	Equivalent to -u.		
3394	ver	bose	Equivalent to $-\mathbf{v}$.		
3395 3396 3397	vi		Allow shell command line editing using the built-in <i>vi</i> editor. Enabling <i>vi</i> mode shall disable any other command line editing mode provided as an implementation extension.		
3398			It need not be possible to set <i>vi</i> mode on for certain block-mode terminals.		
3399	xtr	race	Equivalent to -x.		
3400 3401			ll write a message to standard error when it tries to expand a variable that is mediately exit. An interactive shell shall not exit.		
3402	−v Th	e shell shal	l write its input to standard error as it is read.		
3403 3404 3405	command and before it executes it. It is unspecified whether the command that turns				
3406 3407			these options shall be off (unset) unless stated otherwise in the description of ss the shell was invoked with them on; see <i>sh</i> .		
3408 3409 3410	parameter '#' shall be set to reflect the number of positional parameters. All positional				
3411 3412 3413 3414	delimit the arguments if the first argument begins with $'+'$ or $'-'$, or to prevent inadvertent listing of all shell variables when there are no arguments. The command $set-$ without $argument$				
3415	OPTIONS				
3416		DESCRIPT	TION.		
3417 3418		DESCRIPT	TION.		
3419 3420	STDIN Not use	ed.		ı	
	INPUT FILES			ı	
3422	None.				

```
3423
     ENVIRONMENT VARIABLES
              None.
3424
3425
     ASYNCHRONOUS EVENTS
              Default.
3426
     STDOUT
3427
              See the DESCRIPTION.
3428
3429
     STDERR
              The standard error shall be used only for diagnostic messages.
3430
3431
     OUTPUT FILES
              None.
3432
     EXTENDED DESCRIPTION
3433
3434
              None.
     EXIT STATUS
3435
              Zero.
3436
     CONSEQUENCES OF ERRORS
3437
              Default.
3438
     APPLICATION USAGE
3439
              None.
3440
     EXAMPLES
3441
              Write out all variables and their values:
3442
3443
              Set $1, $2, and $3 and set "$#" to 3:
3444
              set c a b
3445
              Turn on the -\mathbf{x} and -\mathbf{v} options:
3446
3447
              set -xv
              Unset all positional parameters:
3448
3449
              Set $1 to the value of x, even if it begins with '-' or '+':
3450
              set -- "$x"
3451
              Set the positional parameters to the expansion of x, even if x expands with a leading '-' or '+':
3452
              set -- $x
3453
     RATIONALE
3454
              The set – form is listed specifically in the SYNOPSIS even though this usage is implied by the
3455
              Utility Syntax Guidelines. The explanation of this feature removes any ambiguity about whether
3456
              the set -- form might be misinterpreted as being equivalent to set without any options or
3457
              arguments. The functionality of this form has been adopted from the KornShell. In System V, set
3458
              -- only unsets parameters if there is at least one argument; the only way to unset all parameters
3459
              is to use shift. Using the KornShell version should not affect System V scripts because there
3460
              should be no reason to issue it without arguments deliberately; if it were issued as, for example:
3461
              set -- "$@"
3462
```

and there were in fact no arguments resulting from "\$@", unsetting the parameters would have no result.

The *set* + form in early proposals was omitted as being an unnecessary duplication of *set* alone and not widespread historical practice.

The *noclobber* option was changed to allow *set* -**C** as well as the *set* -**o** *noclobber* option. The single-letter version was added so that the historical "\$-" paradigm would not be broken; see Section 2.5.2 (on page 34).

The $-\mathbf{h}$ flag is related to command name hashing and is only required on XSI-conformant systems.

The following *set* flags were omitted intentionally with the following rationale:

-k The -k flag was originally added by the author of the Bourne shell to make it easier for users of pre-release versions of the shell. In early versions of the Bourne shell the construct set name=value had to be used to assign values to shell variables. The problem with -k is that the behavior affects parsing, virtually precluding writing any compilers. To explain the behavior of -k, it is necessary to describe the parsing algorithm, which is implementation-defined. For example:

```
set -k; echo name=value
and:
set -k
echo name=value
```

behave differently. The interaction with functions is even more complex. What is more, the $-\mathbf{k}$ flag is never needed, since the command line could have been reordered.

-t The -t flag is hard to specify and almost never used. The only known use could be done with here-documents. Moreover, the behavior with ksh and sh differs. The reference page says that it exits after reading and executing one command. What is one command? If the input is date; date, sh executes both date commands while ksh does only the first.

Consideration was given to rewriting *set* to simplify its confusing syntax. A specific suggestion was that the *unset* utility should be used to unset options instead of using the non-*getopt()*-able + *option* syntax. However, the conclusion was reached that the historical practice of using + *option* was satisfactory and that there was no compelling reason to modify such widespread historical practice.

The $-\mathbf{o}$ option was adopted from the KornShell to address user needs. In addition to its generally friendly interface, $-\mathbf{o}$ is needed to provide the vi command line editing mode, for which historical practice yields no single-letter option name. (Although it might have been possible to invent such a letter, it was recognized that other editing modes would be developed and $-\mathbf{o}$ provides ample name space for describing such extensions.)

Historical implementations are inconsistent in the format used for $-\mathbf{o}$ option status reporting. The $+\mathbf{o}$ format without an option-argument was added to allow portable access to the options that can be saved and then later restored using, for instance, a dot script.

Historically, sh did trace the command set + x, but ksh did not.

The *ignoreeof* setting prevents accidental logouts when the end-of-file character (typically <control>-D) is entered. A user shall explicitly *exit* to leave the interactive shell.

The set – \mathbf{m} option was added to apply only to the UPE because it applies primarily to interactive use, not shell script applications.

3507 The ability to do asynchronous notification became available in the 1988 version of the 3508 KornShell. To have it occur, the user had to issue the command: 3509 trap "jobs -n" CLD 3510 The C shell provides two different levels of an asynchronous notification capability. The 3511 environment variable *notify* is analogous to what is done in set -**b** or set -**o** notify. When set, it notifies the user immediately of background job completions. When unset, this capability is 3512 turned off. 3513 The other notification ability comes through the built-in utility *notify*. The syntax is: 3514 3515 notify [%job ...] By issuing *notify* with no operands, it causes the C shell to notify the user asynchronously when 3516 the state of the current job changes. If given operands, *notify* asynchronously informs the user of 3517 3518 changes in the states of the specified jobs. To add asynchronous notification to the POSIX shell, neither the KornShell extensions to trap, 3519 nor the C shell notify environment variable seemed appropriate (notify is not a proper POSIX 3520 environment variable name). 3521 The set –**b** option was selected as a compromise. 3522 The *notify* built-in was considered to have more functionality than was required for simple 3523 3524 asynchronous notification. 3525 **FUTURE DIRECTIONS** None. 3526 3527 **SEE ALSO** Section 2.14 (on page 64) 3528 3529 **CHANGE HISTORY** 3530 Issue 6 3531 The obsolescent *set* command name followed by '-' has been removed. 3532 The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification: 3533 3534 • The *nolog* option is added to $set - \mathbf{o}$. IEEE PASC Interpretation 1003.2 #167 is applied, clarifying that the options default also takes 3535 3536 into account the description of the option. IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page 3537 sections use terms as described in the Utility Description Defaults (Section 1.11). No change in 3538 behavior is intended. 3539 IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/8 is applied, changing the square brackets 3540

in the example in RATIONALE to be in bold, which is the typeface used for optional items.

shift Shell Command Language

3542 **NAME** shift — shift positional parameters 3543 3544 **SYNOPSIS** shift [n] 3545 **DESCRIPTION** 3546 The positional parameters shall be shifted. Positional parameter 1 shall be assigned the value of 3547 parameter (1+n), parameter 2 shall be assigned the value of parameter (2+n), and so on. The 3548 parameters represented by the numbers "\$#" down to "\$#-n+1" shall be unset, and the 3549 parameter '#' is updated to reflect the new number of positional parameters. 3550 3551 The value *n* shall be an unsigned decimal integer less than or equal to the value of the special parameter ' #'. If n is not given, it shall be assumed to be 1. If n is 0, the positional and special 3552 3553 parameters are not changed. **OPTIONS** 3554 None. 3555 **OPERANDS** 3556 See the DESCRIPTION. 3557 **STDIN** 3558 Not used. 3559 **INPUT FILES** 3560 3561 None. **ENVIRONMENT VARIABLES** 3562 None. 3563 ASYNCHRONOUS EVENTS 3564 3565 Default. **STDOUT** 3566 3567 Not used. **STDERR** 3568 3569 The standard error shall be used only for diagnostic messages. **OUTPUT FILES** 3570 None. **EXTENDED DESCRIPTION** 3572 3573 None. **EXIT STATUS** 3574 The exit status is >0 if n>\$#; otherwise, it is zero. 3575 **CONSEQUENCES OF ERRORS** 3576 Default. 3577

shift

3578	APPLICATION USAGE				
3579	None.				
3580	EXAMPLES				
3581	\$ set a b c d e				
3582	\$ shift 2				
3583	\$ echo \$*				
3584	c d e				
3585	RATIONALE				
3586	None.				
3587 3588					
3589	SEE ALSO				
3590	Section 2.14 (on page 64)				
3591	CHANGE HISTORY				
3592	Issue 6				
3593	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page				
3594	sections use terms as described in the Utility Description Defaults (Section 1.11). No change in				
3595	behavior is intended.				

times

```
3596
    NAME
             times — write process times
3597
3598
    SYNOPSIS
             times
3599
3600
    DESCRIPTION
             The times utility shall write the accumulated user and system times for the shell and for all of its
3601
             child processes, in the following POSIX locale format:
3602
             "%dm%fs %dm%fs\n%dm%fs %dm%fs\n", <shell user minutes>,
3603
                  <shell user seconds>, <shell system minutes>,
3604
                  <shell system seconds>, <children user minutes>,
3605
                  <children user seconds>, <children system minutes>,
3606
                  <children system seconds>
3607
             The four pairs of times shall correspond to the members of the <sys/times.h> tms structure
3608
3609
             (defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 13, Headers) as
             returned by times(): tms_utime, tms_stime, tms_cutime, and tms_cstime, respectively.
3610
    OPTIONS
3611
             None.
3612
    OPERANDS
3613
             None.
3614
    STDIN
3615
             Not used.
3616
    INPUT FILES
3617
             None.
3618
    ENVIRONMENT VARIABLES
3619
             None.
3620
    ASYNCHRONOUS EVENTS
3621
             Default.
3622
    STDOUT
3623
             See the DESCRIPTION.
3624
    STDERR
3625
             The standard error shall be used only for diagnostic messages.
3626
3627
    OUTPUT FILES
             None.
3628
    EXTENDED DESCRIPTION
3629
             None.
3630
    EXIT STATUS
3631
             Zero.
3632
    CONSEQUENCES OF ERRORS
3633
             Default.
3634
```

Shell Command Language times

3635	APPLICATION USAGE
3636	None.
3637	EXAMPLES
3638	\$ times
3639	0m0.43s 0m1.11s
3640	8m44.18s 1m43.23s
3641	RATIONALE
3642	The times special built-in from the Single UNIX Specification is now required for all conforming
3643	shells.
3644	FUTURE DIRECTIONS
3645	None.
3646	SEE ALSO
3647	Section 2.14 (on page 64)
3648	CHANGE HISTORY
3649	Issue 6
3650	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/9 is applied, changing text in the
3651	DESCRIPTION from: "Write the accumulated user and system times for the shell and for all of
3652	its child processes" to: "The times utility shall write the accumulated user and system times
2052	for the shell and for all of its shild processes.

```
3654
    NAME
             trap — trap signals
3655
3656
    SYNOPSIS
             trap [action condition ...]
3657
3658
```

DESCRIPTION

3659

3660

3661 3662

3663

3664

3665

3666

3667

3668

3669

3670

3671

3672

3674

3675

3676 3677

3679

3680 3681

3682

3683

3684

3688

XSI

If action is '-', the shell shall reset each condition to the default value. If action is null (""), the shell shall ignore each specified *condition* if it arises. Otherwise, the argument *action* shall be read and executed by the shell when one of the corresponding conditions arises. The action of trap shall override a previous action (either default action or one explicitly set). The value of "\$?" after the *trap* action completes shall be the value it had before *trap* was invoked.

The condition can be EXIT, 0 (equivalent to EXIT), or a signal specified using a symbolic name, without the SIG prefix, as listed in the tables of signal names in the <signal.h> header defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 13, Headers; for example, HUP, INT, QUIT, TERM. Implementations may permit names with the SIG prefix or ignore case in signal names as an extension. Setting a trap for SIGKILL or SIGSTOP produces undefined results.

The environment in which the shell executes a *trap* on EXIT shall be identical to the environment immediately after the last command executed before the trap on EXIT was taken.

Each time *trap* is invoked, the *action* argument shall be processed in a manner equivalent to:

```
3673
            eval action
```

Signals that were ignored on entry to a non-interactive shell cannot be trapped or reset, although no error need be reported when attempting to do so. An interactive shell may reset or catch signals ignored on entry. Traps shall remain in place for a given shell until explicitly changed with another *trap* command.

When a subshell is entered, traps that are not being ignored are set to the default actions. This does not imply that the *trap* command cannot be used within the subshell to set new traps.

The trap command with no arguments shall write to standard output a list of commands associated with each condition. The format shall be:

```
"trap -- %s %s ...\n", <action>, <condition> ...
```

The shell shall format the output, including the proper use of quoting, so that it is suitable for reinput to the shell as commands that achieve the same trapping results. For example:

```
save traps=$(trap)
3685
3686
3687
            eval "$save traps"
```

XSI-conformant systems also allow numeric signal numbers for the conditions corresponding to the following signal names:

```
3689
                  SIGHUP
3690
              2
                  SIGINT
3691
              3
                  SIGQUIT
3692
              6
                  SIGABRT
3693
              9
                  SIGKILL
3694
                  SIGALRM
3695
```

trap

```
3696
                 SIGTERM
3697
             The trap special built-in shall conform to the Base Definitions volume of IEEE Std 1003.1-2001,
3698
             Section 12.2, Utility Syntax Guidelines.
3699
     OPTIONS
3700
             None.
3701
     OPERANDS
3702
             See the DESCRIPTION.
3703
3704
     STDIN
             Not used.
3705
     INPUT FILES
3706
             None.
3707
    ENVIRONMENT VARIABLES
3708
             None.
3709
    ASYNCHRONOUS EVENTS
3710
             Default.
3711
    STDOUT
3712
             See the DESCRIPTION.
3713
     STDERR
3714
             The standard error shall be used only for diagnostic messages.
3715
     OUTPUT FILES
3716
             None.
     EXTENDED DESCRIPTION
3718
             None.
3719
    EXIT STATUS
3720
             If the trap name or number is invalid, a non-zero exit status shall be returned; otherwise, zero
3721
     XSI
             shall be returned. For both interactive and non-interactive shells, invalid signal names or
3722
             numbers shall not be considered a syntax error and do not cause the shell to abort.
3723
     CONSEQUENCES OF ERRORS
3724
             Default.
3725
     APPLICATION USAGE
3726
             None.
3727
     EXAMPLES
3728
             Write out a list of all traps and actions:
3729
3730
             trap
3731
             Set a trap so the logout utility in the directory referred to by the HOME environment variable
             executes when the shell terminates:
3732
3733
             trap '$HOME/logout' EXIT
3734
             or:
3735
             trap '$HOME/logout' 0
3736
             Unset traps on INT, QUIT, TERM, and EXIT:
```

```
3737
              trap - INT QUIT TERM EXIT
     RATIONALE
3738
3739
              Implementations may permit lowercase signal names as an extension. Implementations may
              also accept the names with the SIG prefix; no known historical shell does so. The trap and kill
3740
3741
              utilities in this volume of IEEE Std 1003.1-2001 are now consistent in their omission of the SIG
              prefix for signal names. Some kill implementations do not allow the prefix, and kill -l lists the
3742
              signals without prefixes.
3743
              Trapping SIGKILL or SIGSTOP is syntactically accepted by some historical implementations, but
3744
3745
              it has no effect. Portable POSIX applications cannot attempt to trap these signals.
              The output format is not historical practice. Since the output of historical trap commands is not
3746
              portable (because numeric signal values are not portable) and had to change to become so, an
3747
              opportunity was taken to format the output in a way that a shell script could use to save and
3748
3749
              then later reuse a trap if it wanted.
              The KornShell uses an ERR trap that is triggered whenever set –e would cause an exit. This is
3750
              allowable as an extension, but was not mandated, as other shells have not used it.
3751
3752
              The text about the environment for the EXIT trap invalidates the behavior of some historical
              versions of interactive shells which, for example, close the standard input before executing a
3753
              trap on 0. For example, in some historical interactive shell sessions the following trap on 0 would
3754
              always print "--":
3755
              trap 'read foo; echo "-$foo-"' 0
3756
     FUTURE DIRECTIONS
3757
              None.
3758
     SEE ALSO
3759
              Section 2.14 (on page 64)
3760
     CHANGE HISTORY
3761
3762
     Issue 6
              XSI-conforming implementations provide the mapping of signal names to numbers given above
3763
3764
              (previously this had been marked obsolescent). Other implementations need not provide this
3765
              optional mapping.
              IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page
3766
              sections use terms as described in the Utility Description Defaults (Section 1.11). No change in
3767
3768
              behavior is intended.
```

```
3769
     NAME
              unset — unset values and attributes of variables and functions
3770
3771
     SYNOPSIS
3772
              unset [-fv] name ...
3773
     DESCRIPTION
              Each variable or function specified by name shall be unset.
3774
3775
              If -v is specified, name refers to a variable name and the shell shall unset it and remove it from
              the environment. Read-only variables cannot be unset.
3776
3777
              If -f is specified, name refers to a function and the shell shall unset the function definition.
              If neither -f nor -v is specified, name refers to a variable; if a variable by that name does not
3778
              exist, it is unspecified whether a function by that name, if any, shall be unset.
3779
              Unsetting a variable or function that was not previously set shall not be considered an error and
3780
              does not cause the shell to abort.
3781
              The unset special built-in shall support the Base Definitions volume of IEEE Std 1003.1-2001,
3782
              Section 12.2, Utility Syntax Guidelines.
3783
              Note that:
3784
              VARIABLE=
3785
              is not equivalent to an unset of VARIABLE; in the example, VARIABLE is set to "". Also, the
3786
              variables that can be unset should not be misinterpreted to include the special parameters (see
3787
3788
              Section 2.5.2 (on page 34)).
     OPTIONS
3789
              See the DESCRIPTION.
3790
     OPERANDS
3791
              See the DESCRIPTION.
3792
     STDIN
3793
3794
              Not used.
     INPUT FILES
3795
              None.
3796
     ENVIRONMENT VARIABLES
3797
3798
              None.
     ASYNCHRONOUS EVENTS
3799
              Default.
3800
     STDOUT
3801
              Not used.
3802
     STDERR
3803
              The standard error shall be used only for diagnostic messages.
3804
     OUTPUT FILES
3805
              None.
3806
     EXTENDED DESCRIPTION
3807
              None.
3808
```

```
3809
     EXIT STATUS
               0 All name operands were successfully unset.
3810
             >0 At least one name could not be unset.
3811
     CONSEQUENCES OF ERRORS
3812
             Default.
3813
     APPLICATION USAGE
3814
             None.
3815
     EXAMPLES
3816
             Unset VISUAL variable:
3817
             unset -v VISUAL
3818
             Unset the functions foo and bar:
3819
3820
             unset -f foo bar
     RATIONALE
3821
             Consideration was given to omitting the -f option in favor of an unfunction utility, but the
3822
3823
             standard developers decided to retain historical practice.
             The -v option was introduced because System V historically used one name space for both
3824
             variables and functions. When unset is used without options, System V historically unset either a
3825
             function or a variable, and there was no confusion about which one was intended. A portable
3826
             POSIX application can use unset without an option to unset a variable, but not a function; the -f
3827
3828
             option must be used.
     FUTURE DIRECTIONS
3829
             None.
3830
     SEE ALSO
3831
             Section 2.14 (on page 64)
3832
     CHANGE HISTORY
3833
     Issue 6
3834
             IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/5 is applied so that the reference page
3835
             sections use terms as described in the Utility Description Defaults (Section 1.11). No change in
```

behavior is intended.

3839 BE

3840 3841

3842

3843

3845

3846

3847

3848

3849

3850

3851

3852 3853

3854

3855

3859

3860

3861 3862

3863

3865 3866

3867 3868

3869

3870

This chapter describes the services and utilities that shall be implemented on all systems that claim conformance to the Batch Environment Services and Utilities option. This functionality is dependent on support of this option (and the rest of this section is not further shaded for this option).

3.1 General Concepts

3844 3.1.1 Batch Client-Server Interaction

Batch jobs are created and managed by batch servers. A batch client interacts with a batch server to access batch services on behalf of the user. In order to use batch services, a user must have access to a batch client.

A batch server is a computational entity, such as a daemon process, that provides batch services. Batch servers route, queue, modify, and execute batch jobs on behalf of batch clients.

The batch utilities described in this volume of IEEE Std 1003.1-2001 (and listed in Table 3-1) are clients of batch services; they allow users to perform actions on the job such as creating, modifying, and deleting batch jobs from a shell command line. Although these batch utilities may be said to accomplish certain services, they actually obtain services on behalf of a user by means of requests to batch servers.

Table 3-1 Batch Utilities

3856	qalter	qmove	qrls	qstat
3857	qdel	qmsg	qselect	qsub
3858	qhold	qrerun	qsig	

Client-server interaction takes place by means of the batch requests defined in this chapter. Because direct access to batch jobs and queues is limited to batch servers, clients and servers of different implementations can interoperate, since dependencies on private structures for batch jobs and queues are limited to batch servers. Also, batch servers may be clients of other batch servers.

3864 3.1.2 Batch Queues

Two types of batch queue are described: routing queues and execution queues. When a batch job is placed in a routing queue, it is a candidate for routing. A batch job is removed from routing queues under the following conditions:

- The batch job has been routed to another queue.
- The batch job has been deleted from the batch queue.
- The batch job has been aborted.
- When a batch job is placed in an execution queue, it is a candidate for execution.
- A batch job is removed from an execution queue under the following conditions:

- The batch job has been executed and exited.
- The batch job has been aborted.
- The batch job has been deleted from the batch queue.
 - The batch job has been moved to another queue.

Access to a batch queue is limited to the batch server that manages the batch queue. Clients never access a batch queue or a batch job directly, either to read or write information; all client access to batch queues or jobs takes place through batch servers.

3880 3.1.3 Batch Job Creation

When a batch server creates a batch job on behalf of a client, it shall assign a batch job identifier to the job. A batch job identifier consists of both a sequence number that is unique among the sequence numbers issued by that server and the name of the server. Since the batch server name is unique within a name space, the job identifier is likewise unique within the name space.

The batch server that creates a batch job shall return the batch server-assigned job identifier to the client that requested the job creation. If the batch server routes or moves the job to another server, it sends the job identifier with the job. Once assigned, the job identifier of a batch job shall never change.

3889 3.1.4 Batch Job Tracking

Since a batch job may be moved after creation, the batch server name component of the job identifier need not indicate the location of the job. An implementation may provide a batch job tracking mechanism, in which case the user generally does not need to know the location of the job. However, an implementation need not provide a batch job tracking mechanism, in which case the user must find routed jobs by probing the possible destinations.

3895 3.1.5 Batch Job Routing

To route a batch job, a batch server either moves the job to some other queue that is managed by the batch server, or requests that some other batch server accept the job.

Each routing queue has one or more queues to which it can route batch jobs. The batch server administrator creates routing queues.

A batch server may route a batch job from a routing queue to another routing queue. Batch servers shall prevent or otherwise handle cases of circular routing paths. As a deferred service, a batch server routes jobs from the routing queues that it manages. The algorithm by which a batch server selects a batch queue to which to route a batch job is implementation-defined.

A batch job need not be eligible for routing to all the batch queues fed by the routing queue from which it is routed. A batch server that has been asked to accept the job may reject the request if the job requires resources that are unavailable to that batch server, or if the client is not authorized to access the batch server.

Batch servers may route high-priority jobs before low-priority jobs, but, on other than overloaded systems, the effect may be imperceptible to the user. If all the batch servers fed by a routing queue reject requests to accept the job for reasons that are permanent, the batch server that manages the job shall abort the job. If all or some rejections are temporary, the batch server should try to route the job again at some later point.

The reasons for rejecting a batch job are implementation-defined.

The reasons for which the routing should be retried later and the reasons for which the job should be aborted are also implementation-defined.

3916 **3.1.6 Batch Job Execution**

To execute a batch job is to create a session leader (a process) that runs the shell program indicated by the *Shell_Path* attribute of the job. The script shall be passed to the program as its standard input. An implementation may pass the script to the program by other implementation-defined means. At the time a batch job begins execution, it is defined to enter the RUNNING state. The primary program that is executed by a batch job is typically, though not necessarily, a shell program.

A batch server shall execute eligible jobs as a deferred service—no client request is necessary once the batch job is created and eligible. However, the attributes of a batch job, such as the job hold type, may render the job ineligible. A batch server shall scan the execution queues that it manages for jobs that are eligible for execution. The algorithm by which the batch server selects eligible jobs for execution is implementation-defined.

As part of creating the process for the batch job, the batch server shall open the standard output and standard error streams of the session.

The attributes of a batch job may indicate that the batch server executing the job shall send mail to a list of users at the time it begins execution of the job.

3932 3.1.7 Batch Job Exit

3923 3924

3925

3926

3927

3928

3929 3930

3931

3933

3934

3935 3936

3940

3941

3942

3944

3945

3946

3947

3948

3949

3950

3951

3952

3953

When the session leader of an executing job terminates, the job exits. As part of exiting a batch job, the batch server that manages the job shall remove the job from the batch queue in which it resides. The server shall transfer output files of the job to a location described by the attributes of the job.

The attributes of a batch job may indicate that the batch server managing the job shall send mail to a list of users at the time the job exits.

3939 3.1.8 Batch Job Abort

A batch server shall abort jobs for which a required deferred service cannot be performed. The attributes of a batch job may indicate that the batch server that aborts the job shall send mail to a list of users at the time it aborts the job.

3943 3.1.9 Batch Authorization

Clients, such as the batch environment utilities (marked BE), access batch services by means of requests to one or more batch servers. To acquire the services of any given batch server, the user identifier under which the client runs must be authorized to use that batch server.

The user with an associated user name that creates a batch job shall own the job and can perform actions such as read, modify, delete, and move.

A user identifier of the same value at a different host need not be the same user. For example, user name *smith* at host **alpha** may or may not represent the same person as user name *smith* at host **beta**. Likewise, the same person may have access to different user names on different hosts.

An implementation may optionally provide an authorization mechanism that permits one user name to access jobs under another user name.

A process on a client host may be authorized to run processes under multiple user names at a batch server host. Where appropriate, the utilities defined in this volume of IEEE Std 1003.1-2001

provide a means for a user to choose from among such user names when creating or modifying a batch job.

3958 3.1.10 Batch Administration

The processing of a batch job by a batch server is affected by the attributes of the job. The processing of a batch job may also be affected by the attributes of the batch queue in which the job resides and by the status of the batch server that manages the job. See also the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 3, Definitions for batch definitions.

3963 3.1.11 Batch Notification

3964

3965

3966

3968 3969

3970

3971

3972

3973

3974

3975

3976 3977

3978

Whereas batch servers are persistent entities, clients are often transient. For example, the *qsub* utility creates a batch job and exits. For this reason, batch servers notify users of batch job events by sending mail to the user that owns the job, or to other designated users.

3967 3.2 Batch Services

The presence of Batch Environment Services and Utilities option services is indicated by the configuration variable POSIX2_PBS. A conforming batch server provides services as defined in this section.

A batch server shall provide batch services in two ways:

- 1. The batch server provides a service at the request of a client.
- 2. The batch server provides a deferred service as a result of a change in conditions monitored by the batch server.

If a batch server cannot complete a request, it shall reject the request. If a batch server cannot complete a deferred service for a batch job, the batch server shall abort the batch job. Table 3-2 (on page 105) is a summary of environment variables that shall be supported by an implementation of the batch server and utilities.

Table 3-2 Environment Variable Summary

3980	Variable	Description
3981	PBS_DPREFIX	Defines the directive prefix (see <i>qsub</i>)
3982	PBS_ENVIRONMENT	Batch Job is batch or interactive (see Section 3.2.2.1)
3983	PBS_JOBID	The <i>job_identifier</i> attribute of job (see Section 3.2.3.8)
3984	PBS_JOBNAME	The <i>job_name</i> attribute of job (see Section 3.2.3.8)
3985	PBS_O_HOME	Defines the <i>HOME</i> of the batch client (see <i>qsub</i>)
3986	PBS_O_HOST	Defines the host name of the batch client (see <i>qsub</i>)
3987	PBS_O_LANG	Defines the <i>LANG</i> of the batch client (see <i>qsub</i>)
3988	PBS_O_LOGNAME	Defines the <i>LOGNAME</i> of the batch client (see <i>qsub</i>)
3989	PBS_O_MAIL	Defines the MAIL of the batch client (see qsub)
3990	PBS_O_PATH	Defines the <i>PATH</i> of the batch client (see <i>qsub</i>)
3991	PBS_O_QUEUE	Defines the submit queue of the batch client (see <i>qsub</i>)
3992	PBS_O_SHELL	Defines the SHELL of the batch client (see qsub)
3993	PBS_O_TZ	Defines the <i>TZ</i> of the batch client (see <i>qsub</i>)
3994	PBS_O_WORKDIR	Defines the working directory of the batch client (see <i>qsub</i>)
3995	PBS_QUEUE	Defines the initial execution queue (see Section 3.2.2.1)

3.2.1 Batch Job States

A batch job shall always be in one of the following states: QUEUED, RUNNING, HELD, WAITING, EXITING, or TRANSITING. The state of a batch job determines the types of requests that the batch server that manages the batch job can accept for the batch job. A batch server shall change the state of a batch job either in response to service requests from clients or as a result of deferred services, such as job execution or job routing.

A batch job that is in the QUEUED state resides in a queue but is still pending either execution or routing, depending on the queue type.

A batch server that queues a batch job in a routing queue shall put the batch job in the QUEUED state. A batch server that puts a batch job in an execution queue, but has not yet executed the batch job, shall put the batch job in the QUEUED state. A batch job that resides in an execution queue and is executing is defined to be in the RUNNING state. While a batch job is in the RUNNING state, a session leader is associated with the batch job.

A batch job that resides in an execution queue, but is ineligible to run because of a hold attribute, is defined to be in the HELD state.

A batch job that is not held, but must wait until a future date and time before executing, is defined to be in the WAITING state.

When the session leader associated with a running job exits, the batch job shall be placed in the EXITING state.

A batch job for which the session leader has terminated is defined to be in the EXITING state, and the batch server that manages such a batch job cannot accept job modification requests that affect the batch job. While a batch job is in the EXITING state, the batch server that manages the batch job is staging output files and notifying clients of job completion. Once a batch job has exited, it no longer exists as an object managed by a batch server.

A batch job that is being moved from a routing queue to another queue is defined to be in the TRANSITING state.

4026

4027

4028

4029

4030 4031

4032

4033 4034 4035

4036

4037

4038

4039

4042 4043

4044

4045

4046

4047

4048 4049

4050

4052

4053

When a batch job in a routing queue has been selected to be moved to a new destination, then the batch job shall be in either the QUEUED state or the TRANSITING state, depending on the batch server implementation.

Batch jobs with either an *Execution_Time* attribute value set in the future or a *Hold_Types* attribute of value not equal to NO_HOLD, or both, may be routed or held in the routing queue. The treatment of jobs with the *Execution_Time* or *Hold_Types* attributes in a routing queue is implementation-defined.

When a batch job in a routing queue has not been selected to be moved to a new destination and the batch job has a *Hold_Types* attribute value of other than NO_HOLD, then the job should be in the HELD state.

Note: The effect of a hold upon a batch job in a routing queue is implementation-defined. The implementation should use the state that matches whether the batch job can route with a hold or not.

When a batch job in a routing queue has not been selected to be moved to a new destination and the batch job has:

- A Hold_Types attribute value of NO_HOLD
- An Execution_Time attribute in the past

then the batch job shall be in the QUEUED state.

When a batch job in a routing queue has not been selected to be moved to a new destination and the batch job has:

- A *Hold Types* attribute value of NO HOLD
- An *Execution_Time* attribute in the future

then the batch job may be in the WAITING state.

Note: The effect of a future execution time upon a batch job in a routing queue is implementation-defined. The implementation should use the state that matches whether the batch job can route with a hold or not.

Table 3-3 (on page 107) describes the next state of a batch job, given the current state of the batch job and the type of request. Table 3-4 (on page 108) describes the response of a batch server to a request, given the current state of the batch job and the type of request.

4051 3.2.2 Deferred Batch Services

This section describes the deferred services performed by batch servers: job execution, job routing, job exit, job abort, and the rerunning of jobs after a restart.

4054 3.2.2.1 Batch Job Execution

To execute a batch job is to create a session leader (a process) that runs the shell program indicated by the *Shell_Path_List* attribute of the batch job. The script is passed to the program as its standard input. An implementation may pass the script to the program by other implementation-defined means. At the time a batch job begins execution, it is defined to enter the RUNNING state.

4060

Table 3-3 Next State Table

4061
4062
4063
4064
4065
4066
4067
4068
4069
4070
4071
4072
4073
4074
4075
4076

4078

4087 4088

4089

4090

4091

4092 4093

4094

4095

4096

4097 4098

4099

4100

4101

4102 4103

			Cı	urrent State			
Request Type	X	Q	R	Н	W	E	T
Queue Batch Job Request	Q	e	e	e	e	e	e
Modify Batch Job Request	e	Q	R	Н	W	e	T
Delete Batch Job Request	e	X	Е	X	X	E	X
Batch Job Message Request	e	Q	R	Н	W	E	T
Rerun Batch Job Request	e	e	Q	e	e	e	e
Signal Batch Job Request	e	e	R	Н	W	e	e
Batch Job Status Request	e	Q	R	Н	W	E	T
Batch Queue Status Request	X	Q	R	Н	W	E	T
Server Status Request	X	Q	R	Н	W	E	T
Select Batch Jobs Request	X	Q	R	Н	W	E	T
Move Batch Job Request	e	Q	R	Н	W	e	T
Hold Batch Job Request	e	Н	R/H	Н	Н	e	T
Release Batch Job Request	e	Q	R	Q/W/H	W	e	T
Server Shutdown Request	X	Q	Q	Н	W	Е	T
Locate Batch Job Request	e	Q	R	Н	W	Е	T

Legend

- 4079 X Nonexistent
- 4080 Q QUEUED
- 4081 R RUNNING
- 4082 H HELD
- 4083 W WAITING
- 4084 E EXITING
- 4085 T TRANSITING
- 4086 e Error

A batch server that has an execution queue containing jobs is said to own the queue and manage the batch jobs in that queue. A batch server that has been started shall execute the batch jobs in the execution queues owned by the batch server. The batch server shall schedule for execution those jobs in the execution queues that are in the QUEUED state. The algorithm for scheduling jobs is implementation-defined.

A batch server that executes a batch job shall create, in the environment of the session leader of the batch job, an environment variable named *PBS_ENVIRONMENT*, the value of which is the string PBS_BATCH encoded in the portable character set.

A batch server that executes a batch job shall create, in the environment of the session leader of the batch job, an environment variable named *PBS_QUEUE*, the value of which is the name of the execution queue of the batch job encoded in the portable character set.

To rerun a batch job is to requeue a batch job that is currently executing and then kill the session leader of the executing job by sending a SIGKILL prior to completion; see Section 3.2.3.11 (on page 120). A batch server that reruns a batch job shall append the standard output and standard error files of the batch job to the corresponding files of the previous execution, if they exist, with appropriate annotation. If either file does not exist, that file shall be created as in normal execution.

4	1	04

Table 3-4 Results/Output Table

4105
4106
4107
4108
4109
4110
4111
4112
4113
4114
4115
4116
4117
4118
4119
4120

4127 4128

4129

4130

4131

4132 4133

4134

4135

4136 4137

4138

4139

4140

4141 4142

4143

4144

4145 4146

			Cur	rent S	State		
Request Type	X	Q	R	Н	W	E	T
Queue Batch Job Request	О	e	e	e	e	e	e
Modify Batch Job Request	e	О	e	О	0	e	e
Delete Batch Job Request	e	О	О	О	0	e	О
Batch Job Message Request	e	e	О	e	e	e	e
Rerun Batch Job Request	e	e	О	e	e	e	e
Signal Batch Job Request	e	e	О	e	e	e	e
Batch Job Status Request	e	О	О	О	0	О	О
Batch Queue Status Request	О	О	Ο	О	0	О	О
Server Status Request	О	О	О	О	0	О	О
Select Batch Job Request	e	О	О	О	O	О	О
Move Batch Job Request	e	О	О	О	О	e	e
Hold Batch Job Request	e	О	Ο	О	O	e	e
Release Batch Job Request	e	О	e	О	O	e	e
Server Shutdown Request	О	О	e	О	0	e	e
Locate Batch Job Request	e	0	О	0	0	О	О

4122 Legend

- 4123 O OK
- 4124 e Error message

The execution of a batch job by a batch server shall be controlled by job, queue, and server attributes, as defined in this section.

Account Name Attribute

Batch accounting is an optional feature of batch servers. If a batch server implements accounting, the statements in this section apply and the configuration variable POSIX2_PBS_ACCOUNTING shall be set to 1.

A batch server that executes a batch job shall charge the account named in the *Account_Name* attribute of the batch job for resources consumed by the batch job.

If the *Account_Name* attribute of the batch job is absent from the batch job attribute list or is altered while the batch job is in execution, the batch server action is implementation-defined.

Checkpoint Attribute

Batch checkpointing is an optional feature of batch servers. If a batch server implements checkpointing, the statements in this section apply and the configuration variable POSIX2_PBS_CHECKPOINT shall be set to 1.

There are two attributes associated with the checkpointing feature: *Checkpoint* and *Minimum_Cpu_Interval*. *Checkpoint* is a batch job attribute, while *Minimum_Cpu_Interval* is a queue attribute. An implementation that does not support checkpointing shall support the *Checkpoint* job attribute to the extent that the batch server shall maintain and pass this attribute to other servers.

The behavior of a batch server that executes a batch job for which the value of the *Checkpoint* attribute is CHECKPOINT_UNSPECIFIED is implementation-defined. A batch server that executes a batch job for which the value of the *Checkpoint* attribute is NO_CHECKPOINT shall

4147 not checkpoint the batch job.

A batch server that executes a batch job for which the value of the *Checkpoint* attribute is CHECKPOINT_AT_SHUTDOWN shall checkpoint the batch job only when the batch server accepts a request to shut down during the time when the batch job is in the RUNNING state.

A batch server that executes a batch job for which the value of the *Checkpoint* attribute is CHECKPOINT_AT_MIN_CPU_INTERVAL shall checkpoint the batch job at the interval specified by the *Minimum_Cpu_Interval* attribute of the queue for which the batch job has been selected. The *Minimum_Cpu_Interval* attribute shall be specified in units of CPU minutes.

A batch server that executes a batch job for which the value of the *Checkpoint* attribute is an unsigned integer shall checkpoint the batch job at an interval that is the value of either the *Checkpoint* attribute, or the *Minimum_Cpu_Interval* attribute of the queue for which the batch job has been selected, whichever is greater. Both intervals shall be in units of CPU minutes. When the *Minimum_Cpu_Interval* attribute is greater than the *Checkpoint* attribute, the batch job shall write a warning message to the standard error stream of the batch job.

Error_Path Attribute

The *Error_Path* attribute of a running job cannot be changed by a *Modify Batch Job Request*. When the *Join_Path* attribute of the batch job is set to the value FALSE and the *Keep_Files* attribute of the batch job does not contain the value KEEP_STD_ERROR, a batch server that executes a batch job shall perform one of the following actions:

- Set the standard error stream of the session leader of the batch job to the path described by the value of the *Error_Path* attribute of the batch job.
- Buffer the standard error of the session leader of the batch job until completion of the batch job, and when the batch job exits return the contents to the destination described by the value of the *Error_Path* attribute of the batch job.

Applications shall not rely on having access to the standard error of a batch job prior to the completion of the batch job.

When the *Error_Path* attribute does not specify a host name, then the batch server shall retain the standard error of the batch job on the host of execution.

When the *Error_Path* attribute does specify a host name and the *Keep_Files* attribute does not contain the value KEEP_STD_ERROR, then the final destination of the standard error of the batch job shall be on the host whose host name is specified.

If the path indicated by the value of the *Error_Path* attribute of the batch job is a relative path, the batch server shall expand the path relative to the home directory of the user on the host to which the file is being returned.

When the batch server buffers the standard error of the batch job and the file cannot be opened for write upon completion of the batch job, then the server shall place the standard error in an implementation-defined location and notify the user of the location via mail. It shall be possible for the user to process this mail using the *mailx* utility.

If a batch server that does not buffer the standard error cannot open the standard error path of the batch job for write access, then the batch server shall abort the batch job.

Execution_Time Attribute

A batch server shall not execute a batch job before the time represented by the value of the *Execution_Time* attribute of the batch job. The *Execution_Time* attribute is defined in seconds since the Epoch.

Hold_Types Attribute

A batch server shall support the following hold types:

- s Can be set or released by a user with at least a privilege level of batch administrator (SYSTEM).
- Can be set or released by a user with at least a privilege level of batch operator (OPERATOR).
- u Can be set or released by the user with at least a privilege level of user, where the user is defined in the *Job_Owner* attribute (USER).
- n Indicates that none of the *Hold_Types* attributes are set (NO_HOLD).

An implementation may define other hold types. Any additional hold types, how they are specified, their internal representation, their behavior, and how they affect the behavior of other utilities are implementation-defined.

The value of the *Hold_Types* attribute shall be the union of the valid hold types ('s', 'o', 'u', and any implementation-defined hold types), or 'n'.

A batch server shall not execute a batch job if the *Hold_Types* attribute of the batch job has a value other than NO_HOLD. If the *Hold_Types* attribute of the batch job has a value other than NO_HOLD, the batch job shall be in the HELD state.

Job_Owner Attribute

The *Job_Owner* attribute consists of a pair of user name and host name values of the form:

username@hostname

A batch server that accepts a *Queue Batch Job Request* shall set the *Job_Owner* attribute to a string that is the *username@hostname* of the user who submitted the job.

Join_Path Attribute

A batch server that executes a batch job for which the value of the *Join_Path* attribute is TRUE shall ignore the value of the *Error_Path* attribute and merge the standard error of the batch job with the standard output of the batch job.

Keep_Files Attribute

A batch server that executes a batch job for which the value of the *Keep_Files* attribute includes the value KEEP_STD_OUTPUT shall retain the standard output of the batch job on the host where execution occurs. The standard output shall be retained in the home directory of the user under whose user ID the batch job is executed and the filename shall be the default filename for the standard output as defined under the **–o** option of the *qsub* utility. The *Output_Path* attribute is not modified.

A batch server that executes a batch job for which the value of the *Keep_Files* attribute includes the value KEEP_STD_ERROR shall retain the standard error of the batch job on the host where execution occurs. The standard error shall be retained in the home directory of the user under whose user ID the batch job is executed and the filename shall be the default filename for

4231 4232

4233

4234 4235

4236

4237

4238

4239

4240

4241

4242 4243

4245

4246

4248 4249

4250 4251

4252

4254

4255

4256

4257 4258

4259

4260

standard error as defined under the **–e** option of the *qsub* utility. The *Error_Path* attribute is not modified.

A batch server that executes a batch job for which the value of the *Keep_Files* attribute includes values other than KEEP_STD_OUTPUT and KEEP_STD_ERROR shall retain these other files on the host where execution occurs. These files (with implementation-defined names) shall be retained in the home directory of the user under whose user identifier the batch job is executed.

Mail_Points and Mail_Users Attributes

A batch server that executes a batch job for which one of the values of the *Mail_Points* attribute is the value MAIL_AT_BEGINNING shall send a mail message to each user account listed in the *Mail_Users* attribute of the batch job.

The mail message shall contain at least the batch job identifier, queue, and server at which the batch job currently resides, and the *Job_Owner* attribute.

Output_Path Attribute

The *Output_Path* attribute of a running job cannot be changed by a *Modify Batch Job Request*. When the *Keep_Files* attribute of the batch job does not contain the value KEEP_STD_OUTPUT, a batch server that executes a batch job shall either:

 Set the standard output stream of the session leader of the batch job to the destination described by the value of the Output_Path attribute of the batch job.

or:

- Buffer the standard output of the session leader of the batch job until completion of the batch job, and when the batch job exits return the contents to the destination described by the value of the *Output_Path* attribute of the batch job.
- When the *Output_Path* attribute does not specify a host name, then the batch server shall retain the standard output of the batch job on the host of execution.
- When the *Keep_Files* attribute does not contain the value KEEP_STD_OUTPUT and the *Output_Path* attribute does specify a host name, then the final destination of the standard output of the batch job shall be on the host specified.
- If the path specified in the *Output_Path* attribute of the batch job is a relative path, the batch server shall expand the path relative to the home directory of the user on the host to which the file is being returned.
- Whether or not the batch server buffers the standard output of the batch job until completion of the batch job is implementation-defined. Applications shall not rely on having access to the standard output of a batch job prior to the completion of the batch job.
- When the batch server does buffer the standard output of the batch job and the file cannot be opened for write upon completion of the batch job, then the batch server shall place the standard output in an implementation-defined location and notify the user of the location via mail. It shall be possible for the user to process this mail using the *mailx* utility.
- If a batch server that does not buffer the standard output cannot open the standard output path of the batch job for write access, then the batch server shall abort the batch job.

Priority Attribute

A batch server implementation may choose to preferentially execute a batch job based on the *Priority* attribute. The interpretation of the batch job *Priority* attribute by a batch server is implementation-defined. If an implementation uses the *Priority* attribute, it shall interpret larger values of the *Priority* attribute to mean the batch job shall be preferentially selected for execution.

Rerunable Attribute

A batch job that began execution but did not complete, because the batch server either shut down or terminated abnormally, shall be requeued if the *Rerunable* attribute of the batch job has the value TRUE.

If a batch job, which was requeued after beginning execution but prior to completion, has a valid checkpoint file and the batch server supports checkpointing, then the batch job shall be restarted from the last valid checkpoint.

If the batch job cannot be restarted from a checkpoint, then when a batch job has a *Rerunable* attribute value of TRUE and was requeued after beginning execution but prior to completion, the batch server shall place the batch job into execution at the beginning of the job.

When a batch job has a *Rerunable* attribute value other than TRUE and was requeued after beginning execution but prior to completion, and the batch job cannot be restarted from a checkpoint, then the batch server shall abort the batch job.

Resource_List Attribute

A batch server that executes a batch job shall establish the resource limits of the session leader of the batch job according to the values of the *Resource_List* attribute of the batch job. Resource limits shall be enforced by an implementation-defined method.

Shell_Path_List Attribute

The *Shell_Path_List* job attribute consists of a list of pairs of pathname and host name values. The host name component can be omitted, in which case the pathname serves as the default pathname when a batch server cannot find the name of the host on which it is running in the list.

A batch server that executes a batch job shall select, from the value of the *Shell_Path_List* attribute of the batch job, a pathname where the shell to execute the batch job shall be found. The batch server shall select the pathname, in order of preference, according to the following methods:

- Select the pathname that contains the name of the host on which the batch server is running.
- Select the pathname for which the host name has been omitted.
- Select the pathname for the login shell of the user under which the batch job is to execute.

If the shell path value selected is an invalid pathname, the batch server shall abort the batch job.

If the value of the selected pathname from the *Shell_Path_List* attribute of the batch job represents a partial path, the batch server shall expand the path relative to a path that is implementation-defined.

The batch server that executes the batch job shall execute the program that was selected from the *Shell_Path_List* attribute of the batch job. The batch server shall pass the path to the script of the batch job as the first argument to the shell program.

4307 User_List Attribute

The *User_List* job attribute consists of a list of pairs of user name and host name values. The host name component can be omitted, in which case the user name serves as a default when a batch server cannot find the name of the host on which it is running in the list.

A batch server that executes a batch job shall select, from the value of the *User_List* attribute of the batch job, a user name under which to create the session leader. The server shall select the user name, in order of preference, according to the following methods:

- Select the user name of a value that contains the name of the host on which the batch server
 executes.
- Select the user name of a value for which the host name has been omitted.
- Select the user name from the *Job_Owner* attribute of the batch job.

Variable_List Attribute

A batch server that executes a batch job shall create, in the environment of the session leader of the batch job, each environment variable listed in the *Variable_List* attribute of the batch job, and set the value of each such environment variable to that of the corresponding variable in the variable list.

3.2.2.2 Batch Job Routing

To route a batch job is to select a queue from a list and move the batch job to that queue.

A batch server that has routing queues, which have been started, shall route the jobs in the routing queues owned by the batch server. A batch server may delay the routing of a batch job. The algorithm for selecting a batch job and the queue to which it will be routed is implementation-defined.

When a routing queue has multiple possible destinations specified, then the precedence of the destinations is implementation-defined.

A batch server that routes a batch job to a queue at another server shall move the batch job into the target queue with a *Queue Batch Job Request*.

If the target server rejects the *Queue Batch Job Request*, the routing server shall retry routing the batch job or abort the batch job. A batch server that retries failed routings shall provide a means for the batch administrator to specify the number of retries and the minimum period of time between retries. The means by which an administrator specifies the number of retries and the delay between retries is implementation-defined. When the number of retries specified by the batch administrator has been exhausted, the batch server shall abort the batch job and perform the functions of *Batch Job Exit*; see Section 3.2.2.3.

4340 3.2.2.3 Batch Job Exit

For each job in the EXITING state, the batch server that exited the batch job shall perform the following deferred services in the order specified:

- 1. If buffering standard error, move that file into the location specified by the *Error_Path* attribute of the batch job.
- 2. If buffering standard output, move that file into the location specified by the *Output_Path* attribute of the batch job.
- 3. If the *Mail_Points* attribute of the batch job includes MAIL_AT_EXIT, send mail to the users listed in the *Mail_Users* attribute of the batch job. The mail message shall contain at least

4352

4353

4354

4355 4356

4357

4358

4359

4360

4361

4362

4363

4364

4365

4373

4374

4376 4377

4378

4379

4380

4382 4383

4384

4385

the batch job identifier, queue, and server at which the batch job currently resides, and the Job_Owner attribute.

4. Remove the batch job from the queue.

If a batch server that buffers the standard error output cannot return the standard error file to the standard error path at the time the batch job exits, the batch server shall do one of the following:

- Mail the standard error file to the batch job owner.
- Save the standard error file and mail the location and name of the file where the standard error is stored to the batch job owner.
 - Save the standard error file and notify the user by other implementation-defined means.

If a batch server that buffers the standard output cannot return the standard output file to the standard output path at the time the batch job exits, the batch server shall do one of the following:

- Mail the standard output file to the batch job owner.
- Save the standard output file and mail the location and name of the file where the standard output is stored to the batch job owner.
- Save the standard output file and notify the user by other implementation-defined means.
- 4366 At the conclusion of job exit processing, the batch job is no longer managed by a batch server.

4367 3.2.2.4 Batch Server Restart

A batch server that has been either shutdown or terminated abnormally, and has returned to operation, is said to have "restarted".

Upon restarting, a batch server shall requeue those jobs managed by the batch server that were in the RUNNING state at the time the batch server shut down and for which the *Rerunable* attribute of the batch job has the value TRUE.

Queues are defined to be non-volatile. A batch server shall store the content of queues that it controls in such a way that server and system shutdowns do not erase the content of the queues.

4375 3.2.2.5 Batch Job Abort

A batch server that cannot perform a deferred service for a batch job shall abort the batch job.

A batch server that aborts a batch job shall perform the following services:

- Delete the batch job from the queue in which it resides.
- If the *Mail_Points* attribute of the batch job includes the value MAIL_AT_ABORT, send mail to the users listed in the value of the *Mail_Users* attribute of the job. The mail message shall contain at least the batch job identifier, queue, and server at which the batch job currently resides, the *Job_Owner* attribute, and the reason for the abort.
- If the batch job was in the RUNNING state, terminate the session leader of the executing job by sending the session leader a SIGKILL, place the batch job in the EXITING state, and perform the actions of *Batch Job Exit*.

3.2.3 Requested Batch Services

This section describes the services provided by batch servers in response to requests from clients. Table 3-5 summarizes the current set of batch service requests and for each gives its type (deferred or not) and whether it is an optional function.

Table 3-5 Batch Services Summary

Batch Service	Deferred	Optional
Batch Job Execution	Yes	No
Batch Job Routing	Yes	No
Batch Job Exit	Yes	No
Batch Server Restart	Yes	No
Batch Job Abort	Yes	No
Delete Batch Job Request	No	No
Hold Batch Job Request	No	No
Batch Job Message Request	No	Yes
Batch Job Status Request	No	No
Locate Batch Job Request	No	Yes
Modify Batch Job Request	No	No
Move Batch Job Request	No	No
Queue Batch Job Request	No	No
Batch Queue Status Request	No	No
Release Batch Job Request	No	No
Rerun Batch Job Request	No	No
Select Batch Jobs Request	No	No
Server Shutdown Request	No	No
Server Status Request	No	No
Signal Batch Job Request	No	No
Track Batch Job Request	No	Yes

If a request is rejected because the batch client is not authorized to perform the action, the batch server shall return the same status as when the batch job does not exist.

3.2.3.1 Delete Batch Job Request

A batch job is defined to have been deleted when it has been removed from the queue in which it resides and not instantiated in another queue. A client requests that the server that manages a batch job delete the batch job. Such a request is called a *Delete Batch Job Request*.

A batch server shall reject a *Delete Batch Job Request* if any of the following statements are true:

- The user of the batch client is not authorized to delete the designated job.
- The designated job is not managed by the batch server.
- The designated job is in a state inconsistent with the delete request.

A batch server may reject a *Delete Batch Job Request* for other implementation-defined reasons. The method used to determine whether the user of a client is authorized to perform the requested action is implementation-defined.

A batch server requested to delete a batch job shall delete the batch job if the batch job exists and is not in the EXITING state.

A batch server that deletes a batch job in the RUNNING state shall send a SIGKILL signal to the session leader of the batch job. It is implementation-defined whether additional signals are sent

- 4430 to the session leader of the job prior to sending the SIGKILL signal.
- A batch server that deletes a batch job in the RUNNING state shall place the batch job in the
- 4432 EXITING state after it has killed the session leader of the batch job and shall perform the actions
- 4433 of Batch Job Exit.

4442

4443

4461

4463

4434 3.2.3.2 Hold Batch Job Request

A batch client can request that the batch server add one or more holds to a batch job. Such a request is called a *Hold Batch Job Request*.

- 4437 A batch server shall reject a *Hold Batch Job Request* if any of the following statements are true:
- The batch server does not support one or more of the requested holds to be added to the batch job.
 - The user of the batch client is not authorized to add one or more of the requested holds to the batch job.
 - The batch server does not manage the specified job.
 - The designated job is in the EXITING state.
- A batch server may reject a *Hold Batch Job Request* for other implementation-defined reasons. The method used to determine whether the user of a client is authorized to perform the requested action is implementation-defined.
- A batch server that accepts a *Hold Batch Job Request* for a batch job in the RUNNING state shall place a hold on the batch job. The effects, if any, the hold will have on a batch job in the RUNNING state are implementation-defined.
- A batch server that accepts a *Hold Batch Job Request* shall add each type of hold listed in the *Hold Batch Job Request*, that is not already present, to the value of the *Hold_Types* attribute of the batch job.

4453 3.2.3.3 Batch Job Message Request

- Batch Job Message Request is an optional feature of batch servers. If an implementation supports
 Batch Job Message Request, the statements in this section apply and the configuration variable
 POSIX2_PBS_MESSAGE shall be set to 1.
- A batch client can request that a batch server write a message into certain output files of a batch job. Such a request is called a *Batch Job Message Request*.
- 4459 A batch server shall reject a *Batch Job Message Request* if any of the following statements are true:
- The batch server does not support sending messages to jobs.
 - The user of the batch client is not authorized to post a message to the designated job.
- The designated job does not exist on the batch server.
 - The designated job is not in the RUNNING state.
- A batch server may reject a *Batch Job Message Request* for other implementation-defined reasons.
 The method used to determine whether the user of a client is authorized to perform the requested action is implementation-defined.
- A batch server that accepts a *Batch Job Message Request* shall write the message sent by the batch client into the files indicated by the batch client.

4469 3.2.3.4 Batch Job Status Request

- A batch client can request that a batch server respond with the status and attributes of a batch job. Such a request is called a *Batch Job Status Request*.
- 4472 A batch server shall reject a *Batch Job Status Request* if any of the following statements are true:
- The user of the batch client is not authorized to query the status of the designated job.
- The designated job is not managed by the batch server.
- A batch server may reject a *Batch Job Status Request* for other implementation-defined reasons.

 The method used to determine whether the user of a client is authorized to perform the requested action is implementation-defined.
- A batch server that accepts a *Batch Job Status Request* shall return a *Batch Job Status Message* to the batch client.
- A batch server may return other information in response to a *Batch Job Status Request*.

4481 3.2.3.5 Locate Batch Job Request

- Locate Batch Job Request is an optional feature of batch servers. If an implementation supports
 Locate Batch Job Request, the statements in this section apply and the configuration variable
 POSIX2 PBS LOCATE shall be set to 1.
- A batch client can ask a batch server to respond with the location of a batch job that was created by the batch server. Such a request is called a *Locate Batch Job Request*.
- A batch server that accepts a *Locate Batch Job Request* shall return a *Batch Job Location Message* to the batch client.
- A batch server may reject a *Locate Batch Job Request* for a batch job that was not created by that server.
- A batch server may reject a *Locate Batch Job Request* for a batch job that is no longer managed by that server; that is, for a batch job that is not in a queue owned by that server.
- 4493 A batch server may reject a *Locate Batch Job Request* for other implementation-defined reasons.

4494 3.2.3.6 Modify Batch Job Request

4500

- Batch clients modify (alter) the attributes of a batch job by making a request to the server that manages the batch job. Such a request is called a *Modify Batch Job Request*.
- 4497 A batch server shall reject a *Modify Batch Job Request* if any of the following statements are true:
- The user of the batch client is not authorized to make the requested modification to the batch job.
 - The designated job is not managed by the batch server.
- The requested modification is inconsistent with the state of the batch job.
- An unrecognized resource is requested for a batch job in an execution queue.
- A batch server may reject a *Modify Batch Job Request* for other implementation-defined reasons.

 The method used to determine whether the user of a client is authorized to perform the requested action is implementation-defined.
- A batch server that accepts a *Modify Batch Job Request* shall modify all the specified attributes of the batch job. A batch server that rejects a *Modify Batch Job Request* shall modify none of the attributes of the batch job.

4519

4526

4531

4533

4534

4536 4537

4538

4539

4540

4541

4542

4543

4545

4547 4548

If the servicing by a batch server of an otherwise valid request would result in no change, then the batch server shall indicate successful completion of the request.

4511 3.2.3.7 Move Batch Job Request

- A batch client can request that a batch server move a batch job to another destination. Such a request is called a *Move Batch Job Request*.
- 4514 A batch server shall reject a *Move Batch Job Request* if any of the following statements are true:
- The user of the batch client is not authorized to remove the designated job from the queue in which the batch job resides.
 - The user of the batch client is not authorized to move the designated job to the destination.
 - The designated job is not managed by the batch server.
 - The designated job is in the EXITING state.
- The destination is inaccessible.
- A batch server can reject a *Move Batch Job Request* for other implementation-defined reasons. The method used to determine whether the user of a client is authorized to perform the requested action is implementation-defined.
- A batch server that accepts a *Move Batch Job Request* shall perform the following services:
- Queue the designated job at the destination.
 - Remove the designated job from the queue in which the batch job resides.
- If the destination resides on another batch server, the batch server shall queue the batch job at the destination by sending a *Queue Batch Job Request* to the other server. If the *Queue Batch Job Request* fails, the batch server shall reject the *Move Batch Job Request*. If the *Queue Batch Job Request* succeeds, the batch server shall remove the batch job from its queue.
 - The batch server shall not modify any attributes of the batch job.

4532 3.2.3.8 Queue Batch Job Request

A batch queue is controlled by one and only one batch server. A batch server is said to own the queues that it controls. Batch clients make requests of batch servers to have jobs queued. Such a request is called a *Queue Batch Job Request*.

A batch server requested to queue a batch job for which the queue is not specified shall select an implementation-defined queue for the batch job. Such a queue is called the "default queue" of the batch server. The implementation shall provide the means for a batch administrator to specify the default queue. The queue, whether specified or defaulted, is called the "target queue".

A batch server shall reject a *Queue Batch Job Request* if any of the following statements are true:

- The client is not authorized to create a batch job in the target queue.
- The request specifies a queue that does not exist on the batch server.
- The target queue is an execution queue and the batch server cannot satisfy a resource requirement of the batch job.
- The target queue is an execution queue and an unrecognized resource is requested.
 - The target queue is an execution queue, the batch server does not support checkpointing, and the value of the *Checkpoint* attribute of the batch job is not NO_CHECKPOINT.

- The job requires access to a user identifier that the batch client is not authorized to access.
- 4550 A batch server may reject a *Queue Batch Job Request* for other implementation-defined reasons.
- A batch server that accepts a *Queue Batch Job Request* for a batch job for which the PBS_O_QUEUE value is missing from the value of the *Variable_List* attribute of the batch job shall add that variable to the list and set the value to the name of the target queue. Once set, no server shall change the value of PBS_O_QUEUE, even if the batch job is moved to another queue.
 - A batch server that accepts a *Queue Batch Job Request* for a batch job for which the PBS_JOBID value is missing from the value of the *Variable_List* attribute shall add that variable to the list and set the value to the batch job identifier assigned by the server in the format:
- 4559 sequence_number.server

A batch server that accepts a *Queue Batch Job Request* for a batch job for which the PBS_JOBNAME value is missing from the value of the *Variable_List* attribute of the batch job shall add that variable to the list and set the value to the *Job_Name* attribute of the batch job.

4563 3.2.3.9 Batch Queue Status Request

4556

4558

4568

4575

4576 4577

4578

4579

4580

A batch client can request that a batch server respond with the status and attributes of a queue.

Such a request is called a *Batch Queue Status Request*.

- 4566 A batch server shall reject a *Batch Queue Status Request* if any of the following statements are true:
- The user of the batch client is not authorized to query the status of the designated queue.
 - The designated queue does not exist on the batch server.
- A batch server may reject a *Batch Queue Status Request* for other implementation-defined reasons.

 The method used to determine whether the user of a client is authorized to perform the requested action is implementation-defined.
- A batch server that accepts a *Batch Queue Status Request* shall return a *Batch Queue Status Reply* to the batch client.

4574 3.2.3.10 Release Batch Job Request

A batch client can request that the server remove one or more holds from a batch job. Such a request is called a *Release Batch Job Request*.

- A batch server shall reject a *Release Batch Job Request* if any of the following statements are true:
- The user of the batch client is not authorized to remove one or more of the requested holds from the batch job.
 - The batch server does not manage the specified job.
- A batch server may reject a *Release Batch Job Request* for other implementation-defined reasons.

 The method used to determine whether the user of a client is authorized to perform the requested action is implementation-defined.
- A batch server that accepts a *Release Batch Job Request* shall remove each type of hold listed in the *Release Batch Job Request*, that is present, from the value of the *Hold_Types* attribute of the batch job.

4593

4595

4601

4602

4603

4587 3.2.3.11 Rerun Batch Job Request

To rerun a batch job is to kill the session leader of the batch job and leave the batch job eligible for re-execution. A batch client can request that a batch server rerun a batch job. Such a request is called *Rerun Batch Job Request*.

4591 A batch server shall reject a *Rerun Batch Job Request* if any of the following statements are true:

- The user of the batch client is not authorized to rerun the designated job.
- The *Rerunable* attribute of the designated job has the value FALSE.
- The designated job is not in the RUNNING state.
- The batch server does not manage the designated job.

A batch server may reject a *Rerun Batch Job Request* for other implementation-defined reasons.

The method used to determine whether the user of a client is authorized to perform the requested action is implementation-defined.

A batch server that rejects a *Rerun Batch Job Request* shall in no way modify the execution of the batch job.

A batch server that accepts a request to rerun a batch job shall perform the following services:

- Requeue the batch job in the execution queue in which it was executing.
- Send a SIGKILL signal to the process group of the session leader of the batch job.

An implementation may indicate to the batch job owner that the batch job has been rerun.
Whether and how the batch job owner is notified that a batch job is rerun is implementationdefined.

A batch server that reruns a batch job may send other implementation-defined signals to the session leader of the batch job prior to sending the SIGKILL signal.

A batch server may preferentially select a rerun job for execution. Whether rerun jobs shall be selected for execution before other jobs is implementation-defined.

4611 3.2.3.12 Select Batch Jobs Request

A batch client can request from a batch server a list of jobs managed by that server that match a list of selection criteria. Such a request is called a *Select Batch Jobs Request*. All the batch jobs managed by the batch server that receives the request are candidates for selection.

A batch server that accepts a *Select Batch Jobs Request* shall return a list of zero or more job identifiers that correspond to jobs that meet the selection criteria.

If the batch client is not authorized to query the status of a batch job, the batch server shall not select the batch job.

4619 3.2.3.13 Server Shutdown Request

A batch server is defined to have shut down when it does not respond to requests from clients and does not perform deferred services for jobs. A batch client can request that a batch server shut down. Such a request is called a *Server Shutdown Request*.

A batch server shall reject a *Server Shutdown Request* from a client that is not authorized to shut down the batch server. The method used to determine whether the user of a client is authorized to perform the requested action is implementation-defined.

- A batch server may reject a *Server Shutdown Request* for other implementation-defined reasons.

 The reasons for which a *Server Shutdown Request* may be rejected are implementation-defined.
- At server shutdown, a batch server shall do, in order of preference, one of the following:
- If checkpointing is implemented and the batch job is checkpointable, then checkpoint the batch job and requeue it.
 - If the batch job is rerunnable, then requeue the batch job to be rerun (restarted from the beginning).
- 4633
 Abort the batch job.

4634 3.2.3.14 Server Status Request

4631

4632

4638

4645

4646

4647

4650

- A batch client can request that a batch server respond with the status and attributes of the batch server. Such a request is called a *Server Status Request*.
- 4637 A batch server shall reject a *Server Status Request* if the following statement is true:
 - The user of the batch client is not authorized to query the status of the designated server.
- A batch server may reject a *Server Status Request* for other implementation-defined reasons. The method used to determine whether the user of a client is authorized to perform the requested action is implementation-defined.
- A batch server that accepts a *Server Status Request* shall return a *Server Status Reply* to the batch client.

4644 3.2.3.15 Signal Batch Job Request

A batch client can request that a batch server signal the session leader of a batch job. Such a request is called a *Signal Batch Job Request*.

- A batch server shall reject a *Signal Batch Job Request* if any of the following statements are true:
- The user of the batch client is not authorized to signal the batch job.
- The job is not in the RUNNING state.
 - The batch server does not manage the designated job.
- The requested signal is not supported by the implementation.
- A batch server may reject a *Signal Batch Job Request* for other implementation-defined reasons.

 The method used to determine whether the user of a client is authorized to perform the requested action is implementation-defined.
- A batch server that accepts a request to signal a batch job shall send the signal requested by the batch client to the process group of the session leader of the batch job.

4657 3.2.3.16 Track Batch Job Request

Track Batch Job Request is an optional feature of batch servers. If an implementation supports
Track Batch Job Request, the statements in this section apply and the configuration variable
POSIX2_PBS_TRACK shall be set to 1.

Track Batch Job Request provides a method for tracking the current location of a batch job. Clients may use the tracking information to determine the batch server that should receive a batch server request.

4669

4670

4671

4672 4673

4676 4677

4682

4683

4684 4685

4686

4687

4688

4691

4692 4693

4694

4695

4696

4697 4698

4699

4700 4701

4702

4664 If *Track Batch Job Request* is supported by a batch server, then when the batch server queues a
4665 batch job as a result of a *Queue Batch Job Request*, and the batch server is not the batch server that
4666 created the batch job, the batch server shall send a *Track Batch Job Request* to the batch server that
4667 created the job.

If *Track Batch Job Request* is supported by a batch server, then the *Track Batch Job Request* may also be sent to other servers as a backup to the primary server. The method by which backup servers are specified is implementation-defined.

If *Track Batch Job Request* is supported by a batch server that receives a *Track Batch Job Request*, then the batch server shall record the current location of the batch job as contained in the request.

3.3 Common Behavior for Batch Environment Utilities

4675 3.3.1 Batch Job Identifier

A utility shall recognize *job_identifiers* of the format:

[sequence number] [.server name] [@server]

4678 where:

sequence_number An integer that, when combined with server_name, provides a batch job identifier that is unique within the batch system.

4681 server_name The name of the batch server to which the batch job was originally submitted.

server The name of the batch server that is currently managing the batch job.

If the application omits the batch *server_name* portion of a batch job identifier, a utility shall use the name of a default batch server.

If the application omits the batch *server* portion of a batch job identifier, a utility shall use:

- The batch server indicated by *server_name*, if present
- The name of the default batch server
 - The name of the batch server that is currently managing the batch job

4689 If only *@server* is specified, then the status of all jobs owned by the user on the requested server 4690 is listed.

The means by which a utility determines the default batch server is implementation-defined.

If the application presents the batch *server* portion of a batch job identifier to a utility, the utility shall send the request to the specified server.

A strictly conforming application shall use the syntax described for the job identifier. Whenever a batch job identifier is specified whose syntax is not recognized by an implementation, then a message for each error that occurs shall be written to standard error and the utility shall exit with an exit status greater than zero.

When a batch job identifier is supplied as an argument to a batch utility and the *server_name* portion of the batch job identifier is omitted, then the utility shall use the name of the default batch server.

When a batch job identifier is supplied as an argument to a batch utility and the batch *server* portion of the batch job identifier is omitted, then the utility shall use either:

• The name of the default batch server

4704 or:

4705 4706

4707

4708

4710

4713

4714

4715

4716 4717

4718

4719

4720

4722

4723

4724 4725

4726 4727

4728

4729 4730

4732

4733

4734

4735 4736

4737

4738

4739 4740

4741

The name of the batch server that is currently managing the batch job

When a batch job identifier is supplied as an argument to a batch utility and the batch *server* portion of the batch job identifier is specified, then the utility shall send the required *Batch Server Request* to the specified server.

4709 **3.3.2 Destination**

The utility shall recognize a *destination* of the format:

4711 [queue] [@server]

4712 where:

queue The name of a valid execution or routing queue at the batch server denoted by @server, defined as a string of up to 15 alphanumeric characters in the portable character set (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.1, Portable Character Set) where the first character is alphabetic.

server The name of a batch server, defined as a string of alphanumeric characters in the portable character set.

If the application omits the batch *server* portion of a destination, then the utility shall use either:

The name of the default batch server

4721 or:

• The name of the batch server that is currently managing the batch job

The means by which a utility determines the default batch server is implementation-defined.

If the application omits the *queue* portion of a destination, then the utility shall use the name of the default queue at the batch server chosen. The means by which a batch server determines its default queue is implementation-defined. If a destination is specified in the *queue@server* form, then the utility shall use the specified queue at the specified server.

A strictly conforming application shall use the syntax described for a destination. Whenever a destination is specified whose syntax is not recognized by an implementation, then a message shall be written to standard error and the utility shall exit with an exit status greater than zero.

4731 3.3.3 Multiple Keyword-Value Pairs

For each option that can have multiple keyword-value pair arguments, the following rules shall apply. Examples of options that can have list-oriented option-arguments are $-\mathbf{u}$ value@keyword and $-\mathbf{l}$ keyword=value.

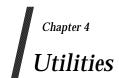
1. If a batch utility is presented with a list-oriented option-argument for which a keyword has a corresponding value that begins with a single or double quote, then the utility shall stop interpreting the input stream for delimiters until a second single or double quote, respectively, is encountered. This feature allows some flexibility for a comma (',') or equals sign ('=') to be part of the value string for a particular keyword; for example:

```
keywd1='val1,val2',keywd2="val3,val4"
```

Note: This may require the user to escape the quotes as in the following command:

4742 foo -xkeywd1='val1,val2',keywd2="val3,val4"

- 2. If a batch server is presented with a list-oriented attribute that has a keyword that was encountered earlier in the list, then the later entry for that keyword shall replace the earlier entry.
- 3. If a batch server is presented with a list-oriented attribute that has a keyword without any corresponding value of the form *keyword*= or *@keyword* and the same keyword was encountered earlier in the list, then the prior entry for that keyword shall be ignored by the batch server.
- 4. If a batch utility is expecting a list-oriented option-argument entry of the form *keyword=value*, but is presented with an entry of the form *keyword* without any corresponding *value*, then the entry shall be treated as though a default value of NULL was assigned (that is, *keyword=NULL*) for entry parsing purposes. The utility shall include only the keyword, not the NULL value, in the associated job attribute.
- 5. If a batch utility is expecting a list-oriented option-argument entry of the form *value@keyword*, but is presented with an entry of the form *value* without any corresponding *keyword*, then the entry shall be treated as though a keyword of NULL was assigned (that is, *value@NULL*) for entry parsing purposes. The utility shall include only the value, not the NULL keyword, in the associated job attribute.
- 6. A batch server shall accept a list-oriented attribute that has multiple occurrences of the same keyword, interpreting the keywords, in order, with the last value encountered taking precedence over prior instances of the same keyword. This rule allows, but does not require, a batch utility to preprocess the attribute to remove duplicate keywords.
- 7. If a batch utility is presented with multiple list-oriented option-arguments on the command line or in script directives, or both, for a single option, then the utility shall concatenate, in order, any command line keyword and value pairs to the end of any directive keyword and value pairs separated by a single comma to produce a single string that is an equivalent, valid option-argument. The resulting string shall be assigned to the associated attribute of the batch job (after optionally removing duplicate entries as described in item 6).



This chapter contains the definitions of the utilities, as follows:

- Mandatory utilities that are present on every conformant system
- Optional utilities that are present only on systems supporting the associated option; see Section 1.8.1 (on page 9) for information on the options in this volume of IEEE Std 1003.1-2001

admin Utilities

```
4777
    NAME
           admin — create and administer SCCS files (DEVELOPMENT)
4778
4779
    SYNOPSIS
            admin -i[name][-n][-a login][-d flag][-e login][-f flag][-m mrlist]
4780
4781
                [-r rel][-t[name][-y[comment]] newfile
            admin -n[-a login][-d flag][-e login][-f flag][-m mrlist][-t[name]]
                [-y[comment]] newfile ...
4783
            admin [-a login] [-d flag] [-m mrlist] [-r rel] [-t [name]] file ...
4784
            admin -h file ...
4785
            admin -z file ...
4786
4787
```

DESCRIPTION

4788

4789

4790

4791

4792 4793

4795 4796

4797

4798

4799

4800

4801

4802

4803

4805

4806

4807

4808

4809

The *admin* utility shall create new SCCS files or change parameters of existing ones. If a named file does not exist, it shall be created, and its parameters shall be initialized according to the specified options. Parameters not initialized by an option shall be assigned a default value. If a named file does exist, parameters corresponding to specified options shall be changed, and other parameters shall be left as is.

All SCCS filenames supplied by the application shall be of the form s. *filename*. New SCCS files shall be given read-only permission mode. Write permission in the parent directory is required to create a file. All writing done by *admin* shall be to a temporary *x-file*, named *x.filename* (see *get*) created with read-only mode if *admin* is creating a new SCCS file, or created with the same mode as that of the SCCS file if the file already exists. After successful execution of *admin*, the SCCS file shall be removed (if it exists), and the *x-file* shall be renamed with the name of the SCCS file. This ensures that changes are made to the SCCS file only if no errors occur.

The *admin* utility shall also use a transient lock file (named z.filename), which is used to prevent simultaneous updates to the SCCS file; see *get*.

OPTIONS

The *admin* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines, except that the $-\mathbf{i}$, $-\mathbf{t}$, and $-\mathbf{y}$ options have optional optionarguments. These optional option-arguments shall not be presented as separate arguments. The following options are supported:

- -**n** Create a new SCCS file. When −**n** is used without −**i**, the SCCS file shall be created with control information but without any file data.
- 4810 —i[name] Specify the name of a file from which the text for a new SCCS file shall be taken.

 4811 The text constitutes the first delta of the file (see the -r option for the delta
 4812 numbering scheme). If the -i option is used, but the name option-argument is
 4813 omitted, the text shall be obtained by reading the standard input. If this option is
 4814 omitted, the SCCS file shall be created with control information but without any
 4815 file data. The -i option implies the -n option.
- Specify the SID of the initial delta to be inserted. This SID shall be a trunk SID; that is, the branch and sequence numbers shall be zero or missing. The level number is optional, and defaults to 1.
- Specify the *name* of a file from which descriptive text for the SCCS file shall be taken. In the case of existing SCCS files (neither –**i** nor –**n** is specified):

Utilities admin

4821 • A -t option without a name option-argument shall cause the removal of 4822 descriptive text (if any) currently in the SCCS file. 4823 • A -t option with a name option-argument shall cause the text (if any) in the named file to replace the descriptive text (if any) currently in the SCCS file. 4824 4825 -f flag Specify a *flag*, and, possibly, a value for the *flag*, to be placed in the SCCS file. Several -f options may be supplied on a single admin command line. 4826 Implementations shall recognize the following flags and associated values: 4827 b Allow use of the $-\mathbf{b}$ option on a *get* command to create branch deltas. 4828 4829 cceil Specify the highest release (that is, ceiling), a number less than or equal to 9 999, which may be retrieved by a get command for editing. The default 4830 value for an unspecified c flag shall be 9 999. 4831 ffloor Specify the lowest release (that is, floor), a number greater than 0 but less 4832 than 9999, which may be retrieved by a get command for editing. The 4833 default value for an unspecified f flag shall be 1. 4834 **d**SID Specify the default delta number (SID) to be used by a *get* command. 4835 Treat the "No ID keywords" message issued by get or delta as a fatal istr 4836 error. In the absence of this flag, the message is only a warning. The 4837 message is issued if no SCCS identification keywords (see get) are found 4838 in the text retrieved or stored in the SCCS file. If a value is supplied, the 4839 4840 application shall ensure that the keywords exactly match the given string; however, the string shall contain a keyword, and no embedded 4841 <newline>s. 4842 j Allow concurrent *get* commands for editing on the same SID of an SCCS 4843 file. This allows multiple concurrent updates to the same version of the 4844 SCCS file. 4845 llist 4846 Specify a *list* of releases to which deltas can no longer be made (that is, *get* -e against one of these locked releases fails). Conforming applications 4847 shall use the following syntax to specify a *list*. Implementations may accept additional forms as an extension: 4849 <list> ::= a | <range-list> 4850 <range-list> ::= <range> | <range-list>, <range> 4851 <range> ::= <SID> 4852 The character a in the list shall be equivalent to specifying all releases for 4853 the named SCCS file. The non-terminal *SID* in range shall be the delta 4854 number of an existing delta associated with the SCCS file. 4855 Cause delta to create a null delta in each of those releases (if any) being 4856 n skipped when a delta is made in a new release (for example, in making 4857 4858 delta 5.1 after delta 2.7, releases 3 and 4 are skipped). These null deltas shall serve as anchor points so that branch deltas may later be created 4859 4860 from them. The absence of this flag shall cause skipped releases to be nonexistent in the SCCS file, preventing branch deltas from being created 4861 from them in the future. During the initial creation of an SCCS file, the **n** 4862 flag may be ignored; that is, if the -r option is used to set the release 4863 number of the initial SID to a value greater than 1, null deltas need not be 4864

created for the "skipped" releases.

4865

admin Utilities

4866 4867		qtext	Substitute user-definable $text$ for all occurrences of the %Q% keyword in the SCCS file text retrieved by get .
4868 4869 4870 4871		mmod	Specify the module name of the SCCS file substituted for all occurrences of the $\%M\%$ keyword in the SCCS file text retrieved by get . If the m flag is not specified, the value assigned shall be the name of the SCCS file with the leading $'$. $'$ removed.
4872 4873		ttype	Specify the <i>type</i> of module in the SCCS file substituted for all occurrences of the % Y % keyword in the SCCS file text retrieved by <i>get</i> .
4874 4875 4876 4877 4878		v pgm	Cause <i>delta</i> to prompt for modification request (MR) numbers as the reason for creating a delta. The optional value specifies the name of an MR number validation program. (If this flag is set when creating an SCCS file, the application shall ensure that the m option is also used even if its value is null.)
4879 4880 4881 4882	− d flag	supplied (The <i>llis</i>	(delete) the specified <i>flag</i> from an SCCS file. Several –d options may be d on a single <i>admin</i> command. See the –f option for allowable <i>flag</i> names. It flag gives a <i>list</i> of releases to be unlocked. See the –f option for further ion of the l flag and the syntax of a <i>list</i> .)
4883 4884 4885 4886 4887 4888 4889	− a login	may ma specifyin used on desired may add	a <i>login</i> name, or numerical group ID, to be added to the list of users who ake deltas (changes) to the SCCS file. A group ID shall be equivalent to a single <i>admin</i> common to that group ID. Several —a options may be a single <i>admin</i> command line. As many <i>logins</i> , or numerical group IDs, as may be on the list simultaneously. If the list of users is empty, then anyone deltas. If <i>login</i> or group ID is preceded by a '!', the users so specified denied permission to make deltas.
4890 4891 4892 4893	−e login	allowed equivale	a <i>login</i> name, or numerical group ID, to be erased from the list of users to make deltas (changes) to the SCCS file. Specifying a group ID is ent to specifying all <i>login</i> names common to that group ID. Several —e may be used on a single <i>admin</i> command line.
4894 4895 4896	-y[comment]	manner	ne <i>comment</i> text into the SCCS file as a comment for the initial delta in a identical to that of <i>delta</i> . In the POSIX locale, omission of the –y option ult in a default comment line being inserted in the form:
4897		"date	and time created %s %s by %s", <date>, <time>, <login></login></time></date>
4898 4899 4900		specifica	date> is expressed in the format of the date utility's $y/m/d$ conversion ation, <time> in the format of the date utility's T conversion specification and <login> is the login name of the user creating the file.</login></time>
4901 4902 4903 4904 4905	-m mrlist	reason f shall ens a value	ne list of modification request (MR) numbers into the SCCS file as the for creating the initial delta in a manner identical to <i>delta</i> . The application sure that the v flag is set and the MR numbers are validated if the v flag has (the name of an MR number validation program). A diagnostic message written if the v flag is not set or MR validation fails.
4906 4907 4908 4909	−h	with the	he structure of the SCCS file and compare the newly computed checksum checksum that is stored in the SCCS file. If the newly computed checksum t match the checksum in the SCCS file, a diagnostic message shall be
4910 4911	- z	_	ute the SCCS file checksum and store it in the first line of the SCCS file (see option above). Note that use of this option on a truly corrupted file may

Utilities admin

4912			prevent future detection of the corruption.
4913	OPERAND	S	
4914	The	e followin	g operands shall be supported:
4915 4916 4917 4918	file		A pathname of an existing SCCS file or a directory. If <i>file</i> is a directory, the <i>admin</i> utility shall behave as though each file in the directory were specified as a named file, except that non-SCCS files (last component of the pathname does not begin with s.) and unreadable files shall be silently ignored.
4919	nev	vfile	A pathname of an SCCS file to be created.
4920 4921 4922	line	e of the st	e file or newfile operand appears, and it is $'-'$, the standard input shall be read; each tandard input shall be taken to be the name of an SCCS file to be processed. Non-ind unreadable files shall be silently ignored.
4923	STDIN		
4924 4925 4926	if a	file or ne	d input shall be a text file used only if $-\mathbf{i}$ is specified without an option-argument or wfile operand is specified as '-'. If the first character of any standard input line is the POSIX locale, the results are unspecified.
4927	INPUT FILI		
4928	The	e existing	SCCS files shall be text files of an unspecified format.
4929 4930 4931 4932	a te uns	ext file; if t specified.	ion shall ensure that the file named by the —i option's <i>name</i> option-argument shall be the first character of any line in this file is <soh> in the POSIX locale, the results are If this file contains more than 99 999 lines, the number of lines recorded in the his file shall be 99 999 for this delta.</soh>
4933	ENVIRON	MENT VA	ARIABLES
4934	The	e followin	g environment variables shall affect the execution of admin:
4935 4936 4937 4938	LA	NG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
			Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
4939 4940	LC	_ALL	•
		_ALL _CTYPE	used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other
4940 4941 4942	LC	_CTYPE	used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. 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).
4940 4941 4942 4943	LC		used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. 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).
4940 4941 4942 4943 4944 4945 4946	LC.	_CTYPE	used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. 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). GES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and the contents of the default -y
4940 4941 4942 4943 4944 4945 4946 4947	LC. LC. NL ASYNCHRO	_CTYPE _MESSAG SPATH	used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. 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). GES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and the contents of the default –y comment. Determine the location of message catalogs for the processing of LC_MESSAGES.

Not used.

4952

admin Utilities

4953 **STDERR** The standard error shall be used only for diagnostic messages. 4954 4955 Any SCCS files created shall be text files of an unspecified format. During processing of a file, a 4956 4957 locking *z-file*, as described in *get* (on page 476), may be created and deleted. **EXTENDED DESCRIPTION** 4958 None. 4959 **EXIT STATUS** 4960 The following exit values shall be returned: 4961 Successful completion. 4962 >0 An error occurred. 4963 **CONSEQUENCES OF ERRORS** 4964 Default. 4965 APPLICATION USAGE 4966 It is recommended that directories containing SCCS files be writable by the owner only, and that 4967 4968 SCCS files themselves be read-only. The mode of the directories should allow only the owner to modify SCCS files contained in the directories. The mode of the SCCS files prevents any 4969 modification at all except by SCCS commands. 4970 **EXAMPLES** 4971 None. 4972 RATIONALE 4973 None. 4974 **FUTURE DIRECTIONS** 4975 None. 4976 **SEE ALSO** 4977 delta, get, prs, what 4978 **CHANGE HISTORY** 4979 First released in Issue 2. 4980 Issue 6 4981 The normative text is reworded to avoid use of the term "must" for application requirements. 4982 The normative text is reworded to emphasize the term "shall" for implementation requirements. 4983 4984 The grammar is updated. The Open Group Base Resolution bwg2001-007 is applied, adding new text to the INPUT FILES 4985

section warning that the maximum lines recorded in the file is 99 999.

option.

4986

4987

4988

The Open Group Base Resolution bwg2001-009 is applied, amending the description of the -h

alias **Utilities**

4989 **NAME** alias — define or display aliases 4990 4991 **SYNOPSIS** alias [alias-name[=string] ...] 4992 UP 4993 **DESCRIPTION** 4994 The alias utility shall create or redefine alias definitions or write the values of existing alias 4995 definitions to standard output. An alias definition provides a string value that shall replace a 4996 4997 command name when it is encountered; see Section 2.3.1 (on page 32). An alias definition shall affect the current shell execution environment and the execution 4998 environments of the subshells of the current shell. When used as specified by this volume of 4999 IEEE Std 1003.1-2001, the alias definition shall not affect the parent process of the current shell 5000 nor any utility environment invoked by the shell; see Section 2.12 (on page 61). 5001 **OPTIONS** 5002 None. 5003 **OPERANDS** 5004 The following operands shall be supported: 5005 Write the alias definition to standard output. 5006 alias-name 5007 alias-name=string 5008 Assign the value of *string* to the alias *alias-name*. 5009 If no operands are given, all alias definitions shall be written to standard output. **STDIN** 5010 Not used. 5011 **INPUT FILES** 5012 None. 5013 **ENVIRONMENT VARIABLES** 5014 5015 The following environment variables shall affect the execution of *alias*: LANG Provide a default value for the internationalization variables that are unset or null. 5016 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 5017 Internationalization Variables for the precedence of internationalization variables 5018 used to determine the values of locale categories.) 5019 5020 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 5021 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 5022 characters (for example, single-byte as opposed to multi-byte characters in 5023 arguments). 5024 LC_MESSAGES 5025 Determine the locale that should be used to affect the format and contents of 5026 diagnostic messages written to standard error.

Determine the location of message catalogs for the processing of *LC_MESSAGES*.

5027

5028 XSI

NLSPATH

alias Utilities

ASYNCHRONOUS EVENTS 5029 Default. 5030 5031 **STDOUT** The format for displaying aliases (when no operands or only *name* operands are specified) shall 5032 5033 "%s=%s\n", name, value 5034 5035 The value string shall be written with appropriate quoting so that it is suitable for reinput to the shell. See the description of shell quoting in Section 2.2 (on page 30). 5036 5037 **STDERR** 5038 The standard error shall be used only for diagnostic messages. **OUTPUT FILES** 5039 None. 5040 **EXTENDED DESCRIPTION** 5041 None 5042 **EXIT STATUS** 5043 The following exit values shall be returned: 5044 Successful completion. 5045 >0 One of the *name* operands specified did not have an alias definition, or an error occurred. 5046 **CONSEQUENCES OF ERRORS** 5047 5048 Default. APPLICATION USAGE 5049 5050 None **EXAMPLES** 5051 5052 1. Change *ls* to give a columnated, more annotated output: alias ls="ls -CF" 5053 2. Create a simple "redo" command to repeat previous entries in the command history file: 5054 alias r='fc -s' 5055 3. Use 1K units for du: 5056 alias du=du\ −k 5057 4. Set up *nohup* so that it can deal with an argument that is itself an alias name: 5058 alias nohup="nohup " 5059 **RATIONALE** 5060 The *alias* description is based on historical KornShell implementations. Known differences exist 5061 between that and the C shell. The KornShell version was adopted to be consistent with all the 5062 other KornShell features in this volume of IEEE Std 1003.1-2001, such as command line editing. 5063 5064 Since *alias* affects the current shell execution environment, it is generally provided as a shell regular built-in. 5065 Historical versions of the KornShell have allowed aliases to be exported to scripts that are 5066 invoked by the same shell. This is triggered by the alias -x flag; it is allowed by this volume of 5067

5068 5069 IEEE Std 1003.1-2001 only when an explicit extension such as -x is used. The standard

developers considered that aliases were of use primarily to interactive users and that they

Utilities alias

5070 5071	should normally not affect shell scripts called by those users; functions are available to such scripts.
5072 5073	Historical versions of the KornShell had not written aliases in a quoted manner suitable for reentry to the shell, but this volume of IEEE Std 1003.1-2001 has made this a requirement for all
5074	similar output. Therefore, consistency with this volume of IEEE Std 1003.1-2001 was chosen over
5075	this detail of historical practice.
5076	FUTURE DIRECTIONS
5077	None.
5078	SEE ALSO
5079	Section 2.9.5 (on page 54)
5080	
5081	First released in Issue 4.
5082	Issue 6
5083	This utility is marked as part of the User Portability Utilities option.
5084	The APPLICATION USAGE section is added.

ar Utilities

```
5085
    NAME
            ar — create and maintain library archives
5086
5087
    SYNOPSIS
            ar -d[-v] archive file ...
5088
    SD
5089
            ar -m [-v] archive file ...
5090
    XSI
            ar -m -a[-v] posname archive file ...
5091
5092
            ar -m -b[-v] posname archive file ...
            ar -m -i[-v] posname archive file ...
5093
5094
            ar -p[-v][-s] archive [file ...]
5095
    XSI
            ar -q[-cv] archive file ...
5096
    XSI
5097
5098
            ar -r[-cuv] archive file ...
            ar -r -a[-cuv] posname archive file ...
5099
    XSI
5100
            ar -r -b[-cuv] posname archive file ...
5101
            ar -r -i[-cuv] posname archive file ...
5102
5103
    XSI
            ar -t[-v][-s] archive [file ...]
            ar -x[-v][-sCT] archive [file ...]
5104
    XSI
```

5105 **DESCRIPTION**

5106 5107

5108 5109

5110

5111

5112

5113 5114

5115

5116

5117

5118

5119

5120 5121

5122

5123

5124

5125

5126

XSI

The *ar* utility is part of the Software Development Utilities option.

The *ar* utility can be used to create and maintain groups of files combined into an archive. Once an archive has been created, new files can be added, and existing files in an archive can be extracted, deleted, or replaced. When an archive consists entirely of valid object files, the implementation shall format the archive so that it is usable as a library for link editing (see *c99* and *fort77*). When some of the archived files are not valid object files, the suitability of the archive for library use is undefined. If an archive consists entirely of printable files, the entire archive shall be printable.

When *ar* creates an archive, it creates administrative information indicating whether a symbol table is present in the archive. When there is at least one object file that *ar* recognizes as such in the archive, an archive symbol table shall be created in the archive and maintained by *ar*; it is used by the link editor to search the archive. Whenever the *ar* utility is used to create or update the contents of such an archive, the symbol table shall be rebuilt. The –**s** option shall force the symbol table to be rebuilt.

All *file* operands can be pathnames. However, files within archives shall be named by a filename, which is the last component of the pathname used when the file was entered into the archive. The comparison of *file* operands to the names of files in archives shall be performed by comparing the last component of the operand to the name of the file in the archive.

It is unspecified whether multiple files in the archive may be identically named. In the case of such files, however, each *file* and *posname* operand shall match only the first file in the archive having a name that is the same as the last component of the operand.

XSI

Utilities ar

5127 5128	OPTION		y shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2,
5129			ax Guidelines.
5130		The following	ng options shall be supported:
5131	XSI	-a	Position new files in the archive after the file named by the <i>posname</i> operand.
5132	XSI	-b	Position new files in the archive before the file named by the <i>posname</i> operand.
5133 5134		- c	Suppress the diagnostic message that is written to standard error by default when the archive is created.
5135 5136 5137	XSI	- С	Prevent extracted files from replacing like-named files in the file system. This option is useful when $-T$ is also used, to prevent truncated filenames from replacing files with the same prefix.
5138		$-\mathbf{d}$	Delete one or more files from archive.
5139 5140	XSI	- i	Position new files in the archive before the file in the archive named by the <i>posname</i> operand (equivalent to $-\mathbf{b}$).
5141 5142 5143	XSI	-m	Move the named files in the archive. The $-\mathbf{a}$, $-\mathbf{b}$, or $-\mathbf{i}$ options with the <i>posname</i> operand indicate the position; otherwise, move the names files in the archive to the end of the archive.
5144 5145 5146		- p	Write the contents of the <i>files</i> in the archive named by <i>file</i> operands from <i>archive</i> to the standard output. If no <i>file</i> operands are specified, the contents of all files in the archive shall be written in the order of the archive.
5147 5148 5149	XSI	–q	Append the named files to the end of the archive. In this case <i>ar</i> does not check whether the added files are already in the archive. This is useful to bypass the searching otherwise done when creating a large archive piece by piece.
5150 5151 5152 5153 5154 5155 5156	XSI	- r	Replace or add <i>files</i> to <i>archive</i> . If the archive named by <i>archive</i> does not exist, a new archive shall be created and a diagnostic message shall be written to standard error (unless the $-\mathbf{c}$ option is specified). If no <i>files</i> are specified and the <i>archive</i> exists, the results are undefined. Files that replace existing files in the archive shall not change the order of the archive. Files that do not replace existing files in the archive shall be appended to the archive unless a $-\mathbf{a}$, $-\mathbf{b}$, or $-\mathbf{i}$ option specifies another position.
5157 5158 5159	XSI	-s	Force the regeneration of the archive symbol table even if <i>ar</i> is not invoked with an option that modifies the archive contents. This option is useful to restore the archive symbol table after it has been stripped; see <i>strip</i> .
5160 5161 5162		−t	Write a table of contents of <i>archive</i> to the standard output. The files specified by the <i>file</i> operands shall be included in the written list. If no <i>file</i> operands are specified, all files in <i>archive</i> shall be included in the order of the archive.
5163 5164 5165 5166	XSI	-T	Allow filename truncation of extracted files whose archive names are longer than the file system can support. By default, extracting a file with a name that is too long shall be an error; a diagnostic message shall be written and the file shall not be extracted.
5167 5168 5169		–u	Update older files in the archive. When used with the -r option, files in the archive shall be replaced only if the corresponding <i>file</i> has a modification time that is at least as new as the modification time of the file in the archive.

ar Utilities

5170 5171 5172		−v	Give verbose output. When used with the option characters $-\mathbf{d}$, $-\mathbf{r}$, or $-\mathbf{x}$, write a detailed file-by-file description of the archive creation and maintenance activity, as described in the STDOUT section.
5173 5174 5175			When used with $-\mathbf{p}$, write the name of the file in the archive to the standard output before writing the file in the archive itself to the standard output, as described in the STDOUT section.
5176 5177			When used with -t, include a long listing of information about the files in the archive, as described in the STDOUT section.
5178 5179 5180 5181		- x	Extract the files in the archive named by the <i>file</i> operands from <i>archive</i> . The contents of the archive shall not be changed. If no <i>file</i> operands are given, all files in the archive shall be extracted. The modification time of each file extracted shall be set to the time the file is extracted from the archive.
5182 5183	OPERA		ng operands shall be supported:
5184		archive	A pathname of the archive.
5185 5186 5187 5188 5189		file	A pathname. Only the last component shall be used when comparing against the names of files in the archive. If two or more <i>file</i> operands have the same last pathname component (basename), the results are unspecified. The implementation's archive format shall not truncate valid filenames of files added to or replaced in the archive.
5190 5191	XSI	posname	The name of a file in the archive, used for relative positioning; see options $-\mathbf{m}$ and $-\mathbf{r}$.
5192	STDIN		
5193		Not used.	
5194 5195	INPUT		named by <i>archive</i> shall be a file in the format created by <i>ar</i> – r .
5196 5197	ENVIRO	ONMENT VA The followin	ARIABLES ag environment variables shall affect the execution of ar:
5198 5199 5200 5201		LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
5202 5203		LC_ALL	If set to a non-empty string value, override the values of all the other
5204			internationalization variables.
5205 5206		LC_CTYPE	internationalization variables. 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).
		LC_CTYPE LC_MESSAC	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). GES
5206			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).
5206 5207 5208			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). GES Determine the locale that should be used to affect the format and contents of
5206 5207 5208 5209	XSI	LC_MESSAC	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). GES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.

Utilities ar

TZ5214 Determine the timezone used to calculate date and time strings written by ar - tv. If TZ is unset or null, an unspecified default timezone shall be used. 5215 5216 ASYNCHRONOUS EVENTS Default. 5217 STDOUT 5218 If the $-\mathbf{d}$ option is used with the $-\mathbf{v}$ option, the standard output format shall be: 5219 5220 $"d - %s\n", < file>$ 5221 where *file* is the operand specified on the command line. If the $-\mathbf{p}$ option is used with the $-\mathbf{v}$ option, ar shall precede the contents of each file with: 5222 $\n<$ s>\n\n", <file> 5223 where *file* is the operand specified on the command line, if *file* operands were specified, and the 5224 name of the file in the archive if they were not. 5225 If the $-\mathbf{r}$ option is used with the $-\mathbf{v}$ option: 5226 • If *file* is already in the archive, the standard output format shall be: 5227 "r - %s\n", <file> 5228 where *<file>* is the operand specified on the command line. 5229 5230 • If *file* is not already in the archive, the standard output format shall be: a - sn'', < file>5231 where *<file>* is the operand specified on the command line. 5232 5233 If the **-t** option is used, *ar* shall write the names of the files in the archive to the standard output 5234 in the format: "%s\n", <file> 5235 where file is the operand specified on the command line, if file operands were specified, or the 5236 5237 name of the file in the archive if they were not. If the -t option is used with the -v option, the standard output format shall be: 5238 "%s %u/%u %u %s %d %d:%d %d %s\n", <member mode>, <user ID>, 5239 <group ID>, <number of bytes in member>, 5240 5241 <abbreviated month>, <day-of-month>, <hour>, 5242 <minute>, <year>, <file> where: 5243 <file> Shall be the operand specified on the command line, if *file* operands were specified, 5244 or the name of the file in the archive if they were not. 5245 5246 <member mode> Shall be formatted the same as the *<file mode>* string defined in the STDOUT 5247

section of *ls*, except that the first character, the *<entry type>*, is not used; the string

represents the file mode of the file in the archive at the time it was added to or

The following represent the last-modification time of a file when it was most recently added to

replaced in the archive.

or replaced in the archive:

5248

5249

5250

5251

5252

ar Utilities

5253 5254	<abbreviated< th=""><th>month> Equivalent to the format of the %b conversion specification format in date.</th></abbreviated<>	month> Equivalent to the format of the %b conversion specification format in date.
5255	<day-of-mon< td=""><td></td></day-of-mon<>	
5256	houm	
5257	<hour></hour>	Equivalent to the format of the %H conversion specification format in <i>date</i> .
5258	<minute></minute>	Equivalent to the format of the %M conversion specification format in <i>date</i> .
5259	<year></year>	Equivalent to the format of the %Y conversion specification format in <i>date</i> .
5260 5261		TME does not specify the POSIX locale, a different format and order of presentation ds relative to each other may be used in a format appropriate in the specified locale.
5262	If the – x opt	tion is used with the $-\mathbf{v}$ option, the standard output format shall be:
5263	"x - %s\r	u", <file></file>
5264 5265		s the operand specified on the command line, if <i>file</i> operands were specified, or the file in the archive if they were not.
5266	STDERR	
5267 5268		d error shall be used only for diagnostic messages. The diagnostic message about ew archive when $-c$ is not specified shall not modify the exit status.
5269	OUTPUT FILES	
5270	Archives ar	e files with unspecified formats.
5270 5271 5272	Archives ar EXTENDED DESCR None.	-
5271	EXTENDED DESCR	-
5271 5272	EXTENDED DESCR None. EXIT STATUS	-
5271 5272 5273	EXTENDED DESCR None. EXIT STATUS The following	EIPTION
5271 5272 5273 5274	EXTENDED DESCR None. EXIT STATUS The following 0 Success	EIPTION ng exit values shall be returned:
5271 5272 5273 5274 5275	EXTENDED DESCR None. EXIT STATUS The following 0 Success	eiption Ing exit values shall be returned: In a completion. In a completion or occurred.
5271 5272 5273 5274 5275 5276 5277	EXTENDED DESCRING. None. EXIT STATUS The following of Successes to An error CONSEQUENCES CONSE	eiption Ing exit values shall be returned: In or occurred. OF ERRORS
5271 5272 5273 5274 5275 5276 5277 5278 5279 5280	EXTENDED DESCRINONE. None. EXIT STATUS The following of Success to Success to An error of Default. APPLICATION USA	eiption Ing exit values shall be returned: In or occurred. OF ERRORS
5271 5272 5273 5274 5275 5276 5277 5278 5279 5280 5281	EXTENDED DESCRENCES None. EXIT STATUS The following of Success >0 An error of An error of Consequences of Default. APPLICATION USANONE. EXAMPLES None. RATIONALE The archives which are made as a file. The the same intended as a file. The the same intended as a file. The the same intended as a file.	eiption Ing exit values shall be returned: In or occurred. OF ERRORS

Utilities such as *pax* (and its forebears *tar* and *cpio*) also provide portable "archives". This is a not a duplication; the *ar* utility is included to provide an interface primarily for *make* and the compilers, based on a historical model.

In historical implementations, the $-\mathbf{q}$ option (available on XSI-conforming systems) is known to execute quickly because ar does not check on whether the added members are already in the archive. This is useful to bypass the searching otherwise done when creating a large archive

5290

5291

5292

5293

5294

Utilities ar

piece-by-piece. These remarks may but need not remain true for a brand new implementation of this utility; hence, these remarks have been moved into the RATIONALE.

 BSD implementations historically required applications to provide the -s option whenever the archive was supposed to contain a symbol table. As in this volume of IEEE Std 1003.1-2001, System V historically creates or updates an archive symbol table whenever an object file is removed from, added to, or updated in the archive.

The OPERANDS section requires what might seem to be true without specifying it: the archive cannot truncate the filenames below {NAME_MAX}. Some historical implementations do so, however, causing unexpected results for the application. Therefore, this volume of IEEE Std 1003.1-2001 makes the requirement explicit to avoid misunderstandings.

According to the System V documentation, the options -**dmpqrtx** are not required to begin with a hyphen ('-'). This volume of IEEE Std 1003.1-2001 requires that a conforming application use the leading hyphen.

The archive format used by the 4.4 BSD implementation is documented in this RATIONALE as an example:

A file created by ar begins with the "magic" string "!<arch>\n". The rest of the archive is made up of objects, each of which is composed of a header for a file, a possible filename, and the file contents. The header is portable between machine architectures, and, if the file contents are printable, the archive is itself printable.

The header is made up of six ASCII fields, followed by a two-character trailer. The fields are the object name (16 characters), the file last modification time (12 characters), the user and group IDs (each 6 characters), the file mode (8 characters), and the file size (10 characters). All numeric fields are in decimal, except for the file mode, which is in octal.

The modification time is the file *st_mtime* field. The user and group IDs are the file *st_uid* and *st_gid* fields. The file mode is the file *st_mode* field. The file size is the file *st_size* field. The two-byte trailer is the string " '<newline>".

Only the name field has any provision for overflow. If any filename is more than 16 characters in length or contains an embedded space, the string "#1/" followed by the ASCII length of the name is written in the name field. The file size (stored in the archive header) is incremented by the length of the name. The name is then written immediately following the archive header.

Any unused characters in any of these fields are written as <space>s. If any fields are their particular maximum number of characters in length, there is no separation between the fields.

Objects in the archive are always an even number of bytes long; files that are an odd number of bytes long are padded with a <newline>, although the size in the header does not reflect this

The *ar* utility description requires that (when all its members are valid object files) *ar* produce an object code library, which the linkage editor can use to extract object modules. If the linkage editor needs a symbol table to permit random access to the archive, *ar* must provide it; however, *ar* does not require a symbol table.

The BSD $-\mathbf{o}$ option was omitted. It is a rare conforming application that uses ar to extract object code from a library with concern for its modification time, since this can only be of importance to make. Hence, since this functionality is not deemed important for applications portability, the modification time of the extracted files is set to the current time.

Utilities ar

5340 There is at least one known implementation (for a small computer) that can accommodate only 5341 object files for that system, disallowing mixed object and other files. The ability to handle any type of file is not only historical practice for most implementations, but is also a reasonable 5342 expectation. 5343

> Consideration was given to changing the output format of ar -tv to the same format as the output of ls –l. This would have made parsing the output of ar the same as that of ls. This was rejected in part because the current ar format is commonly used and changes would break historical usage. Second, ar gives the user ID and group ID in numeric format separated by a slash. Changing this to be the user name and group name would not be correct if the archive were moved to a machine that contained a different user database. Since ar cannot know whether the archive was generated on the same machine, it cannot tell what to report.

> The text on the -ur option combination is historical practice—since one filename can easily represent two different files (for example, /a/foo and /b/foo), it is reasonable to replace the file in the archive even when the modification time in the archive is identical to that in the file system.

FUTURE DIRECTIONS

None. 5355

SEE ALSO 5356

5344

5345

5346

5347

5348

5350

5351

5352

5353

5354

5357 5358

5360

5362

5364

5367

5368

5369

5371 5372

c99, date, fort77, pax, strip the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 13, Headers, <unistd.h> description of {POSIX_NO_TRUNC}

CHANGE HISTORY 5359

First released in Issue 2.

5361 Issue 5

The FUTURE DIRECTIONS section is added.

5363 Issue 6

This utility is marked as part of the Software Development Utilities option.

The STDOUT description is changed for the -v option to align with the IEEE P1003.2b draft 5365 5366 standard.

The normative text is reworded to avoid use of the term "must" for application requirements.

The *TZ* entry is added to the ENVIRONMENT VARIABLES section.

IEEE PASC Interpretation 1003.2 #198 is applied, changing the description to consistently use "file" to refer to a file in the file system hierarchy, "archive" to refer to the archive being 5370 operated upon by the ar utility, and "file in the archive" to refer to a copy of a file that is contained in the archive.

IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/10 is applied, making corrections to the 5373 SYNOPSIS. The change was needed since the $-\mathbf{a}$, $-\mathbf{b}$, and $-\mathbf{i}$ options are mutually-exclusive, and 5374 posname is required if any of these options is specified. 5375

IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/11 is applied, correcting the description 5376 of the two-byte trailer in RATIONALE which had missed out a backquote. The correct trailer is a 5377 backquote followed by a <newline>. 5378

Utilities asa

5379 NAME $as a - interpret\ carriage\text{-}control\ characters$ 5380 5381 **SYNOPSIS** [file ...] 5382 FR asa 5383 DESCRIPTION 5384 The asa utility shall write its input files to standard output, mapping carriage-control characters 5385 from the text files to line-printer control sequences in an implementation-defined manner. 5386 The first character of every line shall be removed from the input, and the following actions are performed. 5388 If the character removed is: 5389 <space> The rest of the line is output without change. 5390 0 A <newline> is output, then the rest of the input line. 5391 1 One or more implementation-defined characters that causes an advance to the next 5392 page shall be output, followed by the rest of the input line. 5393 The <newline> of the previous line shall be replaced with one or more 5394 + implementation-defined characters that causes printing to return to column position 1, 5395 followed by the rest of the input line. If the '+' is the first character in the input, it shall 5396 be equivalent to <space>. 5397 The action of the asa utility is unspecified upon encountering any character other than those 5398 listed above as the first character in a line. 5399 **OPTIONS** 5400 None 5401 **OPERANDS** 5402 5403 file A pathname of a text file used for input. If no file operands are specified, the standard input shall be used. 5404 5405 STDIN The standard input shall be used only if no file operands are specified; see the INPUT FILES 5406 section. 5407 **INPUT FILES** 5408 The input files shall be text files. 5409 **ENVIRONMENT VARIABLES** 5410 The following environment variables shall affect the execution of asa: 5411 LANG Provide a default value for the internationalization variables that are unset or null. 5412 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 5413 Internationalization Variables for the precedence of internationalization variables 5414 used to determine the values of locale categories.) 5415 LC_ALL If set to a non-empty string value, override the values of all the other 5416 internationalization variables. 5417

arguments and input files).

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

LC_CTYPE

5418

5419 5420 **asa** Utilities

5421 LC_MESSAGES 5422 Determine the locale that should be used to affect the format and contents of 5423 diagnostic messages written to standard error. NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 5424 XSI ASYNCHRONOUS EVENTS 5425 Default. 5426 **STDOUT** 5427 The standard output shall be the text from the input file modified as described in the 5428 DESCRIPTION section. STDERR 5430 None. 5431 **OUTPUT FILES** 5432 None. 5433 EXTENDED DESCRIPTION 5434 None. 5435 **EXIT STATUS** 5436 The following exit values shall be returned: 5437 5438 All input files were output successfully. 5439 An error occurred. **CONSEQUENCES OF ERRORS** 5440 Default. 5441 APPLICATION USAGE 5449 None. 5443 **EXAMPLES** 5444 The following command: 5445 5446 asa file permits the viewing of file (created by a program using FORTRAN-style carriage-control 5447 characters) on a terminal. 5448 2. The following command: 5449 5450 a.out | asa | lp formats the FORTRAN output of **a.out** and directs it to the printer. 5451 **RATIONALE** The asa utility is needed to map "standard" FORTRAN 77 output into a form acceptable to 5453 contemporary printers. Usually, asa is used to pipe data to the *lp* utility; see *lp*. 5454 This utility is generally used only by FORTRAN programs. The standard developers decided to 5455 retain asa to avoid breaking the historical large base of FORTRAN applications that put 5456 carriage-control characters in their output files. There is no requirement that a system have a FORTRAN compiler in order to run applications that need asa. 5458 Historical implementations have used an ASCII <form-feed> in response to a 1 and an ASCII 5459 <carriage-return> in response to a '+'. It is suggested that implementations treat characters 5460 other than 0, 1, and '+' as <space> in the absence of any compelling reason to do otherwise. 5461 However, the action is listed here as "unspecified", permitting an implementation to provide 5462

Utilities asa

5463	extensions to access fast multiple-line slewing and channel seeking in a non-portable manner.
5464 5465	FUTURE DIRECTIONS None.
5466 5467	SEE ALSO fort77, lp
5468 5469	CHANGE HISTORY First released in Issue 4.
5470 5471	Issue 6 This utility is marked as part of the FORTRAN Runtime Utilities option.
5472	The normative text is reworded to avoid use of the term "must" for application requirements.

at Utilities

```
5473 NAME
5474 at — execute commands at a later time
5475 SYNOPSIS
5476 UP at [-m] [-f file] [-q queuename
```

```
5476 UP at [-m] [-f file] [-q queuename] -t time_arg

5477 at [-m] [-f file] [-q queuename] timespec ...

5478 at -r at_job_id ...

5479 at -l -q queuename

5480 at -l [at job id ...]
```

XSI

DESCRIPTION

The *at* utility shall read commands from standard input and group them together as an *at-job*, to be executed at a later time.

The at-job shall be executed in a separate invocation of the shell, running in a separate process group with no controlling terminal, except that the environment variables, current working directory, file creation mask, and other implementation-defined execution-time attributes in effect when the *at* utility is executed shall be retained and used when the at-job is executed.

When the at-job is submitted, the at_job_id and scheduled time shall be written to standard error. The at_job_id is an identifier that shall be a string consisting solely of alphanumeric characters and the period character. The at_job_id shall be assigned by the system when the job is scheduled such that it uniquely identifies a particular job.

User notification and the processing of the job's standard output and standard error are described under the -m option.

Users shall be permitted to use *at* if their name appears in the file /usr/lib/cron/at.allow. If that file does not exist, the file /usr/lib/cron/at.deny shall be checked to determine whether the user shall be denied access to *at*. If neither file exists, only a process with the appropriate privileges shall be allowed to submit a job. If only at.deny exists and is empty, global usage shall be permitted. The at.allow and at.deny files shall consist of one user name per line.

OPTIONS

-m

The *at* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported:

5504	− f file	Specify the pathname of a file to be used as the source of the at-job, instead of
5505		standard input.

-I (The letter ell.) Report all jobs scheduled for the invoking user if no at_job_id operands are specified. If at_job_ids are specified, report only information for these jobs. The output shall be written to standard output.

Send mail to the invoking user after the at-job has run, announcing its completion. Standard output and standard error produced by the at-job shall be mailed to the user as well, unless redirected elsewhere. Mail shall be sent even if the job produces no output.

If -m is not used, the job's standard output and standard error shall be provided to the user by means of mail, unless they are redirected elsewhere; if there is no such output to provide, the implementation need not notify the user of the job's completion.

Utilities at

5517	5517a auguanama						
5517 5518 5519 5520 5521 5522 5523	− q queuenam	Specify in which queue to schedule a job for submission. When used with the $-\mathbf{l}$ option, limit the search to that particular queue. By default, at-jobs shall be scheduled in queue a . In contrast, queue b shall be reserved for batch jobs; see $batch$. The meanings of all other $queuenames$ are implementation-defined. If $-\mathbf{q}$ is specified along with either of the $-\mathbf{t}$ $time_arg$ or $timespec$ arguments, the results are unspecified.					
5524 5525	- r		jobs with ty the <i>at</i> utility	the specified <i>at_job_id</i> operands that were previously .			
5526 5527	-t time_arg			at the time specified by the <i>time</i> option-argument, which re has the format as specified by the <i>touch</i> –t <i>time</i> utility.			
5528 5529	OPERANDS The following	ng operands sh	nall be suppor	rted:			
5530 5531	at_job_id	The name re	ported by a p	revious invocation of the <i>at</i> utility at the time the job was			
5532 5533 5534 5535 5536	timespec	are interpret be parsed as shall be inter	Submit the job to be run at the date and time specified. All of the <i>timespec</i> operands are interpreted as if they were separated by <space>s and concatenated, and shall be parsed as described in the grammar at the end of this section. The date and time shall be interpreted as being in the timezone of the user (as determined by the <i>TZ</i> variable), unless a timezone name appears as part of <i>time</i>, below.</space>				
5537 5538 5539		In the POSIX locale, the following describes the three parts of the time specification string. All of the values from the <i>LC_TIME</i> categories in the POSIX locale shall be recognized in a case-insensitive manner.					
5540 5541 5542 5543 5544 5545 5546 5547 5548 5549 5550 5551		time	two-digit nu be hours and numbers se indication (<i>LC_TIME</i> lo clock time s further qua implementa and the strin	m be specified as one, two, or four digits. One-digit and ambers shall be taken to be hours; four-digit numbers to diminutes. The time can alternatively be specified as two parated by a colon, meaning hour:minute. An AM/PM one of the values from the am_pm keywords in the cale category) can follow the time; otherwise, a 24-hour hall be understood. A timezone name can also follow to diffy the time. The acceptable timezone names are tion-defined, except that they shall be case-insensitive ing utc is supported to indicate the time is in Coordinated time. In the POSIX locale, the time field can also be one of g tokens:			
5552			midnight	Indicates the time 12:00 am (00:00).			
5553			noon	Indicates the time 12:00 pm.			
5554			now	Indicates the current day and time. Invoking at <now></now>			
5555 5556 5557				shall submit an at-job for potentially immediate execution (that is, subject only to unspecified scheduling delays).			
5558 5559 5560 5561 5562 5563		date	values from category) for preceded by the day or a	date can be specified as either a month name (one of the a the mon or abmon keywords in the <i>LC_TIME</i> locale ollowed by a day number (and possibly year number a comma), or a day of the week (one of the values from abday keywords in the <i>LC_TIME</i> locale category). In the e, two special days shall be recognized:			

at Utilities

```
5564
                                    today
                                                Indicates the current day.
                                                Indicates the day following the current day.
5565
                                    tomorrow
                                    If no date is given, today shall be assumed if the given time is greater
5566
5567
                                     than the current time, and tomorrow shall be assumed if it is less. If
                                     the given month is less than the current month (and no year is given),
5568
                                     next year shall be assumed.
5569
5570
                        increment
                                    The optional increment shall be a number preceded by a plus sign
                                     ('+') and suffixed by one of the following: minutes, hours, days,
5571
                                    weeks, months, or years. (The singular forms shall also be
5572
                                     accepted.) The keyword next shall be equivalent to an increment
5573
                                     number of +1. For example, the following are equivalent commands:
5574
                                     at 2pm + 1 week
5575
5576
                                     at 2pm next week
            The following grammar describes the precise format of timespec in the POSIX locale. The general
5577
            conventions for this style of grammar are described in Section 1.10 (on page 19). This formal
5578
            syntax shall take precedence over the preceding text syntax description. The longest possible
5579
            token or delimiter shall be recognized at a given point. When used in a timespec, white space
5580
            shall also delimit tokens.
5581
5582
             %token hr24clock hr min
             %token hr24clock hour
5583
             /*
5584
               An hr24clock hr min is a one, two, or four-digit number. A one-digit
5585
               or two-digit number constitutes an hr24clock hour. An hr24clock hour
5586
               may be any of the single digits [0,9], or may be double digits, ranging
5587
               from [00,23]. If an hr24clock hr min is a four-digit number, the
5588
               first two digits shall be a valid hr24clock hour, while the last two
5589
               represent the number of minutes, from [00,59].
5590
5591
5592
             %token wallclock hr min
             %token wallclock hour
5593
             /*
5594
               A wallclock_hr_min is a one, two-digit, or four-digit number.
5595
               A one-digit or two-digit number constitutes a wallclock hour.
5596
5597
               A wallclock hour may be any of the single digits [1,9], or may
5598
               be double digits, ranging from [01,12]. If a wallclock hr min
               is a four-digit number, the first two digits shall be a valid
5599
               wallclock_hour, while the last two represent the number of
5600
5601
               minutes, from [00,59].
5602
             */
5603
             %token minute
5604
5605
               A minute is a one or two-digit number whose value can be [0,9]
               or [00,59].
5606
             */
5607
5608
             %token day number
5609
               A day number is a number in the range appropriate for the particular
5610
```

5611

month and year specified by month name and year number, respectively.

Utilities at

```
5612
              If no year number is given, the current year is assumed if the given
5613
              date and time are later this year. If no year_number is given and
5614
              the date and time have already occurred this year and the month is
              not the current month, next year is the assumed year.
5615
5616
           */
           %token year number
5617
5618
             A year number is a four-digit number representing the year A.D., in
5619
             which the at job is to be run.
5620
5621
5622
           %token inc number
5623
              The inc number is the number of times the succeeding increment
5624
             period is to be added to the specified date and time.
5625
5626
5627
            %token timezone name
5628
              The name of an optional timezone suffix to the time field, in an
5629
              implementation-defined format.
5630
5631
5632
           %token month name
5633
5634
              One of the values from the mon or abmon keywords in the LC TIME
              locale category.
5635
           * /
5636
5637
           %token day of week
5638
              One of the values from the day or abday keywords in the LC TIME
5639
              locale category.
5640
5641
            */
           %token am pm
5642
5643
5644
              One of the values from the am pm keyword in the LC TIME locale
              category.
5645
5646
5647
           %start timespec
           응응
5648
5649
           timespec
                         : time
                         | time date
5650
                          time increment
5651
                          time date increment
5652
5653
                         nowspec
5654
                         : "now"
5655
           nowspec
                          "now" increment
5656
5657
           time
                         : hr24clock hr min
5658
5659
                         hr24clock hr min timezone name
```

at Utilities

```
5660
                           hr24clock hour ":" minute
5661
                           hr24clock hour ":" minute timezone_name
5662
                           wallclock hr min am pm
                           wallclock hr min am pm timezone name
5663
5664
                           wallclock hour ": " minute am pm
                           wallclock hour ":" minute am pm timezone name
5665
                            "noon"
5666
                            "midnight"
5667
5668
                           month name day_number
            date
5669
                           month name day number "," year number
5670
                           day of week
5671
                            "today"
5672
                            "tomorrow"
5673
5674
                            "+" inc number inc period
            increment
5675
                            "next" inc period
5676
5677
                         : "minute" | "minutes"
            inc period
5678
                            "hour" | "hours"
5679
                            "day" | "days"
5680
5681
                            "week" | "weeks"
                            "month" | "months"
5682
                            "year" | "years"
5683
5684
```

5685 **STDIN**

5686

5687 5688

5690

5691 5692

5693

5695

5700

5701

5702

5703

5704

5705

5706

The standard input shall be a text file consisting of commands acceptable to the shell command language described in Chapter 2 (on page 29). The standard input shall only be used if no –f file option is specified.

5689 INPUT FILES

See the STDIN section.

The text files /usr/lib/cron/at.allow and /usr/lib/cron/at.deny shall contain zero or more user names, one per line, of users who are, respectively, authorized or denied access to the *at* and *batch* utilities.

5694 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *at*:

5696 LANG Provide a default value for the internationalization variables that are unset or null.
5697 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
5698 Internationalization Variables for the precedence of internationalization variables
5699 used to determine the values of locale categories.)

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

LC_CTYPE 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).

LC_MESSAGES

Determine the locale that should be used to affect the format and contents of

Utilities at

5707 5708		diagnostic messages written to standard error and informative messages written to standard output.
5709 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
5710 5711	LC_TIME	Determine the format and contents for date and time strings written and accepted by <i>at</i> .
5712 5713 5714 5715 5716	SHELL	Determine a name of a command interpreter to be used to invoke the at-job. If the variable is unset or null, <i>sh</i> shall be used. If it is set to a value other than a name for <i>sh</i> , the implementation shall do one of the following: use that shell; use <i>sh</i> ; use the login shell from the user database; or any of the preceding accompanied by a warning diagnostic about which was chosen.
5717 5718 5719 5720 5721	TZ	Determine the timezone. The job shall be submitted for execution at the time specified by <i>timespec</i> or —t <i>time</i> relative to the timezone specified by the <i>TZ</i> variable. If <i>timespec</i> specifies a timezone, it shall override <i>TZ</i> . If <i>timespec</i> does not specify a timezone and <i>TZ</i> is unset or null, an unspecified default timezone shall be used.
5722 ASYNO	CHRONOUS	EVENTS

ASYNCHRONOUS EVENTS

5723 Default.

STDOUT 5724

When standard input is a terminal, prompts of unspecified format for each line of the user input 5725 5726 described in the STDIN section may be written to standard output.

In the POSIX locale, the following shall be written to the standard output for each job when jobs 5727 are listed in response to the **–l** option: 5728

"%s\t%s\n", at job id, <date> 5729

5730 where *date* shall be equivalent in format to the output of:

date +"%a %b %e %T %Y" 5731

The date and time written shall be adjusted so that they appear in the timezone of the user (as 5732 5733 determined by the *TZ* variable).

STDERR 5734

5735

5736

5741

5743

5747

In the POSIX locale, the following shall be written to standard error when a job has been successfully submitted:

```
5737
           "job %s at %sn", at job id, <date>
```

where date has the same format as that described in the STDOUT section. Neither this, nor 5738 warning messages concerning the selection of the command interpreter, shall be considered a 5739 diagnostic that changes the exit status. 5740

Diagnostic messages, if any, shall be written to standard error.

OUTPUT FILES 5742

None.

EXTENDED DESCRIPTION 5744

None. 5745

EXIT STATUS 5746

The following exit values shall be returned:

The at utility successfully submitted, removed, or listed a job or jobs. 5748

at Utilities

```
5749 >0 An error occurred.
```

CONSEQUENCES OF ERRORS

The job shall not be scheduled, removed, or listed.

APPLICATION USAGE

5750

5751

5752 5753

5754

5755

5756

5757

5758

5759

5760

5761

5768

5769

5770

57715772

5773 5774

5775 5776

5777

5778 5779

5780 5781

5782

5783

5784

The format of the *at* command line shown here is guaranteed only for the POSIX locale. Other cultures may be supported with substantially different interfaces, although implementations are encouraged to provide comparable levels of functionality.

Since the commands run in a separate shell invocation, running in a separate process group with no controlling terminal, open file descriptors, traps, and priority inherited from the invoking environment are lost.

Some implementations do not allow substitution of different shells using *SHELL*. System V systems, for example, have used the login shell value for the user in /etc/passwd. To select reliably another command interpreter, the user must include it as part of the script, such as:

5767 EXAMPLES

1. This sequence can be used at a terminal:

```
at -m 0730 tomorrow
sort < file >outfile
EOT
```

2. This sequence, which demonstrates redirecting standard error to a pipe, is useful in a command procedure (the sequence of output redirection specifications is significant):

```
at now + 1 hour <<!
diff file1 file2 2>&1 >outfile | mailx mygroup
!
```

3. To have a job reschedule itself, *at* can be invoked from within the at-job. For example, this daily processing script named **my.daily** runs every day (although *crontab* is a more appropriate vehicle for such work):

```
# my.daily runs every day
daily processing
at now tomorrow < my.daily</pre>
```

4. The spacing of the three portions of the POSIX locale *timespec* is quite flexible as long as there are no ambiguities. Examples of various times and operand presentation include:

```
5785 at 0815am Jan 24

5786 at 8:15amjan24

5787 at now "+ 1day"

5788 at 5 pm FRIday

5789 at '17

5790 utc+

5791 30minutes'
```

Utilities at

RATIONALE

The *at* utility reads from standard input the commands to be executed at a later time. It may be useful to redirect standard output and standard error within the specified commands.

The -t *time* option was added as a new capability to support an internationalized way of specifying a time for execution of the submitted job.

Early proposals added a "jobname" concept as a way of giving submitted jobs names that are meaningful to the user submitting them. The historical, system-specified <code>at_job_id</code> gives no indication of what the job is. Upon further reflection, it was decided that the benefit of this was not worth the change in historical interface. The <code>at</code> functionality is useful in simple environments, but in large or complex situations, the functionality provided by the Batch Services option is more suitable.

The **-q** option historically has been an undocumented option, used mainly by the *batch* utility.

The System V –m option was added to provide a method for informing users that an at-job had completed. Otherwise, users are only informed when output to standard error or standard output are not redirected.

The behavior of *at* <**now**> was changed in an early proposal from being unspecified to submitting a job for potentially immediate execution. Historical BSD *at* implementations support this. Historical System V implementations give an error in that case, but a change to the System V versions should have no backwards-compatibility ramifications.

On BSD-based systems, a –**u** *user* option has allowed those with appropriate privileges to access the work of other users. Since this is primarily a system administration feature and is not universally implemented, it has been omitted. Similarly, a specification for the output format for a user with appropriate privileges viewing the queues of other users has been omitted.

The **–f** *file* option from System V is used instead of the BSD method of using the last operand as the pathname. The BSD method is ambiguous—does:

at 1200 friday

mean the same thing if there is a file named **friday** in the current directory?

The *at_job_id* is composed of a limited character set in historical practice, and it is mandated here to invalidate systems that might try using characters that require shell quoting or that could not be easily parsed by shell scripts.

The *at* utility varies between System V and BSD systems in the way timezones are used. On System V systems, the *TZ* variable affects the at-job submission times and the times displayed for the user. On BSD systems, *TZ* is not taken into account. The BSD behavior is easily achieved with the current specification. If the user wishes to have the timezone default to that of the system, they merely need to issue the *at* command immediately following an unsetting or null assignment to *TZ*. For example:

TZ= at noon ...

gives the desired BSD result.

While the *yacc*-like grammar specified in the OPERANDS section is lexically unambiguous with respect to the digit strings, a lexical analyzer would probably be written to look for and return digit strings in those cases. The parser could then check whether the digit string returned is a valid *day_number*, *year_number*, and so on, based on the context.

at Utilities

5834	FUTURE DIRECTIONS
5835	None.
5836 5837	SEE ALSO batch, crontab
5838 5839	CHANGE HISTORY First released in Issue 2.
5840 5841	Issue 6 This utility is marked as part of the User Portability Utilities option.
5842 5843	The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:
5844 5845	• If $-\mathbf{m}$ is not used, the job's standard output and standard error are provided to the user by mail.
5846 5847	The effects of using the $-\mathbf{q}$ and $-\mathbf{t}$ options as defined in the IEEE P1003.2b draft standard are specified.
5848	The normative text is reworded to avoid use of the term "must" for application requirements.

5849 NAME awk — pattern scanning and processing language 5850 5851 **SYNOPSIS** awk [-F ERE] [-v assignment] ... program [argument ...] 5852 awk [-F ERE] -f progfile ... [-v assignment] ...[argument ...] 5853 DESCRIPTION 5854 The awk utility shall execute programs written in the awk programming language, which is 5855 specialized for textual data manipulation. An awk program is a sequence of patterns and 5856 corresponding actions. When input is read that matches a pattern, the action associated with that pattern is carried out. 5858 Input shall be interpreted as a sequence of records. By default, a record is a line, less its 5859 terminating <newline>, but this can be changed by using the **RS** built-in variable. Each record of input shall be matched in turn against each pattern in the program. For each pattern matched, 5861 the associated action shall be executed. 5862 The awk utility shall interpret each input record as a sequence of fields where, by default, a field 5863 is a string of non-<blank>s. This default white-space field delimiter can be changed by using the 5864 **FS** built-in variable or -**F** *ERE*. The *awk* utility shall denote the first field in a record \$1, the 5865 second \$2, and so on. The symbol \$0 shall refer to the entire record; setting any other field causes 5866 the re-evaluation of \$0. Assigning to \$0 shall reset the values of all other fields and the NF built-5867 in variable. 5868 **OPTIONS** 5869 The awk utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines. 5871 5872 The following options shall be supported: -F ERE Define the input field separator to be the extended regular expression *ERE*, before any input is read; see **Regular Expressions** (on page 161). 5874 **−f** progfile Specify the pathname of the file progfile containing an awk program. If multiple 5875 instances of this option are specified, the concatenation of the files specified as progfile in the order specified shall be the awk program. The awk program can 5877 alternatively be specified in the command line as a single argument. 5878 v assignment 5879 The application shall ensure that the assignment argument is in the same form as an 5880 assignment operand. The specified variable assignment shall occur prior to 5881 executing the awk program, including the actions associated with **BEGIN** patterns 5882 (if any). Multiple occurrences of this option can be specified. 5883 **OPERANDS** 5884 The following operands shall be supported: 5885 If no -f option is specified, the first operand to awk shall be the text of the awk 5886 program program. The application shall supply the *program* operand as a single argument to 5887 5888 awk. If the text does not end in a <newline>, awk shall interpret the text as if it did. Either of the following two types of *argument* can be intermixed: 5889 argument file A pathname of a file that contains the input to be read, which is 5890 matched against the set of patterns in the program. If no file operands 5891

used.

are specified, or if a *file* operand is '-', the standard input shall be

5892

5893

assignment

An operand that begins with an underscore or alphabetic character from the portable character set (see the table in the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.1, Portable Character Set), followed by a sequence of underscores, digits, and alphabetics from the portable character set, followed by the '=' character, shall specify a variable assignment rather than a pathname. The characters before the '=' represent the name of an awk variable; if that name is an awk reserved word (see Grammar (on page 170)) the behavior is undefined. The characters following the equal sign shall be interpreted as if they appeared in the awk program preceded and followed by a double-quote (' "') character, as a STRING token (see Grammar (on page 170)), except that if the last character is an unescaped backslash, it shall be interpreted as a literal backslash rather than as the first character of the sequence "\"". The variable shall be assigned the value of that STRING token and, if appropriate, shall be considered a numeric string (see Expressions in awk (on page 156)), the variable shall also be assigned its numeric value. Each such variable assignment shall occur just prior to the processing of the following file, if any. Thus, an assignment before the first file argument shall be executed after the **BEGIN** actions (if any), while an assignment after the last file argument shall occur before the END actions (if any). If there are no file arguments, assignments shall be executed before processing the standard input.

5917 **STDIN**

5894

5895

5896

5897 5898

5899

5900

5901

5902

5903 5904

5905

5906

5908

5909

5910

5911

5912

5913

5914

5915

5916

5918

5919

5920 5921

59225923

5924

5925

5926

5927

5928

5929

5930

5931

5932

5933

5934 5935

5936

5937

5938

5939

The standard input shall be used only if no *file* operands are specified, or if a *file* operand is '-'; see the INPUT FILES section. If the *awk* program contains no actions and no patterns, but is otherwise a valid *awk* program, standard input and any *file* operands shall not be read and *awk* shall exit with a return status of zero.

INPUT FILES

Input files to the awk program from any of the following sources shall be text files:

- Any file operands or their equivalents, achieved by modifying the awk variables ARGV and ARGC
- Standard input in the absence of any file operands
- Arguments to the **getline** function

Whether the variable **RS** is set to a value other than a <newline> or not, for these files, implementations shall support records terminated with the specified separator up to {LINE_MAX} bytes and may support longer records.

If **–f** *progfile* is specified, the application shall ensure that the files named by each of the *progfile* option-arguments are text files and their concatenation, in the same order as they appear in the arguments, is an *awk* program.

ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *awk*:

LANG

Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)

	<i>LC_ALL</i> If set to a non-empty string value, override the values of all the other internationalization variables.					
	LC_COLLATE					
	Determine the locale for the behavior of ranges, equivalence classes, and multi-					
		character collating elements within regular expressions and in comparisons of string values.				
	LC_CTYPE	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), the behavior of character classes within regular expressions, the identification of characters as letters, and the mapping of				
		uppercase and lowercase characters for the toupper and tolower functions.				
	LC_MESSAC					
		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.				
	LC_NUMER					
		Determine the radix character used when interpreting numeric input, performing conversions between numeric and string values, and formatting numeric output.				
		Regardless of locale, the period character (the decimal-point character of the				
		POSIX locale) is the decimal-point character recognized in processing awk				
		programs (including assignments in command line arguments).				
XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .				
	PATH	Determine the search path when looking for commands executed by <i>system(expr)</i> ,				
		or input and output pipes; see the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables.				
	In addition, all environment variables shall be visible via the awk variable ENVIRON.					
965 ASYNCHRONOUS EVENTS 966 Default.						
7 STDOUT 8 The nature of the output files depends on the <i>awk</i> program.						
STDER	R					
	The standard	d error shall be used only for diagnostic messages.				
OUTPU						
The nature of the output files depends on the awk program.						
EXTENI	DED DESCRI	IPTION				
	Overall Prog	gram Structure				
	An awk prog	ram is composed of pairs of the form:				
<pre>pattern { action }</pre>						
	Either the pa	ttern or the action (including the enclosing brace characters) can be omitted.				
	A missing pa	attern shall match any record of input, and a missing action shall be equivalent to:				
	{ print }					
Execution of the <i>awk</i> program shall start by first executing the actions associated with all BEGIN patterns in the order they occur in the program. Then each <i>file</i> operand (or standard input if no						
	ASYNC STDOU STDER!	LC_COLLAT LC_COLLAT LC_CTYPE LC_MESSAC LC_NUMER LC_NUMER ASYNCHRONOUS I Default. STDOUT The nature o STDERR The standard OUTPUT FILES The nature o EXTENDED DESCRI Overall Program An awk program pattern { Either the path A missing path of print } Execution of				

files were specified) shall be processed in turn by reading data from the file until a record separator is seen (<newline> by default). Before the first reference to a field in the record is evaluated, the record shall be split into fields, according to the rules in **Regular Expressions** (on page 161), using the value of **FS** that was current at the time the record was read. Each pattern in the program then shall be evaluated in the order of occurrence, and the action associated with each pattern that matches the current record executed. The action for a matching pattern shall be executed before evaluating subsequent patterns. Finally, the actions associated with all **END** patterns shall be executed in the order they occur in the program.

Expressions in awk

Expressions describe computations used in *patterns* and *actions*. In the following table, valid expression operations are given in groups from highest precedence first to lowest precedence last, with equal-precedence operators grouped between horizontal lines. In expression evaluation, where the grammar is formally ambiguous, higher precedence operators shall be evaluated before lower precedence operators. In this table *expr*, *expr1*, *expr2*, and *expr3* represent any expression, while lvalue represents any entity that can be assigned to (that is, on the left side of an assignment operator). The precise syntax of expressions is given in **Grammar** (on page 170).

Table 4-1 Expressions in Decreasing Precedence in awk

Syntax	Name	Type of Result	Associativity
(expr)	Grouping	Type of expr	N/A
\$expr	Field reference	String	N/A
++ lvalue	Pre-increment	Numeric	N/A
lvalue	Pre-decrement	Numeric	N/A
lvalue ++	Post-increment	Numeric	N/A
lvalue	Post-decrement	Numeric	N/A
expr ^ expr	Exponentiation	Numeric	Right
! expr	Logical not	Numeric	N/A
+ expr	Unary plus	Numeric	N/A
- expr	Unary minus	Numeric	N/A
expr * expr	Multiplication	Numeric	Left
expr / expr	Division	Numeric	Left
expr % expr	Modulus	Numeric	Left
expr + expr	Addition	Numeric	Left
expr - expr	Subtraction	Numeric	Left
expr expr	String concatenation	String	Left
expr < expr	Less than	Numeric	None
expr <= expr	Less than or equal to	Numeric	None
expr != expr	Not equal to	Numeric	None
expr == expr	Equal to	Numeric	None
expr > expr	Greater than	Numeric	None
expr >= expr	Greater than or equal to	Numeric	None

6024
6025
6026
6027
6028
0020
6029
6030
6031
6032
6033
6034
6035
6036
6037
6038
6039
6040

Syntax	Name	Type of Result	Associativity
expr ~ expr	ERE match	Numeric	None
expr !~ expr	ERE non-match	Numeric	None
expr in array	Array membership	Numeric	Left
(index) in array	Multi-dimension array membership	Numeric	Left
expr && expr	Logical AND	Numeric	Left
expr expr	Logical OR	Numeric	Left
expr1 ? expr2 : expr3	Conditional expression	Type of selected expr2 or expr3	Right
lvalue ^= expr	Exponentiation assignment	Numeric	Right
lvalue %= expr	Modulus assignment	Numeric	Right
lvalue *= expr	Multiplication assignment	Numeric	Right
lvalue /= expr	Division assignment	Numeric	Right
lvalue += expr	Addition assignment	Numeric	Right
lvalue -= expr	Subtraction assignment	Numeric	Right
lvalue = expr	Assignment	Type of expr	Right

Each expression shall have either a string value, a numeric value, or both. Except as stated for specific contexts, the value of an expression shall be implicitly converted to the type needed for the context in which it is used. A string value shall be converted to a numeric value by the equivalent of the following calls to functions defined by the ISO C standard:

```
setlocale(LC_NUMERIC, "");
numeric value = atof(string value);
```

A numeric value that is exactly equal to the value of an integer (see Section 1.7.2 (on page 7)) shall be converted to a string by the equivalent of a call to the **sprintf** function (see **String Functions** (on page 167)) with the string "%d" as the *fint* argument and the numeric value being converted as the first and only *expr* argument. Any other numeric value shall be converted to a string by the equivalent of a call to the **sprintf** function with the value of the variable **CONVFMT** as the *fint* argument and the numeric value being converted as the first and only *expr* argument. The result of the conversion is unspecified if the value of **CONVFMT** is not a floating-point format specification. This volume of IEEE Std 1003.1-2001 specifies no explicit conversions between numbers and strings. An application can force an expression to be treated as a number by adding zero to it, or can force it to be treated as a string by concatenating the null string (" ") to it.

A string value shall be considered a *numeric string* if it comes from one of the following:

- 1. Field variables
- Input from the getline() function
- 3. FILENAME
- 4. **ARGV** array elements
- 5. **ENVIRON** array elements
- 6. Array elements created by the *split()* function
 - 7. A command line variable assignment

8. Variable assignment from another numeric string variable

and after all the following conversions have been applied, the resulting string would lexically be recognized as a **NUMBER** token as described by the lexical conventions in **Grammar** (on page 170):

- All leading and trailing <blank>s are discarded.
- If the first non-<blank> is '+' or '-', it is discarded.
- Changing each occurrence of the decimal point character from the current locale to a period.

If a '-' character is ignored in the preceding description, the numeric value of the *numeric string* shall be the negation of the numeric value of the recognized **NUMBER** token. Otherwise, the numeric value of the *numeric string* shall be the numeric value of the recognized **NUMBER** token. Whether or not a string is a *numeric string* shall be relevant only in contexts where that term is used in this section.

When an expression is used in a Boolean context, if it has a numeric value, a value of zero shall be treated as false and any other value shall be treated as true. Otherwise, a string value of the null string shall be treated as false and any other value shall be treated as true. A Boolean context shall be one of the following:

- The first subexpression of a conditional expression
- An expression operated on by logical NOT, logical AND, or logical OR
- The second expression of a **for** statement
- The expression of an **if** statement
- The expression of the **while** clause in either a **while** or **do...while** statement
 - An expression used as a pattern (as in Overall Program Structure)

All arithmetic shall follow the semantics of floating-point arithmetic as specified by the ISO C standard (see Section 1.7.2 (on page 7)).

The value of the expression:

6092 expr1 ^ expr2

6068

6069

6070 6071

6072

6073

6074

6075

6076

6077

6078

6079

6080

6081

6082

6083 6084

6088

6089

6090

shall be equivalent to the value returned by the ISO C standard function call:

6094 pow(expr1, expr2)

The expression:

6096 lvalue ^= expr

shall be equivalent to the ISO C standard expression:

6098 lvalue = pow(lvalue, expr)

except that Ivalue shall be evaluated only once. The value of the expression:

6100 expr1 % expr2

shall be equivalent to the value returned by the ISO C standard function call:

fmod(expr1, expr2)

6103 The expression:

1value %= expr

shall be equivalent to the ISO C standard expression:

```
6106 lvalue = fmod(lvalue, expr)
```

except that Ivalue shall be evaluated only once.

Variables and fields shall be set by the assignment statement:

```
6109 lvalue = expression
```

and the type of *expression* shall determine the resulting variable type. The assignment includes the arithmetic assignments ("+=", "-=", "*=", "/=", "%=", "^=", "++", "--") all of which shall produce a numeric result. The left-hand side of an assignment and the target of increment and decrement operators can be one of a variable, an array with index, or a field selector.

The *awk* language supplies arrays that are used for storing numbers or strings. Arrays need not be declared. They shall initially be empty, and their sizes shall change dynamically. The subscripts, or element identifiers, are strings, providing a type of associative array capability. An array name followed by a subscript within square brackets can be used as an Ivalue and thus as an expression, as described in the grammar; see **Grammar** (on page 170). Unsubscripted array names can be used in only the following contexts:

- A parameter in a function definition or function call
- The **NAME** token following any use of the keyword **in** as specified in the grammar (see **Grammar** (on page 170)); if the name used in this context is not an array name, the behavior is undefined

A valid array *index* shall consist of one or more comma-separated expressions, similar to the way in which multi-dimensional arrays are indexed in some programming languages. Because *awk* arrays are really one-dimensional, such a comma-separated list shall be converted to a single string by concatenating the string values of the separate expressions, each separated from the other by the value of the **SUBSEP** variable. Thus, the following two index operations shall be equivalent:

```
6130 var[expr1, expr2, ... exprn]
6131 var[expr1 SUBSEP expr2 SUBSEP ... SUBSEP exprn]
```

The application shall ensure that a multi-dimensioned *index* used with the **in** operator is parenthesized. The **in** operator, which tests for the existence of a particular array element, shall not cause that element to exist. Any other reference to a nonexistent array element shall automatically create it.

Comparisons (with the '<', "<=", "!=", "==", '>', and ">=" operators) shall be made numerically if both operands are numeric, if one is numeric and the other has a string value that is a numeric string, or if one is numeric and the other has the uninitialized value. Otherwise, operands shall be converted to strings as required and a string comparison shall be made using the locale-specific collation sequence. The value of the comparison expression shall be 1 if the relation is true, or 0 if the relation is false.

Variables and Special Variables

Variables can be used in an *awk* program by referencing them. With the exception of function parameters (see **User-Defined Functions** (on page 169)), they are not explicitly declared. Function parameter names shall be local to the function; all other variable names shall be global. The same name shall not be used as both a function parameter name and as the name of a function or a special *awk* variable. The same name shall not be used both as a variable name with global scope and as the name of a function. The same name shall not be used within the same scope both as a scalar variable and as an array. Uninitialized variables, including scalar variables, array elements, and field variables, shall have an uninitialized value. An uninitialized value shall have both a numeric value of zero and a string value of the empty string. Evaluation of variables with an uninitialized value, to either string or numeric, shall be determined by the context in which they are used.

Field variables shall be designated by a '\$' followed by a number or numerical expression. The effect of the field number *expression* evaluating to anything other than a non-negative integer is unspecified; uninitialized variables or string values need not be converted to numeric values in this context. New field variables can be created by assigning a value to them. References to nonexistent fields (that is, fields after \$NF), shall evaluate to the uninitialized value. Such references shall not create new fields. However, assigning to a nonexistent field (for example, \$(NF+2)=5) shall increase the value of NF; create any intervening fields with the uninitialized value; and cause the value of \$0 to be recomputed, with the fields being separated by the value of OFS. Each field variable shall have a string value or an uninitialized value when created. Field variables shall have the uninitialized value when created from \$0 using FS and the variable does not contain any characters. If appropriate, the field variable shall be considered a numeric string (see Expressions in awk (on page 156)).

Implementations shall support the following other special variables that are set by awk:

ARGC The number of elements in the **ARGV** array.

ARGV An array of command line arguments, excluding options and the *program* argument, numbered from zero to ARGC-1.

The arguments in **ARGV** can be modified or added to; **ARGC** can be altered. As each input file ends, *awk* shall treat the next non-null element of **ARGV**, up to the current value of **ARGC**-1, inclusive, as the name of the next input file. Thus, setting an element of **ARGV** to null means that it shall not be treated as an input file. The name '-' indicates the standard input. If an argument matches the format of an *assignment* operand, this argument shall be treated as an *assignment* rather than a *file* argument.

CONVFMT The **printf** format for converting numbers to strings (except for output statements, where **OFMT** is used); "% .6g" by default.

An array representing the value of the environment, as described in the *exec* functions defined in the System Interfaces volume of IEEE Std 1003.1-2001. The indices of the array shall be strings consisting of the names of the environment variables, and the value of each array element shall be a string consisting of the value of that variable. If appropriate, the environment variable shall be considered a *numeric string* (see **Expressions in awk** (on page 156)); the array element shall also have its numeric value.

In all cases where the behavior of *awk* is affected by environment variables (including the environment of any commands that *awk* executes via the **system** function or via pipeline redirections with the **print** statement, the **printf** statement, or the **getline** function), the environment used shall be the environment at the time

6168 ARG

ENVIRON

6190 6191		<i>awk</i> began executing; it is implementation-defined whether any modification of ENVIRON affects this environment.	
6192 6193 6194	FILENAME	A pathname of the current input file. Inside a BEGIN action the value is undefined. Inside an END action the value shall be the name of the last input file processed.	
6195 6196 6197	FNR	The ordinal number of the current record in the current file. Inside a BEGIN action the value shall be zero. Inside an END action the value shall be the number of the last record processed in the last file processed.	
6198	FS	Input field separator regular expression; a <space> by default.</space>	
6199 6200 6201 6202 6203	NF	The number of fields in the current record. Inside a BEGIN action, the use of NF is undefined unless a getline function without a <i>var</i> argument is executed previously. Inside an END action, NF shall retain the value it had for the last record read, unless a subsequent, redirected, getline function without a <i>var</i> argument is performed prior to entering the END action.	
6204 6205 6206	NR	The ordinal number of the current record from the start of input. Inside a BEGIN action the value shall be zero. Inside an END action the value shall be the number of the last record processed.	
6207 6208 6209 6210	OFMT	The printf format for converting numbers to strings in output statements (see Output Statements (on page 165)); "%.6g" by default. The result of the conversion is unspecified if the value of OFMT is not a floating-point format specification.	
6211	OFS	The print statement output field separation; <space> by default.</space>	
6212	ORS	The print statement output record separator; a <newline> by default.</newline>	
6213	RLENGTH	The length of the string matched by the match function.	
6214 6215 6216 6217 6218 6219	RS	The first character of the string value of RS shall be the input record separator; a <newline> by default. If RS contains more than one character, the results are unspecified. If RS is null, then records are separated by sequences consisting of a <newline> plus one or more blank lines, leading or trailing blank lines shall not result in empty records at the beginning or end of the input, and a <newline> shall always be a field separator, no matter what the value of FS is.</newline></newline></newline>	
6220 6221	RSTART	The starting position of the string matched by the match function, numbering from 1. This shall always be equivalent to the return value of the match function.	
6222 6223	SUBSEP	The subscript separator string for multi-dimensional arrays; the default value is implementation-defined.	
6224	Regular Expressions		
6225 6226 6227	Definitions vehicles that it shall a	lity shall make use of the extended regular expression notation (see the Base volume of IEEE Std 1003.1-2001, Section 9.4, Extended Regular Expressions) except allow the use of C-language conventions for escaping special characters within the selfied in the table in the Base Definitions volume of IEEE Std 1003.1.2001. Chapter 5	

EREs, as specified in the table in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5,

File Format Notation ('\\', '\a', '\b', '\f', '\n', '\r', '\t', '\v') and the following table; these escape sequences shall be recognized both inside and outside bracket expressions.

Note that records need not be separated by <newline>s and string constants can contain

<newline>s, so even the "\n" sequence is valid in awk EREs. Using a slash character within an

ERE requires the escaping shown in the following table.

6228 6229

6230

6231

6232

6233

6234

Table 4-2 Escape Sequences in awk

6235
6236
6237
6238
6239
6240
6241
6242
6243
6244
6245
6246
6247
6248
6249
6250
6251

6252

6253

6254 6255

6256

6257

6258

6259 6260

6261

6262

6263

6264 6265

6266 6267

6268

6269

6270

6271

6272

6273

6274

6275 6276

6277

6278 6279

Escape Sequence **Description** Meaning Backslash quotation-mark Quotation-mark character \/ Backslash slash Slash character \ddd A backslash character followed by the The character whose encoding is longest sequence of one, two, or three represented by the one, two, or threeoctal-digit characters (01234567). If all digit octal integer. Multi-byte of the digits are 0 (that is, characters require multiple, representation of the NUL character), concatenated escape sequences of this the behavior is undefined. type, including the leading $' \setminus '$ for each byte. A backslash character followed by any Undefined \c character not described in this table or in the table in the Base Definitions volume of IEEE Std 1003.1-2001. Chapter 5, File Format Notation ('\\','\a','\b','\f','\n', '\r','\t','\v').

A regular expression can be matched against a specific field or string by using one of the two regular expression matching operators, '~' and "!~". These operators shall interpret their right-hand operand as a regular expression and their left-hand operand as a string. If the regular expression matches the string, the '~' expression shall evaluate to a value of 1, and the "!~" expression shall evaluate to a value of 0. (The regular expression matching operation is as defined by the term matched in the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.1, Regular Expression Definitions, where a match occurs on any part of the string unless the regular expression is limited with the circumflex or dollar sign special characters.) If the regular expression does not match the string, the '~' expression shall evaluate to a value of 0, and the "!~" expression shall evaluate to a value of 1. If the right-hand operand is any expression other than the lexical token ERE, the string value of the expression shall be interpreted as an extended regular expression, including the escape conventions described above. Note that these same escape conventions shall also be applied in determining the value of a string literal (the lexical token STRING), and thus shall be applied a second time when a string literal is used in this context.

When an **ERE** token appears as an expression in any context other than as the right-hand of the '~' or "!~" operator or as one of the built-in function arguments described below, the value of the resulting expression shall be the equivalent of:

\$0 ~ /ere/

The *ere* argument to the **gsub**, **match**, **sub** functions, and the *fs* argument to the **split** function (see **String Functions** (on page 167)) shall be interpreted as extended regular expressions. These can be either **ERE** tokens or arbitrary expressions, and shall be interpreted in the same manner as the right-hand side of the ' ~' or "! ~" operator.

An extended regular expression can be used to separate fields by using the –**F** *ERE* option or by assigning a string containing the expression to the built-in variable **FS**. The default value of the **FS** variable shall be a single <space>. The following describes **FS** behavior:

1. If **FS** is a null string, the behavior is unspecified.

2. If **FS** is a single character:

a. If **FS** is <space>, skip leading and trailing <blank>s; fields shall be delimited by sets of one or more <blank>s.

- b. Otherwise, if **FS** is any other character *c*, fields shall be delimited by each single occurrence of *c*.
- 3. Otherwise, the string value of **FS** shall be considered to be an extended regular expression. Each occurrence of a sequence matching the extended regular expression shall delimit fields.

Except for the ' $^{\sim}$ ' and "! $^{\sim}$ " operators, and in the **gsub**, **match**, **split**, and **sub** built-in functions, ERE matching shall be based on input records; that is, record separator characters (the first character of the value of the variable **RS**, <newline> by default) cannot be embedded in the expression, and no expression shall match the record separator character. If the record separator is not <newline>, <newline>s embedded in the expression can be matched. For the ' $^{\sim}$ ' and "! $^{\sim}$ " operators, and in those four built-in functions, ERE matching shall be based on text strings; that is, any character (including <newline> and the record separator) can be embedded in the pattern, and an appropriate pattern shall match any character. However, in all *awk* ERE matching, the use of one or more NUL characters in the pattern, input record, or text string produces undefined results.

Patterns

A *pattern* is any valid *expression*, a range specified by two expressions separated by a comma, or one of the two special patterns **BEGIN** or **END**.

Special Patterns

The *awk* utility shall recognize two special patterns, **BEGIN** and **END**. Each **BEGIN** pattern shall be matched once and its associated action executed before the first record of input is read (except possibly by use of the **getline** function—see **Input/Output and General Functions** (on page 168)—in a prior **BEGIN** action) and before command line assignment is done. Each **END** pattern shall be matched once and its associated action executed after the last record of input has been read. These two patterns shall have associated actions.

BEGIN and **END** shall not combine with other patterns. Multiple **BEGIN** and **END** patterns shall be allowed. The actions associated with the **BEGIN** patterns shall be executed in the order specified in the program, as are the **END** actions. An **END** pattern can precede a **BEGIN** pattern in a program.

If an *awk* program consists of only actions with the pattern **BEGIN**, and the **BEGIN** action contains no **getline** function, *awk* shall exit without reading its input when the last statement in the last **BEGIN** action is executed. If an *awk* program consists of only actions with the pattern **END** or only actions with the patterns **BEGIN** and **END**, the input shall be read before the statements in the **END** actions are executed.

Expression Patterns

An expression pattern shall be evaluated as if it were an expression in a Boolean context. If the result is true, the pattern shall be considered to match, and the associated action (if any) shall be executed. If the result is false, the action shall not be executed.

Pattern Ranges

A pattern range consists of two expressions separated by a comma; in this case, the action shall be performed for all records between a match of the first expression and the following match of the second expression, inclusive. At this point, the pattern range can be repeated starting at input records subsequent to the end of the matched range.

Actions

An action is a sequence of statements as shown in the grammar in **Grammar** (on page 170). Any single statement can be replaced by a statement list enclosed in braces. The application shall ensure that statements in a statement list are separated by <newline>s or semicolons. Statements in a statement list shall be executed sequentially in the order that they appear.

The *expression* acting as the conditional in an **if** statement shall be evaluated and if it is non-zero or non-null, the following statement shall be executed; otherwise, if **else** is present, the statement following the **else** shall be executed.

The **if**, **while**, **do**...**while**, **for**, **break**, and **continue** statements are based on the ISO C standard (see Section 1.7.2 (on page 7)), except that the Boolean expressions shall be treated as described in **Expressions in awk** (on page 156), and except in the case of:

```
for (variable in array)
```

which shall iterate, assigning each *index* of *array* to *variable* in an unspecified order. The results of adding new elements to *array* within such a **for** loop are undefined. If a **break** or **continue** statement occurs outside of a loop, the behavior is undefined.

The **delete** statement shall remove an individual array element. Thus, the following code deletes an entire array:

```
for (index in array)
    delete array[index]
```

The **next** statement shall cause all further processing of the current input record to be abandoned. The behavior is undefined if a **next** statement appears or is invoked in a **BEGIN** or **END** action.

The **exit** statement shall invoke all **END** actions in the order in which they occur in the program source and then terminate the program without reading further input. An **exit** statement inside an **END** action shall terminate the program without further execution of **END** actions. If an expression is specified in an **exit** statement, its numeric value shall be the exit status of *awk*, unless subsequent errors are encountered or a subsequent **exit** statement with an expression is executed.

Output Statements

Both **print** and **printf** statements shall write to standard output by default. The output shall be written to the location specified by *output_redirection* if one is supplied, as follows:

```
> expression
>> expression
| expression
```

In all cases, the *expression* shall be evaluated to produce a string that is used as a pathname into which to write (for '>' or ">>") or as a command to be executed (for $'\mid '$). Using the first two forms, if the file of that name is not currently open, it shall be opened, creating it if necessary and using the first form, truncating the file. The output then shall be appended to the file. As long as the file remains open, subsequent calls in which *expression* evaluates to the same string value shall simply append output to the file. The file remains open until the **close** function (see **Input/Output and General Functions** (on page 168)) is called with an expression that evaluates to the same string value.

The third form shall write output onto a stream piped to the input of a command. The stream shall be created if no stream is currently open with the value of *expression* as its command name. The stream created shall be equivalent to one created by a call to the *popen()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001 with the value of *expression* as the *command* argument and a value of *w* as the *mode* argument. As long as the stream remains open, subsequent calls in which *expression* evaluates to the same string value shall write output to the existing stream. The stream shall remain open until the **close** function (see **Input/Output and General Functions** (on page 168)) is called with an expression that evaluates to the same string value. At that time, the stream shall be closed as if by a call to the *pclose()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001.

As described in detail by the grammar in **Grammar** (on page 170), these output statements shall take a comma-separated list of *expressions* referred to in the grammar by the non-terminal symbols **expr_list**, **print_expr_list**, or **print_expr_list_opt**. This list is referred to here as the *expression list*, and each member is referred to as an *expression argument*.

The **print** statement shall write the value of each expression argument onto the indicated output stream separated by the current output field separator (see variable **OFS** above), and terminated by the output record separator (see variable **ORS** above). All expression arguments shall be taken as strings, being converted if necessary; this conversion shall be as described in **Expressions in awk** (on page 156), with the exception that the **printf** format in **OFMT** shall be used instead of the value in **CONVFMT**. An empty expression list shall stand for the whole input record (\$0).

The **printf** statement shall produce output based on a notation similar to the File Format Notation used to describe file formats in this volume of IEEE Std 1003.1-2001 (see the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5, File Format Notation). Output shall be produced as specified with the first *expression* argument as the string *format* and subsequent *expression* arguments as the strings *arg1* to *argn*, inclusive, with the following exceptions:

- 1. The *format* shall be an actual character string rather than a graphical representation. Therefore, it cannot contain empty character positions. The <space> in the *format* string, in any context other than a *flag* of a conversion specification, shall be treated as an ordinary character that is copied to the output.
- 2. If the character set contains a ' Δ ' character and that character appears in the *format* string, it shall be treated as an ordinary character that is copied to the output.

- 3. The *escape sequences* beginning with a backslash character shall be treated as sequences of ordinary characters that are copied to the output. Note that these same sequences shall be interpreted lexically by *awk* when they appear in literal strings, but they shall not be treated specially by the **printf** statement.
- 4. A *field width* or *precision* can be specified as the '*' character instead of a digit string. In this case the next argument from the expression list shall be fetched and its numeric value taken as the field width or precision.
- 5. The implementation shall not precede or follow output from the d or u conversion specifier characters with

blank>s not specified by the *format* string.
- 6. The implementation shall not precede output from the o conversion specifier character with leading zeros not specified by the *format* string.
- 7. For the c conversion specifier character: if the argument has a numeric value, the character whose encoding is that value shall be output. If the value is zero or is not the encoding of any character in the character set, the behavior is undefined. If the argument does not have a numeric value, the first character of the string value shall be output; if the string does not contain any characters, the behavior is undefined.
- 8. For each conversion specification that consumes an argument, the next expression argument shall be evaluated. With the exception of the c conversion specifier character, the value shall be converted (according to the rules specified in **Expressions in awk** (on page 156)) to the appropriate type for the conversion specification.
- 9. If there are insufficient expression arguments to satisfy all the conversion specifications in the *format* string, the behavior is undefined.
- 10. If any character sequence in the *format* string begins with a '%' character, but does not form a valid conversion specification, the behavior is unspecified.

Both **print** and **printf** can output at least {LINE_MAX} bytes.

Functions

The awk language has a variety of built-in functions: arithmetic, string, input/output, and general.

Arithmetic Functions

The arithmetic functions, except for **int**, shall be based on the ISO C standard (see Section 1.7.2 (on page 7)). The behavior is undefined in cases where the ISO C standard specifies that an error be returned or that the behavior is undefined. Although the grammar (see **Grammar** (on page 170)) permits built-in functions to appear with no arguments or parentheses, unless the argument or parentheses are indicated as optional in the following list (by displaying them within the "[]" brackets), such use is undefined.

- **atan2**(y,x) Return arctangent of y/x in radians in the range $[-\pi,\pi]$.
- $\cos(x)$ Return cosine of x, where x is in radians.
 - sin(x) Return sine of x, where x is in radians.
- $\exp(x)$ Return the exponential function of x.
- $\log(x)$ Return the natural logarithm of x.
- $\mathbf{sqrt}(x)$ Return the square root of x.

int(x) Return the argument truncated to an integer. Truncation shall be toward 0 when x>0.

rand() Return a random number n, such that $0 \le n < 1$.

srand([*expr*]) Set the seed value for *rand* to *expr* or use the time of day if *expr* is omitted. The previous seed value shall be returned.

String Functions

The string functions in the following list shall be supported. Although the grammar (see **Grammar** (on page 170)) permits built-in functions to appear with no arguments or parentheses, unless the argument or parentheses are indicated as optional in the following list (by displaying them within the "[]" brackets), such use is undefined.

gsub(ere, repl[, in])

Behave like **sub** (see below), except that it shall replace all occurrences of the regular expression (like the *ed* utility global substitute) in \$0 or in the *in* argument, when specified.

- **index**(*s*, *t*) Return the position, in characters, numbering from 1, in string *s* where string *t* first occurs, or zero if it does not occur at all.
- **length**[([s])] Return the length, in characters, of its argument taken as a string, or of the whole record, \$0, if there is no argument.
- **match**(*s*, *ere*) Return the position, in characters, numbering from 1, in string *s* where the extended regular expression *ere* occurs, or zero if it does not occur at all. RSTART shall be set to the starting position (which is the same as the returned value), zero if no match is found; RLENGTH shall be set to the length of the matched string, –1 if no match is found.

split(*s*, *a*[, *fs*])

Split the string s into array elements a[1], a[2], ..., a[n], and return n. All elements of the array shall be deleted before the split is performed. The separation shall be done with the ERE fs or with the field separator fS if fs is not given. Each array element shall have a string value when created and, if appropriate, the array element shall be considered a numeric string (see **Expressions in awk** (on page 156)). The effect of a null string as the value of fs is unspecified.

sprintf(fmt, expr, expr, ...)

Format the expressions according to the **printf** format given by *fmt* and return the resulting string.

sub(ere, repl[, in])

Substitute the string *repl* in place of the first instance of the extended regular expression *ERE* in string *in* and return the number of substitutions. An ampersand ('&') appearing in the string *repl* shall be replaced by the string from *in* that matches the ERE. An ampersand preceded with a backslash ('\') shall be interpreted as the literal ampersand character. An occurrence of two consecutive backslashes shall be interpreted as just a single literal backslash character. Any other occurrence of a backslash (for example, preceding any other character) shall be treated as a literal backslash character. Note that if *repl* is a string literal (the lexical token **STRING**; see **Grammar** (on page 170)), the handling of the ampersand character occurs after any lexical processing, including any lexical backslash escape sequence processing. If *in* is specified and it is not an Ivalue (see **Expressions in awk** (on page 156)), the behavior is undefined. If *in* is omitted, *awk*

6487 shall use the current record (\$0) in its place. 6488 $\mathbf{substr}(s, m[, n])$ 6489 Return the at most *n*-character substring of s that begins at position m, numbering from 1. If n is omitted, or if n specifies more characters than are left in the string, 6490 the length of the substring shall be limited by the length of the string *s*. 6491 tolower(s) Return a string based on the string s. Each character in s that is an uppercase letter 6492 specified to have a tolower mapping by the LC_CTYPE category of the current 6493 locale shall be replaced in the returned string by the lowercase letter specified by 6494 the mapping. Other characters in s shall be unchanged in the returned string. 6495 toupper(s) Return a string based on the string s. Each character in s that is a lowercase letter 6496 specified to have a **toupper** mapping by the *LC_CTYPE* category of the current 6497 locale is replaced in the returned string by the uppercase letter specified by the 6498 mapping. Other characters in s are unchanged in the returned string. 6499 All of the preceding functions that take *ERE* as a parameter expect a pattern or a string valued 6500 expression that is a regular expression as defined in **Regular Expressions** (on page 161). 6501 **Input/Output and General Functions** 6502 The input/output and general functions are: 6503 6504 close (expression) Close the file or pipe opened by a **print** or **printf** statement or a call to **getline** with 6505 the same string-valued expression. The limit on the number of open expression 6506 arguments is implementation-defined. If the close was successful, the function 6507 shall return zero; otherwise, it shall return non-zero. 6508 6509 expression | getline [var] Read a record of input from a stream piped from the output of a command. The 6510 stream shall be created if no stream is currently open with the value of *expression* as 6511 6512 its command name. The stream created shall be equivalent to one created by a call to the popen() function with the value of expression as the command argument and a 6513 value of r as the *mode* argument. As long as the stream remains open, subsequent calls in which expression evaluates to the same string value shall read subsequent 6515 records from the stream. The stream shall remain open until the close function is 6516 called with an expression that evaluates to the same string value. At that time, the 6517 stream shall be closed as if by a call to the pclose() function. If var is omitted, \$0 6518 and NF shall be set; otherwise, var shall be set and, if appropriate, it shall be 6519 6520 considered a numeric string (see Expressions in awk (on page 156)). The getline operator can form ambiguous constructs when there are 6521 unparenthesized operators (including concatenate) to the left of the '|' (to the 6522 beginning of the expression containing **getline**). In the context of the '\$' 6523 operator, ' | ' shall behave as if it had a lower precedence than '\$'. The result of 6524 evaluating other operators is unspecified, and conforming applications shall 6525 parenthesize properly all such usages. 6526 getline Set \$0 to the next input record from the current input file. This form of getline shall 6527 set the NF, NR, and FNR variables. 6528 getline var Set variable var to the next input record from the current input file and, if 6529 appropriate, var shall be considered a numeric string (see Expressions in awk (on 6530

6531

page 156)). This form of **getline** shall set the **FNR** and **NR** variables.

getline [var] < expression

 Read the next record of input from a named file. The *expression* shall be evaluated to produce a string that is used as a pathname. If the file of that name is not currently open, it shall be opened. As long as the stream remains open, subsequent calls in which *expression* evaluates to the same string value shall read subsequent records from the file. The file shall remain open until the **close** function is called with an expression that evaluates to the same string value. If *var* is omitted, \$0 and **NF** shall be set; otherwise, *var* shall be set and, if appropriate, it shall be considered a numeric string (see **Expressions in awk** (on page 156)).

The **getline** operator can form ambiguous constructs when there are unparenthesized binary operators (including concatenate) to the right of the '<' (up to the end of the expression containing the **getline**). The result of evaluating such a construct is unspecified, and conforming applications shall parenthesize properly all such usages.

system(expression)

Execute the command given by *expression* in a manner equivalent to the *system()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001 and return the exit status of the command.

All forms of **getline** shall return 1 for successful input, zero for end-of-file, and -1 for an error.

Where strings are used as the name of a file or pipeline, the application shall ensure that the strings are textually identical. The terminology "same string value" implies that "equivalent strings", even those that differ only by <space>s, represent different files.

User-Defined Functions

The awk language also provides user-defined functions. Such functions can be defined as:

```
function name([parameter, ...]) { statements }
```

A function can be referred to anywhere in an *awk* program; in particular, its use can precede its definition. The scope of a function is global.

Function parameters, if present, can be either scalars or arrays; the behavior is undefined if an array name is passed as a parameter that the function uses as a scalar, or if a scalar expression is passed as a parameter that the function uses as an array. Function parameters shall be passed by value if scalar and by reference if array name.

The number of parameters in the function definition need not match the number of parameters in the function call. Excess formal parameters can be used as local variables. If fewer arguments are supplied in a function call than are in the function definition, the extra parameters that are used in the function body as scalars shall evaluate to the uninitialized value until they are otherwise initialized, and the extra parameters that are used in the function body as arrays shall be treated as uninitialized arrays where each element evaluates to the uninitialized value until otherwise initialized.

When invoking a function, no white space can be placed between the function name and the opening parenthesis. Function calls can be nested and recursive calls can be made upon functions. Upon return from any nested or recursive function call, the values of all of the calling function's parameters shall be unchanged, except for array parameters passed by reference. The **return** statement can be used to return a value. If a **return** statement appears outside of a function definition, the behavior is undefined.

In the function definition, <newline>s shall be optional before the opening brace and after the closing brace. Function definitions can appear anywhere in the program where a *pattern-action*

6578 pair is allowed.

6579

6580

6581

6582 6583

6584

Grammar

The grammar in this section and the lexical conventions in the following section shall together describe the syntax for *awk* programs. The general conventions for this style of grammar are described in Section 1.10 (on page 19). A valid program can be represented as the non-terminal symbol *program* in the grammar. This formal syntax shall take precedence over the preceding text syntax description.

```
%token NAME NUMBER STRING ERE
6585
                                /* Name followed by '(' without white space. */
6586
            %token FUNC NAME
            /* Keywords
                          * /
6587
            %token
                          Begin
6588
                                   End
                         'BEGIN'
                                 'END'
6589
            %token
                          Break
                                   Continue
                                               Delete
                                                         Do
                                                              Else
6590
                         'break' 'continue' 'delete' 'do' 'else'
6591
6592
            %token
                          Exit
                                 For
                                        Function
                                                    Ιf
                                                          Tn
6593
                         'exit' 'for' 'function' 'if' 'in'
6594
            %token
                          Next
                                  Print
                                          Printf
                                                    Return
6595
                         'next' 'print' 'printf' 'return' 'while' */
6596
            /* Reserved function names */
            %token BUILTIN FUNC NAME
6597
                         /* One token for the following:
6598
                          * atan2 cos sin exp log sqrt int rand srand
6599
6600
                          * gsub index length match split sprintf sub
6601
                          * substr tolower toupper close system
6602
6603
            %token GETLINE
                         /* Syntactically different from other built-ins. */
6604
6605
            /* Two-character tokens. */
            token ADD ASSIGN SUB ASSIGN MUL ASSIGN DIV ASSIGN MOD ASSIGN POW ASSIGN
6606
            /*
                   ' +='
                                                        '/='
                                                                    ′%=′
6607
6608
            %token OR
                        AND NO MATCH
                                          ΕQ
                                               _{
m LE}
                                                     GΕ
                                                           NE
                                                                INCR DECR
                                                                             APPEND
                   ' | | ' ' &&' '!~' '==' '<=' '>=' '!=' '++'
                                                                '--'
6609
                                                                      ' >> '
                                                                               * /
            /* One-character tokens. */
6610
            %token '{' '}' '(' ')' '[' ']' ',' ';' NEWLINE
6611
            %token '+' '-' '*' '%' '^' '!' '>' '<' '|' '?' ':' '~' '$' '='
6612
6613
            %start program
6614
6615
           program
                              : item list
6616
                                actionless_item_list
6617
           item list
                              : newline opt
6618
6619
                              actionless item list item terminator
6620
                                item list
                                                        item terminator
6621
                                item list
                                                    action terminator
6622
```

```
6623
           actionless item list : item list
                                                           pattern terminator
6624
                              actionless_item_list pattern terminator
6625
                              : pattern action
6626
           item
                                                     '(' param_list_opt ')'
6627
                                Function NAME
                                     newline opt action
6628
                                Function FUNC_NAME '(' param_list_opt ')'
6629
6630
                                    newline opt action
6631
6632
           param list opt
                              : /* empty */
6633
                              param list
6634
6635
           param list
                              : NAME
                                param list ',' NAME
6636
6637
                              : Begin
6638
           pattern
                                End
6639
6640
                                expr
                                expr ',' newline opt expr
6641
6642
                                                                                 1 } 1
6643
           action
                              : '{' newline opt
                                '{' newline opt terminated statement list
                                                                                 1 } 1
6644
                                '{' newline opt unterminated statement list '}'
6645
6646
                              : terminator ';'
6647
           terminator
                                terminator NEWLINE
6648
                                            ';'
6649
6650
                                            NEWLINE
6651
6652
           terminated statement list : terminated statement
                              | terminated_statement_list terminated_statement
6653
6654
6655
           unterminated statement list : unterminated statement
6656
                                terminated statement list unterminated statement
6657
           terminated_statement : action newline_opt
6658
                               If '(' expr ')' newline opt terminated statement
6659
                                If '(' expr ')' newline_opt terminated_statement
6660
                                    Else newline opt terminated statement
6661
                               While '(' expr ')' newline opt terminated statement
6662
                              | For '(' simple statement opt ';'
6663
                                   expr opt ';' simple statement opt ')' newline opt
6664
                                   terminated statement
6665
                              For '(' NAME In NAME ')' newline opt
6666
6667
                                   terminated statement
                                ';' newline opt
6668
6669
                                terminatable_statement NEWLINE newline_opt
6670
                                terminatable statement ';'
                                                                  newline opt
```

```
6671
           unterminated statement : terminatable statement
6672
6673
                               If '(' expr ')' newline opt unterminated statement
                              | If '(' expr ')' newline opt terminated statement
6674
                                    Else newline opt unterminated statement
6675
                                While '(' expr ')' newline opt unterminated statement
6676
6677
                              | For '(' simple_statement_opt ';'
6678
                               expr opt ';' simple statement opt ')' newline opt
                                   unterminated statement
6679
6680
                               For '(' NAME In NAME ')' newline opt
6681
                                   unterminated statement
6682
6683
           terminatable statement : simple statement
6684
                                Break
6685
                                Continue
                                Next
6686
6687
                                Exit expr opt
                                Return expr opt
6688
                                Do newline opt terminated statement While '(' expr ')'
6689
6690
           simple statement opt : /* empty */
6691
6692
                              | simple statement
6693
6694
           simple statement : Delete NAME '[' expr list ']'
6695
                                expr
                                print statement
6696
6697
6698
           print statement
                              : simple print statement
                                simple print statement output redirection
6699
6700
           simple print statement : Print print expr list opt
6701
6702
                               Print '(' multiple_expr_list ')'
6703
                               Printf print expr list
6704
                                Printf '(' multiple expr list ')'
6705
           output redirection : '>'
6706
6707
                                APPEND expr
                                ′ ′
6708
                                        expr
6709
6710
           expr list opt
                              : /* empty */
6711
                              expr list
6712
6713
           expr list
                              : expr
6714
                                multiple expr list
6715
6716
           multiple_expr_list : expr ',' newline_opt expr
                              | multiple expr list ',' newline opt expr
6717
```

```
6718
6719
                               : /* empty */
            expr opt
6720
                                expr
6721
6722
                               : unary expr
            expr
6723
                                non_unary_expr
6724
6725
                                '+' expr
            unary expr
6726
                                 '-' expr
                                 unary_expr '^'
6727
                                                        expr
6728
                                 unary expr '*'
                                                        expr
6729
                                 unary expr '/'
                                                        expr
6730
                                 unary_expr '%'
                                                        expr
6731
                                 unary expr '+'
                                                        expr
6732
                                 unary expr '-'
                                                       expr
6733
                                 unary expr
                                                       non unary expr
6734
                                 unary_expr '<'
                                                       expr
6735
                                 unary expr LE
                                                       expr
6736
                                 unary expr NE
                                                        expr
6737
                                 unary expr EQ
                                                        expr
6738
                                 unary_expr '>'
                                                        expr
6739
                                 unary expr GE
                                                        expr
                                 unary expr '~'
6740
                                                        expr
6741
                                 unary expr NO MATCH expr
                                 unary_expr In NAME
6742
6743
                                 unary_expr AND newline_opt expr
6744
                                 unary expr OR newline opt expr
6745
                                 unary_expr '?' expr ':' expr
6746
                                 unary input function
6747
            non_unary_expr
                                '(' expr ')'
6748
                                 '!' expr
6749
                                 non_unary_expr '^'
6750
                                                            expr
6751
                                 non unary expr '*'
                                                            expr
6752
                                 non_unary_expr '/'
                                                            expr
6753
                                 non unary expr '%'
                                                            expr
6754
                                 non unary expr '+'
                                                            expr
6755
                                 non_unary_expr '-'
                                                            expr
6756
                                 non unary expr
                                                            non unary expr
                                 non_unary_expr '<'</pre>
6757
                                                            expr
6758
                                 non unary expr LE
                                                            expr
6759
                                 non_unary_expr NE
                                                            expr
6760
                                 non unary expr EQ
                                                            expr
6761
                                 non_unary_expr '>'
                                                            expr
6762
                                 non unary expr GE
                                                            expr
                                 non_unary_expr '~'
6763
                                                            expr
6764
                                 non unary expr NO MATCH expr
6765
                                 non unary expr In NAME
6766
                                 '(' multiple expr list ')' In NAME
6767
                                 non unary expr AND newline opt expr
```

```
6768
                                non unary expr OR newline opt expr
6769
                                non unary expr '?' expr ':' expr
                                NUMBER
6770
6771
                                STRING
6772
                                lvalue
6773
                                ERE
6774
                                lvalue INCR
                                lvalue DECR
6775
                                INCR lvalue
6776
                                DECR lvalue
6777
6778
                                lvalue POW ASSIGN expr
6779
                                lvalue MOD ASSIGN expr
6780
                                lvalue MUL ASSIGN expr
                                lvalue DIV ASSIGN expr
6781
                                lvalue ADD ASSIGN expr
6782
6783
                                lvalue SUB ASSIGN expr
                                lvalue '=' expr
6784
                                FUNC_NAME '(' expr_list_opt ')'
6785
                                    /* no white space allowed before '(' */
6786
                                BUILTIN FUNC NAME '(' expr list opt ')'
6787
                                BUILTIN FUNC NAME
6788
6789
                                non_unary_input_function
6790
           print expr list opt : /* empty */
6791
6792
                               print expr list
6793
6794
           print expr list
                              : print expr
6795
                                print_expr_list ',' newline_opt print_expr
6796
6797
           print expr
                              : unary print expr
6798
                                non unary print expr
6799
6800
           unary_print_expr : '+' print_expr
6801
                                '-' print expr
                                unary_print_expr '^'
6802
                                                             print expr
6803
                                unary print expr '*'
                                                             print expr
                                unary_print_expr '/'
                                                             print expr
6804
                                unary print expr '%'
6805
                                                             print expr
6806
                                unary_print_expr '+'
                                                             print expr
6807
                                unary print expr '-'
                                                             print expr
                                unary_print_expr
6808
                                                             non unary print expr
                                unary_print_expr '~'
6809
                                                             print expr
6810
                                unary print expr NO MATCH print expr
6811
                                unary_print_expr In NAME
6812
                                unary print expr AND newline opt print expr
6813
                                unary print expr OR newline opt print expr
6814
                                unary print expr '?' print expr ':' print expr
6815
           non_unary_print_expr : '(' expr ')'
6816
6817
                              '!' print expr
```

```
non_unary print expr '^'
6818
                                                                 print expr
6819
                                non_unary_print_expr '*'
                                                                 print_expr
                                non unary print expr '/'
6820
                                                                 print expr
                                non unary print expr '%'
6821
                                                                 print expr
6822
                                non unary print expr '+'
                                                                 print expr
6823
                                non unary print expr '-'
                                                                 print expr
6824
                                non_unary_print_expr
                                                                 non_unary_print_expr
                                non_unary_print expr '~'
6825
                                                                 print expr
6826
                                non unary print expr NO MATCH print expr
                                non unary print expr In NAME
6827
6828
                                '(' multiple expr list ')' In NAME
6829
                                non unary print expr AND newline opt print expr
6830
                                non_unary_print_expr OR newline_opt print_expr
6831
                                non unary print expr '?' print expr ':' print expr
                                NUMBER
6832
6833
                                STRING
                                lvalue
6834
                                ERE
6835
                                lvalue INCR
6836
                                lvalue DECR
6837
                                INCR lvalue
6838
                                DECR lvalue
6839
                                lvalue POW ASSIGN print expr
6840
6841
                                lvalue MOD_ASSIGN print_expr
6842
                                lvalue MUL ASSIGN print expr
6843
                                lvalue DIV ASSIGN print expr
6844
                                lvalue ADD ASSIGN print expr
                                lvalue SUB ASSIGN print expr
6845
                                lvalue '=' print_expr
6846
                                FUNC NAME '(' expr list opt ')'
6847
6848
                                   /* no white space allowed before '(' */
                                BUILTIN FUNC NAME '(' expr list opt ')'
6849
6850
                                BUILTIN FUNC NAME
6851
6852
           lvalue
                              : NAME
                                NAME '[' expr list ']'
6853
6854
                                '$' expr
6855
           non unary input function : simple get
6856
                               simple get '<' expr
6857
                                non unary expr '|' simple get
6858
6859
           unary input function : unary expr '|' simple get
6860
6861
6862
            simple get
                              : GETLINE
                                GETLINE lvalue
6863
6864
6865
           newline opt
                              : /* empty */
6866
                                newline_opt NEWLINE
6867
```

This grammar has several ambiguities that shall be resolved as follows:

- Operator precedence and associativity shall be as described in Table 4-1 (on page 156).
- In case of ambiguity, an **else** shall be associated with the most immediately preceding **if** that would satisfy the grammar.
- In some contexts, a slash ('/') that is used to surround an ERE could also be the division operator. This shall be resolved in such a way that wherever the division operator could appear, a slash is assumed to be the division operator. (There is no unary division operator.)

One convention that might not be obvious from the formal grammar is where <newline>s are acceptable. There are several obvious placements such as terminating a statement, and a backslash can be used to escape <newline>s between any lexical tokens. In addition, <newline>s without backslashes can follow a comma, an open brace, logical AND operator ("&&"), logical OR operator (" | | "), the **do** keyword, the **else** keyword, and the closing parenthesis of an **if**, **for**, or **while** statement. For example:

```
{ print $1, $2 }
```

Lexical Conventions

The lexical conventions for *awk* programs, with respect to the preceding grammar, shall be as follows:

- 1. Except as noted, *awk* shall recognize the longest possible token or delimiter beginning at a given point.
- 2. A comment shall consist of any characters beginning with the number sign character and terminated by, but excluding the next occurrence of, a <newline>. Comments shall have no effect, except to delimit lexical tokens.
- The <newline> shall be recognized as the token NEWLINE.
- 4. A backslash character immediately followed by a <newline> shall have no effect.
- 5. The token **STRING** shall represent a string constant. A string constant shall begin with the character '"'. Within a string constant, a backslash character shall be considered to begin an escape sequence as specified in the table in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5, File Format Notation ('\\', '\a', '\b', '\f', '\n', '\r', '\t', '\t', '\v'). In addition, the escape sequences in Table 4-2 (on page 162) shall be recognized. A <newline> shall not occur within a string constant. A string constant shall be terminated by the first unescaped occurrence of the character '"' after the one that begins the string constant. The value of the string shall be the sequence of all unescaped characters and values of escape sequences between, but not including, the two delimiting '"' characters.
- 6. The token **ERE** represents an extended regular expression constant. An ERE constant shall begin with the slash character. Within an ERE constant, a backslash character shall be considered to begin an escape sequence as specified in the table in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5, File Format Notation. In addition, the escape sequences in Table 4-2 (on page 162) shall be recognized. The application shall ensure that a <newline> does not occur within an ERE constant. An ERE constant shall be terminated by the first unescaped occurrence of the slash character after the one that begins the ERE constant. The extended regular expression represented by the ERE constant shall be the sequence of all unescaped characters and values of escape sequences between, but not including, the two delimiting slash characters.

A <blank> shall have no effect, except to delimit lexical tokens or within STRING or ERE tokens.

- 8. The token **NUMBER** shall represent a numeric constant. Its form and numeric value shall be equivalent to either of the tokens **floating-constant** or **integer-constant** as specified by the ISO C standard, with the following exceptions:
 - a. An integer constant cannot begin with 0x or include the hexadecimal digits 'a', 'b', 'c', 'd', 'e', 'f', 'A', 'B', 'C', 'D', 'E', or 'F'.
 - b. The value of an integer constant beginning with 0 shall be taken in decimal rather than octal.
 - c. An integer constant cannot include a suffix ('u', 'U', 'l', or 'L').
 - d. A floating constant cannot include a suffix ('f', 'F', 'l', or 'L').

If the value is too large or too small to be representable (see Section 1.7.2 (on page 7)), the behavior is undefined.

- 9. A sequence of underscores, digits, and alphabetics from the portable character set (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.1, Portable Character Set), beginning with an underscore or alphabetic, shall be considered a word.
- 10. The following words are keywords that shall be recognized as individual tokens; the name of the token is the same as the keyword:

BEGIN	delete	END	function	in	printf
break	do	exit	getline	next	return
continue	else	for	if	print	while

11. The following words are names of built-in functions and shall be recognized as the token **BUILTIN_FUNC_NAME**:

atan2	gsub	log	split	sub	toupper
close	index	match	sprintf	substr	
cos	int	rand	sqrt	system	
exp	length	sin	srand	tolower	

The above-listed keywords and names of built-in functions are considered reserved words.

- 12. The token **NAME** shall consist of a word that is not a keyword or a name of a built-in function and is not followed immediately (without any delimiters) by the ' (' character.
- 13. The token **FUNC_NAME** shall consist of a word that is not a keyword or a name of a built-in function, followed immediately (without any delimiters) by the '(' character. The '(' character shall not be included as part of the token.
- 14. The following two-character sequences shall be recognized as the named tokens:

Token Name	Sequence	Token Name	Sequence
ADD_ASSIGN	+=	NO_MATCH	!~
SUB_ASSIGN	-=	EQ	==
MUL_ASSIGN	*=	LE	<=
DIV_ASSIGN	/=	GE	>=
MOD_ASSIGN	%=	NE	! =
POW_ASSIGN	^=	INCR	++
OR		DECR	
AND	&&	APPEND	>>

15. The following single characters shall be recognized as tokens whose names are the character:

```
<newline> { } ( ) [ ] , ; + - * % ^ ! > < | ? : ~ $ =
```

There is a lexical ambiguity between the token **ERE** and the tokens '/' and **DIV_ASSIGN**. When an input sequence begins with a slash character in any syntactic context where the token '/' or **DIV_ASSIGN** could appear as the next token in a valid program, the longer of those two tokens that can be recognized shall be recognized. In any other syntactic context where the token **ERE** could appear as the next token in a valid program, the token **ERE** shall be recognized.

EXIT STATUS

6956

6957

6958

6959

6960

6961

6962

6963

6964 6965

6966

6967

6968

6970

6971

6972

6973 6974

6975

6976

6977

6978

6979

6980

6982

6983

6984

6985

6986 6987

6988

6989

6990 6991

6992

6993

6994

6995 6996 The following exit values shall be returned:

- 0 All input files were processed successfully.
- >0 An error occurred.

The exit status can be altered within the program by using an **exit** expression.

6969 CONSEQUENCES OF ERRORS

If any *file* operand is specified and the named file cannot be accessed, *awk* shall write a diagnostic message to standard error and terminate without any further action.

If the program specified by either the *program* operand or a *progfile* operand is not a valid *awk* program (as specified in the EXTENDED DESCRIPTION section), the behavior is undefined.

APPLICATION USAGE

The **index**, **length**, **match**, and **substr** functions should not be confused with similar functions in the ISO C standard; the *awk* versions deal with characters, while the ISO C standard deals with bytes.

Because the concatenation operation is represented by adjacent expressions rather than an explicit operator, it is often necessary to use parentheses to enforce the proper evaluation precedence.

6981 EXAMPLES

The *awk* program specified in the command line is most easily specified within single-quotes (for example, '*program*') for applications using *sh*, because *awk* programs commonly contain characters that are special to the shell, including double-quotes. In the cases where an *awk* program contains single-quote characters, it is usually easiest to specify most of the program as strings within single-quotes concatenated by the shell with quoted single-quote characters. For example:

```
awk '/'\''/ { print "quote:", $0 }'
```

prints all lines from the standard input containing a single-quote character, prefixed with quote:.

The following are examples of simple *awk* programs:

1. Write to the standard output all input lines for which field 3 is greater than 5:

```
$3 > 5
```

2. Write every tenth line:

```
(NR % 10) == 0
```

3. Write any line with a substring matching the regular expression:

```
/(G|D)(2[0-9][[:alpha:]]*)/
```

Utilities awk

4. Print any line with a substring containing a 'G' or 'D', followed by a sequence of digits and characters. This example uses character classes **digit** and **alpha** to match language-independent digit and alphabetic characters respectively:

```
/(G|D)([[:digit:][:alpha:]]*)/
```

5. Write any line in which the second field matches the regular expression and the fourth field does not:

```
$2 ~ /xyz/ && $4 !~ /xyz/
```

6. Write any line in which the second field contains a backslash:

```
$2 ~ /\\/
```

6997

6998

6999

7000

7002

7003

7004

7005

7006

7007

7008

7009

70107011

7012

7013

7014

7015

7016 7017

7018

7019

7020

7021

70227023

7024

7025

7026

7027

7028

7029

7030

7031

7032

7. Write any line in which the second field contains a backslash. Note that backslash escapes are interpreted twice; once in lexical processing of the string and once in processing the regular expression:

```
$2 ~ "\\\"
```

8. Write the second to the last and the last field in each line. Separate the fields by a colon:

```
{OFS=":";print $(NF-1), $NF}
```

9. Write the line number and number of fields in each line. The three strings representing the line number, the colon, and the number of fields are concatenated and that string is written to standard output:

```
{print NR ":" NF}
```

10. Write lines longer than 72 characters:

```
length(\$0) > 72
```

11. Write the first two fields in opposite order separated by **OFS**:

```
{ print $2, $1 }
```

12. Same, with input fields separated by a comma or <space>s and <tab>s, or both:

```
BEGIN { FS = ", [ \t] * | [ \t] + " } 
 { print $2, $1 }
```

13. Add up the first column, print sum, and average:

```
\{s += \$1 \} END \{print "sum is ", s, " average is", s/NR\}
```

14. Write fields in reverse order, one per line (many lines out for each line in):

```
{ for (i = NF; i > 0; --i) print $i }
```

15. Write all lines between occurrences of the strings **start** and **stop**:

```
/start/, /stop/
```

16. Write all lines whose first field is different from the previous one:

```
$1 != prev { print; prev = $1 }
```

17. Simulate *echo*:

```
7033 BEGIN {
7034 for (i = 1; i < ARGC; ++i)
7035 printf("%s%s", ARGV[i], i==ARGC-1?"\n":" ")
```

awk Utilities

```
7036
                   }
7037
              18. Write the path prefixes contained in the PATH environment variable, one per line:
                   BEGIN
7038
7039
                             n = split (ENVIRON["PATH"], path, ":")
                             for (i = 1; i <= n; ++i)
7040
                             print path[i]
7041
7042
              19. If there is a file named input containing page headers of the form:
7043
7044
                      Page #
                   and a file named program that contains:
7045
                               \{ \$2 = n++; \}
7046
                               { print }
7047
                   then the command line:
7048
                   awk -f program n=5 input
7049
```

prints the file **input**, filling in page numbers starting at 5.

7051 RATIONALE

7050

7052

7053

7054

7055

7056

7057

7058

7059

7060 7061

7062

7063

7064

7065

7066

7067

7068

7069

7070

7071

7072

7073 7074

7075

This description is based on the new awk, "nawk", (see the referenced *The AWK Programming Language*), which introduced a number of new features to the historical awk:

- 1. New keywords: delete, do, function, return
- 2. New built-in functions: atan2, close, cos, gsub, match, rand, sin, srand, sub, system
- 3. New predefined variables: FNR, ARGC, ARGV, RSTART, RLENGTH, SUBSEP
- 4. New expression operators: ?, :, ,, ^
- 5. The **FS** variable and the third argument to **split**, now treated as extended regular expressions.
- 6. The operator precedence, changed to more closely match the C language. Two examples of code that operate differently are:

```
while (n /= 10 > 1) ...
if (!"wk" \sim /bwk/) ...
```

Several features have been added based on newer implementations of awk:

- Multiple instances of –**f** *progfile* are permitted.
- The new option –v assignment.
- The new predefined variable ENVIRON.
- New built-in functions toupper and tolower.
- More formatting capabilities are added to **printf** to match the ISO C standard.

The overall *awk* syntax has always been based on the C language, with a few features from the shell command language and other sources. Because of this, it is not completely compatible with any other language, which has caused confusion for some users. It is not the intent of the standard developers to address such issues. A few relatively minor changes toward making the language more compatible with the ISO C standard were made; most of these changes are based on similar changes in recent implementations, as described above. There remain several C-

Utilities awk

language conventions that are not in *awk*. One of the notable ones is the comma operator, which is commonly used to specify multiple expressions in the C language **for** statement. Also, there are various places where *awk* is more restrictive than the C language regarding the type of expression that can be used in a given context. These limitations are due to the different features that the *awk* language does provide.

Regular expressions in awk have been extended somewhat from historical implementations to make them a pure superset of extended regular expressions, as defined by IEEE Std 1003.1-2001 (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.4, Extended Regular Expressions). The main extensions are internationalization features and interval expressions. Historical implementations of awk have long supported backslash escape sequences as an extension to extended regular expressions, and this extension has been retained despite inconsistency with other utilities. The number of escape sequences recognized in both extended regular expressions and strings has varied (generally increasing with time) among implementations. The set specified by IEEE Std 1003.1-2001 includes most sequences known to be supported by popular implementations and by the ISO C standard. One sequence that is not supported is hexadecimal value escapes beginning with '\x'. This would allow values expressed in more than 9 bits to be used within awk as in the ISO C standard. However, because this syntax has a non-deterministic length, it does not permit the subsequent character to be a hexadecimal digit. This limitation can be dealt with in the C language by the use of lexical string concatenation. In the awk language, concatenation could also be a solution for strings, but not for extended regular expressions (either lexical ERE tokens or strings used dynamically as regular expressions). Because of this limitation, the feature has not been added to IEEE Std 1003.1-2001.

When a string variable is used in a context where an extended regular expression normally appears (where the lexical token ERE is used in the grammar) the string does not contain the literal slashes.

Some versions of awk allow the form:

```
func name(args, ...) { statements }
```

This has been deprecated by the authors of the language, who asked that it not be specified.

Historical implementations of *awk* produce an error if a **next** statement is executed in a **BEGIN** action, and cause *awk* to terminate if a **next** statement is executed in an **END** action. This behavior has not been documented, and it was not believed that it was necessary to standardize it

The specification of conversions between string and numeric values is much more detailed than in the documentation of historical implementations or in the referenced *The AWK Programming Language*. Although most of the behavior is designed to be intuitive, the details are necessary to ensure compatible behavior from different implementations. This is especially important in relational expressions since the types of the operands determine whether a string or numeric comparison is performed. From the perspective of an application writer, it is usually sufficient to expect intuitive behavior and to force conversions (by adding zero or concatenating a null string) when the type of an expression does not obviously match what is needed. The intent has been to specify historical practice in almost all cases. The one exception is that, in historical implementations, variables and constants maintain both string and numeric values after their original value is converted by any use. This means that referencing a variable or constant can have unexpected side effects. For example, with historical implementations the following program:

awk Utilities

```
7124 if (NR % 2)

7125 c = a + b

7126 if (a == b)

7127 print "numeric comparison"

7128 else

7129 print "string comparison"

7130 }
```

would perform a numeric comparison (and output numeric comparison) for each odd-numbered line, but perform a string comparison (and output string comparison) for each even-numbered line. IEEE Std 1003.1-2001 ensures that comparisons will be numeric if necessary. With historical implementations, the following program:

```
BEGIN {
    OFMT = "%e"
    print 3.14
    OFMT = "%f"
    print 3.14
}
```

would output "3.140000e+00" twice, because in the second **print** statement the constant "3.14" would have a string value from the previous conversion. IEEE Std 1003.1-2001 requires that the output of the second **print** statement be "3.140000". The behavior of historical implementations was seen as too unintuitive and unpredictable.

It was pointed out that with the rules contained in early drafts, the following script would print nothing:

```
BEGIN {
    y[1.5] = 1
    OFMT = "%e"
    print y[1.5]
}
```

Therefore, a new variable, **CONVFMT**, was introduced. The **OFMT** variable is now restricted to affecting output conversions of numbers to strings and **CONVFMT** is used for internal conversions, such as comparisons or array indexing. The default value is the same as that for **OFMT**, so unless a program changes **CONVFMT** (which no historical program would do), it will receive the historical behavior associated with internal string conversions.

The POSIX *awk* lexical and syntactic conventions are specified more formally than in other sources. Again the intent has been to specify historical practice. One convention that may not be obvious from the formal grammar as in other verbal descriptions is where <newline>s are acceptable. There are several obvious placements such as terminating a statement, and a backslash can be used to escape <newline>s between any lexical tokens. In addition, <newline>s without backslashes can follow a comma, an open brace, a logical AND operator ("&&"), a logical OR operator ("| | | "), the **do** keyword, the **else** keyword, and the closing parenthesis of an **if**, **for**, or **while** statement. For example:

```
{ print $1, $2
```

The requirement that *awk* add a trailing <newline> to the program argument text is to simplify the grammar, making it match a text file in form. There is no way for an application or test suite to determine whether a literal <newline> is added or whether *awk* simply acts as if it did.

Utilities awk

IEEE Std 1003.1-2001 requires several changes from historical implementations in order to support internationalization. Probably the most subtle of these is the use of the decimal-point character, defined by the *LC_NUMERIC* category of the locale, in representations of floating-point numbers. This locale-specific character is used in recognizing numeric input, in converting between strings and numeric values, and in formatting output. However, regardless of locale, the period character (the decimal-point character of the POSIX locale) is the decimal-point character recognized in processing *awk* programs (including assignments in command line arguments). This is essentially the same convention as the one used in the ISO C standard. The difference is that the C language includes the *setlocale()* function, which permits an application to modify its locale. Because of this capability, a C application begins executing with its locale set to the C locale, and only executes in the environment-specified locale after an explicit call to *setlocale()*. However, adding such an elaborate new feature to the *awk* language was seen as inappropriate for IEEE Std 1003.1-2001. It is possible to execute an *awk* program explicitly in any desired locale by setting the environment in the shell.

The undefined behavior resulting from NULs in extended regular expressions allows future extensions for the GNU *gawk* program to process binary data.

The behavior in the case of invalid *awk* programs (including lexical, syntactic, and semantic errors) is undefined because it was considered overly limiting on implementations to specify. In most cases such errors can be expected to produce a diagnostic and a non-zero exit status. However, some implementations may choose to extend the language in ways that make use of certain invalid constructs. Other invalid constructs might be deemed worthy of a warning, but otherwise cause some reasonable behavior. Still other constructs may be very difficult to detect in some implementations. Also, different implementations might detect a given error during an initial parsing of the program (before reading any input files) while others might detect it when executing the program after reading some input. Implementors should be aware that diagnosing errors as early as possible and producing useful diagnostics can ease debugging of applications, and thus make an implementation more usable.

The unspecified behavior from using multi-character **RS** values is to allow possible future extensions based on extended regular expressions used for record separators. Historical implementations take the first character of the string and ignore the others.

Unspecified behavior when *split*(*string*,*array*,<null>) is used is to allow a proposed future extension that would split up a string into an array of individual characters.

In the context of the **getline** function, equally good arguments for different precedences of the | and < operators can be made. Historical practice has been that:

```
qetline < "a" "b"</pre>
7204
7205
              is parsed as:
              ( getline < "a" ) "b"
7206
              although many would argue that the intent was that the file ab should be read. However:
7207
7208
              getline < "x" + 1
7209
              parses as:
              getline < ( "x" + 1 )
7210
              Similar problems occur with the | version of getline, particularly in combination with $. For
7211
7212
              example:
```

\$"echo hi" | getline

awk Utilities

(This situation is particularly problematic when used in a **print** statement, where the |**getline** part might be a redirection of the **print**.)

Since in most cases such constructs are not (or at least should not) be used (because they have a natural ambiguity for which there is no conventional parsing), the meaning of these constructs has been made explicitly unspecified. (The effect is that a conforming application that runs into the problem must parenthesize to resolve the ambiguity.) There appeared to be few if any actual uses of such constructs.

Grammars can be written that would cause an error under these circumstances. Where backwards-compatibility is not a large consideration, implementors may wish to use such grammars.

Some historical implementations have allowed some built-in functions to be called without an argument list, the result being a default argument list chosen in some "reasonable" way. Use of length as a synonym for length(\$0) is the only one of these forms that is thought to be widely known or widely used; this particular form is documented in various places (for example, most historical awk reference pages, although not in the referenced The AWK Programming Language) as legitimate practice. With this exception, default argument lists have always been undocumented and vaguely defined, and it is not at all clear how (or if) they should be generalized to user-defined functions. They add no useful functionality and preclude possible future extensions that might need to name functions without calling them. Not standardizing them seems the simplest course. The standard developers considered that length merited special treatment, however, since it has been documented in the past and sees possibly substantial use in historical programs. Accordingly, this usage has been made legitimate, but Issue 5 removed the obsolescent marking for XSI-conforming implementations and many otherwise conforming applications depend on this feature.

In **sub** and **gsub**, if *repl* is a string literal (the lexical token **STRING**), then two consecutive backslash characters should be used in the string to ensure a single backslash will precede the ampersand when the resultant string is passed to the function. (For example, to specify one literal ampersand in the replacement string, use **gsub(ERE**, " $\$ \&").)

Historically the only special character in the *repl* argument of **sub** and **gsub** string functions was the ampersand ($^{\prime}$ & $^{\prime}$) character and preceding it with the backslash character was used to turn off its special meaning.

The description in the ISO POSIX-2: 1993 standard introduced behavior such that the backslash character was another special character and it was unspecified whether there were any other special characters. This description introduced several portability problems, some of which are described below, and so it has been replaced with the more historical description. Some of the problems include:

- Historically, to create the replacement string, a script could use <code>gsub(ERE, "\\&")</code>, but with the ISO POSIX-2: 1993 standard wording, it was necessary to use <code>gsub(ERE, "\\\&")</code>. Backslash characters are doubled here because all string literals are subject to lexical analysis, which would reduce each pair of backslash characters to a single backslash before being passed to <code>gsub</code>.
- Since it was unspecified what the special characters were, for portable scripts to guarantee that characters are printed literally, each character had to be preceded with a backslash. (For example, a portable script had to use **gsub(ERE**, "\\h\\i") to produce a replacement string of "hi".)

The description for comparisons in the ISO POSIX-2:1993 standard did not properly describe historical practice because of the way numeric strings are compared as numbers. The current rules cause the following code:

 Utilities awk

```
7262 if (0 == "000")
7263 print "strange, but true"
7264 else
7265 print "not true"
```

to do a numeric comparison, causing the **if** to succeed. It should be intuitively obvious that this is incorrect behavior, and indeed, no historical implementation of *awk* actually behaves this way.

To fix this problem, the definition of *numeric string* was enhanced to include only those values obtained from specific circumstances (mostly external sources) where it is not possible to determine unambiguously whether the value is intended to be a string or a numeric.

Variables that are assigned to a numeric string shall also be treated as a numeric string. (For example, the notion of a numeric string can be propagated across assignments.) In comparisons, all variables having the uninitialized value are to be treated as a numeric operand evaluating to the numeric value zero.

Uninitialized variables include all types of variables including scalars, array elements, and fields. The definition of an uninitialized value in **Variables and Special Variables** (on page 160) is necessary to describe the value placed on uninitialized variables and on fields that are valid (for example, < \$NF) but have no characters in them and to describe how these variables are to be used in comparisons. A valid field, such as \$1, that has no characters in it can be obtained from an input line of "\t\t" when $FS='\t'$. Historically, the comparison (\$1<10) was done numerically after evaluating \$1 to the value zero.

The phrase "... also shall have the numeric value of the numeric string" was removed from several sections of the ISO POSIX-2:1993 standard because is specifies an unnecessary implementation detail. It is not necessary for IEEE Std 1003.1-2001 to specify that these objects be assigned two different values. It is only necessary to specify that these objects may evaluate to two different values depending on context.

The description of numeric string processing is based on the behavior of the *atof()* function in the ISO C standard. While it is not a requirement for an implementation to use this function, many historical implementations of *awk* do. In the ISO C standard, floating-point constants use a period as a decimal point character for the language itself, independent of the current locale, but the *atof()* function and the associated *strtod()* function use the decimal point character of the current locale when converting strings to numeric values. Similarly in *awk*, floating-point constants in an *awk* script use a period independent of the locale, but input strings use the decimal point character of the locale.

FUTURE DIRECTIONS

None.

7297 SEE ALSO

Section 1.10 (on page 19), grep, lex, sed, the System Interfaces volume of IEEE Std 1003.1-2001, atof(), exec, popen(), setlocale(), strtod()

7300 CHANGE HISTORY

First released in Issue 2.

Issue 5

The FUTURE DIRECTIONS section is added.

Issue 6

The *awk* utility is aligned with the IEEE P1003.2b draft standard.

The normative text is reworded to avoid use of the term "must" for application requirements.

awk Utilities

7307 7308 7309 IEEE PASC Interpretation 1003.2 #211 is applied, adding the sentence "An occurrence of two consecutive backslashes shall be interpreted as just a single literal backslash character." into the description of the **sub** string function.

Utilities basename

7310 **NAME**

7313

7315

7316

7317

7318 7319

7320

7321

7322 7323

7324

7325

7326

7327

7328

7329 7330

7331

73327333

7311 basename — return non-directory portion of a pathname

7312 SYNOPSIS

basename string [suffix]

7314 **DESCRIPTION**

The *string* operand shall be treated as a pathname, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.266, Pathname. The string *string* shall be converted to the filename corresponding to the last pathname component in *string* and then the suffix string *suffix*, if present, shall be removed. This shall be done by performing actions equivalent to the following steps in order:

- 1. If *string* is a null string, it is unspecified whether the resulting string is '.' or a null string. In either case, skip steps 2 through 6.
- 2. If *string* is "//", it is implementation-defined whether steps 3 to 6 are skipped or processed.
- 3. If *string* consists entirely of slash characters, *string* shall be set to a single slash character. In this case, skip steps 4 to 6.
- 4. If there are any trailing slash characters in *string*, they shall be removed.
- 5. If there are any slash characters remaining in *string*, the prefix of *string* up to and including the last slash character in *string* shall be removed.
- 6. If the *suffix* operand is present, is not identical to the characters remaining in *string*, and is identical to a suffix of the characters remaining in *string*, the suffix *suffix* shall be removed from *string*. Otherwise, *string* is not modified by this step. It shall not be considered an error if *suffix* is not found in *string*.

The resulting string shall be written to standard output.

7334 **OPTIONS**

7335 None.

7336 **OPERANDS**

7337 The following operands shall be supported:

7338 string A string. 7339 suffix A string.

7340 **STDIN**

7345

7341 Not used.

7342 INPUT FILES

7343 None.

7344 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *basename*:

7346 LANG Provide a default value for the internationalization variables that are unset or null.
7347 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
7348 Internationalization Variables for the precedence of internationalization variables
7349 used to determine the values of locale categories.)
7350 LC_ALL If set to a non-empty string value, override the values of all the other

7350 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

basename Utilities

7352 LC_CTYPE 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 7353 arguments). 7354 LC MESSAGES 7355 Determine the locale that should be used to affect the format and contents of 7356 diagnostic messages written to standard error. 7357 Determine the location of message catalogs for the processing of *LC_MESSAGES*. **NLSPATH** 7358 XSI ASYNCHRONOUS EVENTS 7359 Default. 7360 **STDOUT** 7361 The *basename* utility shall write a line to the standard output in the following format: 7362 7363 "%s\n", <resulting string> STDERR 7364 The standard error shall be used only for diagnostic messages. 7365 **OUTPUT FILES** 7366 None. 7367 **EXTENDED DESCRIPTION** 7368 7369 None **EXIT STATUS** 7370 The following exit values shall be returned: 7371 Successful completion. 7372 7373 >0 An error occurred. **CONSEQUENCES OF ERRORS** 7374 Default. 7375 APPLICATION USAGE 7376 The definition of pathname specifies implementation-defined behavior for pathnames starting 7377 with two slash characters. Therefore, applications shall not arbitrarily add slashes to the 7378 beginning of a pathname unless they can ensure that there are more or less than two or are 7379 prepared to deal with the implementation-defined consequences. 7380 **EXAMPLES** 7381 If the string *string* is a valid pathname: 7382 7383 \$(basename "string") produces a filename that could be used to open the file named by string in the directory returned 7384 7385 by: \$(dirname "string") 7386 If the string string is not a valid pathname, the same algorithm is used, but the result need not be 7387 7388 a valid filename. The basename utility is not expected to make any judgements about the validity of string as a pathname; it just follows the specified algorithm to produce a result string. 7389 The following shell script compiles /usr/src/cmd/cat.c and moves the output to a file named cat 7390

/usr/src/cmd/cat.c:

7391

7392

in the current directory when invoked with the argument /usr/src/cmd/cat or with the argument

Utilities basename

```
7393
             c99 $(dirname "$1")/$(basename "$1" .c).c
7394
             mv a.out $(basename "$1" .c)
7395
     RATIONALE
             The behaviors of basename and dirname have been coordinated so that when string is a valid
7396
7397
             pathname:
              $(basename "string")
7398
7399
             would be a valid filename for the file in the directory:
7400
              $(dirname "string")
             This would not work for the early proposal versions of these utilities due to the way it specified
7401
7402
             handling of trailing slashes.
             Since the definition of pathname specifies implementation-defined behavior for pathnames
7403
             starting with two slash characters, this volume of IEEE Std 1003.1-2001 specifies similar
7404
7405
             implementation-defined behavior for the basename and dirname utilities.
     FUTURE DIRECTIONS
7406
             None.
7407
     SEE ALSO
7408
             Section 2.5 (on page 33), dirname
7409
     CHANGE HISTORY
7410
             First released in Issue 2.
7411
     Issue 6
7412
7413
             IEEE PASC Interpretation 1003.2 #164 is applied.
             The normative text is reworded to avoid use of the term "must" for application requirements.
7414
```

batch Utilities

7415 NAME batch — schedule commands to be executed in a batch queue 7416 7417 **SYNOPSIS** 7418 UP batch 7419 DESCRIPTION 7420 The batch utility shall read commands from standard input and schedule them for execution in a 7421 batch queue. It shall be the equivalent of the command: 7422 at -q b -m now 7423 where queue b is a special at queue, specifically for batch jobs. Batch jobs shall be submitted to 7424 7425 the batch queue with no time constraints and shall be run by the system using algorithms, based on unspecified factors, that may vary with each invocation of *batch*. 7426 Users shall be permitted to use *batch* if their name appears in the file /usr/lib/cron/at.allow. If 7427 XSI 7428 that file does not exist, the file /usr/lib/cron/at.deny shall be checked to determine whether the user shall be denied access to batch. If neither file exists, only a process with the appropriate 7429 7430 privileges shall be allowed to submit a job. If only at.deny exists and is empty, global usage shall be permitted. The at.allow and at.deny files shall consist of one user name per line. 7431 **OPTIONS** 7432 7433 None. **OPERANDS** 7434 None. 7435 **STDIN** 7436 The standard input shall be a text file consisting of commands acceptable to the shell command 7437 language described in Chapter 2 (on page 29). 7438 **INPUT FILES** 7439 The text files /usr/lib/cron/at.allow and /usr/lib/cron/at.deny shall contain zero or more user 7440 names, one per line, of users who are, respectively, authorized or denied access to the at and 7441 batch utilities. ENVIRONMENT VARIABLES 7443 The following environment variables shall affect the execution of *batch*: 7444 LANG Provide a default value for the internationalization variables that are unset or null. 7445 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 7446 Internationalization Variables for the precedence of internationalization variables 7447 used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other 7449 internationalization variables. 7450 LC_CTYPE 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 7452

7458 LC TIMI

LC_MESSAGES

LC_TIME Determine the format and contents for date and time strings written by *batch*.

arguments and input files).

standard output.

7453

7454

7455

7456

7457

Determine the locale that should be used to affect the format and contents of

diagnostic messages written to standard error and informative messages written to

Utilities batch

7459	XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .
7460 7461 7462 7463 7464		SHELL	Determine the name of a command interpreter to be used to invoke the at-job. If the variable is unset or null, <i>sh</i> shall be used. If it is set to a value other than a name for <i>sh</i> , the implementation shall do one of the following: use that shell; use <i>sh</i> ; use the login shell from the user database; any of the preceding accompanied by a warning diagnostic about which was chosen.
7465 7466 7467 7468 7469		TZ	Determine the timezone. The job shall be submitted for execution at the time specified by <i>timespec</i> or -t <i>time</i> relative to the timezone specified by the <i>TZ</i> variable. If <i>timespec</i> specifies a timezone, it overrides <i>TZ</i> . If <i>timespec</i> does not specify a timezone and <i>TZ</i> is unset or null, an unspecified default timezone shall be used.
7470 7471			
7472			
7473 7474			
7475 7476			
7477	"job %s at %s\n",		
7478	where date shall be equivalent in format to the output of:		
7479	date +"%a %b %e %T %Y"		
7480 7481	The date and time written shall be adjusted so that they appear in the timezone of the user (as determined by the TZ variable).		
7482 7483	Neither this, nor warning messages concerning the selection of the command interpreter, are considered a diagnostic that changes the exit status.		
7484	Diagnostic messages, if any, shall be written to standard error.		
7485 7486	OUTPU	T FILES None.	
7487 7488	EXTENI	DED DESCR None.	IPTION
7489 7490	EXIT ST		ng exit values shall be returned:
7491		0 Success	ful completion.
7492		>0 An erro	r occurred.
7493	CONSEQUENCES OF ERRORS The job shall not be scheduled		

The job shall not be scheduled.

7494

batch Utilities

7495 APPLICATION USAGE

It may be useful to redirect standard output within the specified commands.

7497 EXAMPLES

7496

7498

7502 7503

7504

7505 7506

7508

7509

7510

7511

7512

7513 7514

7515

7521

7524

1. This sequence can be used at a terminal:

2. This sequence, which demonstrates redirecting standard error to a pipe, is useful in a command procedure (the sequence of output redirection specifications is significant):

```
batch <<
! diff file1 file2 2>&1 >outfile | mailx mygroup
!
```

7507 RATIONALE

Early proposals described *batch* in a manner totally separated from at, even though the historical model treated it almost as a synonym for at – \mathbf{qb} . A number of features were added to list and control batch work separately from those in at. Upon further reflection, it was decided that the benefit of this did not merit the change to the historical interface.

The **-m** option was included on the equivalent *at* command because it is historical practice to mail results to the submitter, even if all job-produced output is redirected. As explained in the RATIONALE for *at*, the **now** keyword submits the job for immediate execution (after scheduling delays), despite some historical systems where *at* **now** would have been considered an error.

7516 FUTURE DIRECTIONS

7517 None.

7518 SEE ALSO

7519 at

7520 CHANGE HISTORY

First released in Issue 2.

7522 Issue 6

7523 This utility is marked as part of the User Portability Utilities option.

The NAME is changed to align with the IEEE P1003.2b draft standard.

7525 The normative text is reworded to avoid use of the term "must" for application requirements.

7526 NAME 7527 bc — arbitrary-precision arithmetic language 7528 SYNOPSIS 7529 bc [-1] [file ...] 7530 DESCRIPTION 7531 The bc utility shall implement an arbitrary pr 7532 given, then read from the standard input. If

The *bc* utility shall implement an arbitrary precision calculator. It shall take input from any files given, then read from the standard input. If the standard input and standard output to *bc* are attached to a terminal, the invocation of *bc* shall be considered to be *interactive*, causing behavioral constraints described in the following sections.

7535 OPTIONS

7533

7534

7536

7537

7538

7542

7546

7548

7549

7551

7561

7562

7563

The *bc* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following option shall be supported:

7539 —I (The letter ell.) Define the math functions and initialize *scale* to 20, instead of the default zero; see the EXTENDED DESCRIPTION section.

7541 **OPERANDS**

The following operand shall be supported:

7543 *file* A pathname of a text file containing *bc* program statements. After all *files* have been read, *bc* shall read the standard input.

7545 **STDIN**

See the INPUT FILES section.

7547 INPUT FILES

Input files shall be text files containing a sequence of comments, statements, and function definitions that shall be executed as they are read.

7550 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *bc*:

7552 LANG Provide a default value for the internationalization variables that are unset or null.
7553 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
7554 Internationalization Variables for the precedence of internationalization variables
7555 used to determine the values of locale categories.)

7556 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

7558 *LC_CTYPE* 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).

LC_MESSAGES

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

7564 XSI NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

7565 ASYNCHRONOUS EVENTS

7566 Default.

bc Utilities

7567 STDOUT

7568

7569 7570

7571

7573

7577

7578

7579

7580

7581

7582

The output of the *bc* utility shall be controlled by the program read, and consist of zero or more lines containing the value of all executed expressions without assignments. The radix and precision of the output shall be controlled by the values of the **obase** and **scale** variables; see the EXTENDED DESCRIPTION section.

7572 STDERR

The standard error shall be used only for diagnostic messages.

7574 **OUTPUT FILES**

7575 None.

EXTENDED DESCRIPTION

Grammar

The grammar in this section and the lexical conventions in the following section shall together describe the syntax for *bc* programs. The general conventions for this style of grammar are described in Section 1.10 (on page 19). A valid program can be represented as the non-terminal symbol **program** in the grammar. This formal syntax shall take precedence over the text syntax description.

```
%token
                       EOF NEWLINE STRING LETTER NUMBER
7583
                       MUL OP
7584
            %token
                       1*1, 1/1, 181
            /*
                                                                      */
7585
                       ASSIGN OP
            %token
7586
            /*
                       '=', '+=', '-=', '*=', '/=', '%=', '^='
7587
            %token
                       REL OP
7588
                        '==', '<=', '>=', '!=', '<', '>'
            /*
                                                                      * /
7589
            %token
                       INCR DECR
7590
                       '++', '--'
            /*
7591
            %token
                       Define
                                                       Length
7592
                                   Break
                                             Quit
7593
                        'define',
                                  'break', 'quit', 'length'
            %token
                       Return
                                   For
                                           Ιf
                                                  While
7594
            /*
                        'return', 'for', 'if', 'while', 'sqrt'
7595
                                  Ibase
                                            Obase
            %token
                       Scale
                                                       Auto
7596
            /*
                        'scale', 'ibase', 'obase', 'auto'
                                                                      * /
7597
7598
            %start
                       program
            응응
7599
7600
            program
                                     : EOF
                                       input item program
7601
7602
7603
            input item
                                    : semicolon list NEWLINE
7604
                                       function
7605
                                    : /* empty */
7606
            semicolon list
                                      statement
7607
                                       semicolon_list ';' statement
7608
                                     | semicolon list ';'
7609
```

```
7610
            statement list
                                   : /* empty */
7611
7612
                                     statement
                                     statement list NEWLINE
7613
7614
                                     statement list NEWLINE statement
                                     statement list ';'
7615
                                     statement_list ';' statement
7616
7617
7618
            statement
                                   : expression
7619
                                     STRING
                                     Break
7620
7621
                                     Quit
7622
                                     Return
                                     Return '(' return expression ')'
7623
7624
                                     For '(' expression ';'
                                          relational expression ';'
7625
                                          expression ')' statement
7626
                                     If '(' relational_expression ')' statement
7627
                                     While '(' relational expression ')' statement
7628
                                     '{' statement_list '}'
7629
7630
7631
            function
                                   : Define LETTER '(' opt parameter list ')'
                                          '{' NEWLINE opt_auto_define list
7632
                                          statement list '}'
7633
7634
            opt parameter list
                                   : /* empty */
7635
                                   | parameter list
7636
7637
            parameter list
                                   : LETTER
7638
7639
                                     define list ',' LETTER
7640
7641
            opt_auto_define_list : /* empty */
7642
                                     Auto define list NEWLINE
                                     Auto define list ';'
7643
7644
            define list
                                   : LETTER
7645
                                     LETTER '[' ']'
7646
                                     define list ',' LETTER
7647
                                     define list ',' LETTER '[' ']'
7648
7649
                                   : /* empty */
7650
            opt argument list
7651
                                   argument_list
7652
            argument list
                                   : expression
7653
                                   | LETTER '[' ']' ',' argument list
7654
7655
```

bc Utilities

```
7656
           relational expression : expression
7657
                                   expression REL_OP expression
7658
                                    : /* empty */
7659
           return expression
7660
                                     expression
7661
            expression
                                   : named expression
7662
                                     NUMBER
7663
                                      '(' expression ')'
7664
7665
                                      LETTER '(' opt argument list ')'
                                      '-' expression
7666
                                      expression '+' expression
7667
                                      expression '-' expression
7668
                                      expression MUL OP expression
7669
                                      expression '^' expression
7670
                                      INCR DECR named expression
7671
                                      named expression INCR DECR
7672
                                      named expression ASSIGN OP expression
7673
                                     Length '(' expression ')'
7674
                                      Sgrt '(' expression ')'
7675
                                      Scale '(' expression ')'
7676
7677
                                    : LETTER
           named expression
7678
                                     LETTER '[' expression ']'
7679
                                      Scale
7680
7681
                                      Ibase
7682
                                      Obase
7683
```

Lexical Conventions in bc

The lexical conventions for *bc* programs, with respect to the preceding grammar, shall be as follows:

- 1. Except as noted, *bc* shall recognize the longest possible token or delimiter beginning at a given point.
- 2. A comment shall consist of any characters beginning with the two adjacent characters "/*" and terminated by the next occurrence of the two adjacent characters "*/". Comments shall have no effect except to delimit lexical tokens.
- 3. The <newline> shall be recognized as the token **NEWLINE**.
- 4. The token **STRING** shall represent a string constant; it shall consist of any characters beginning with the double-quote character ('"') and terminated by another occurrence of the double-quote character. The value of the string is the sequence of all characters between, but not including, the two double-quote characters. All characters shall be taken literally from the input, and there is no way to specify a string containing a double-quote character. The length of the value of each string shall be limited to {BC_STRING_MAX} bytes.
- 5. A <blank> shall have no effect except as an ordinary character if it appears within a **STRING** token, or to delimit a lexical token other than **STRING**.

7684

7685

7686 7687

7688 7689

7690

7691

7692

7693

7694

7695

7696

7697

7698 7699

7700

7701

6. The combination of a backslash character immediately followed by a <newline> shall have no effect other than to delimit lexical tokens with the following exceptions:

- It shall be interpreted as the character sequence "\<newline>" in **STRING** tokens.
- It shall be ignored as part of a multi-line NUMBER token.
- 7. The token **NUMBER** shall represent a numeric constant. It shall be recognized by the following grammar:

- 8. The value of a **NUMBER** token shall be interpreted as a numeral in the base specified by the value of the internal register **ibase** (described below). Each of the **digit** characters shall have the value from 0 to 15 in the order listed here, and the period character shall represent the radix point. The behavior is undefined if digits greater than or equal to the value of **ibase** appear in the token. However, note the exception for single-digit values being assigned to **ibase** and **obase** themselves, in **Operations in bc** (on page 198).
- 9. The following keywords shall be recognized as tokens:

```
auto ibase length return while
break if obase scale
define for quit sqrt
```

10. Any of the following characters occurring anywhere except within a keyword shall be recognized as the token **LETTER**:

```
abcdefghijklmnopqrstuvwxyz
```

11. The following single-character and two-character sequences shall be recognized as the token **ASSIGN_OP**:

```
= += -= *= /= %= ^=
```

- 12. If an '=' character, as the beginning of a token, is followed by a '-' character with no intervening delimiter, the behavior is undefined.
 - 13. The following single-characters shall be recognized as the token MUL_OP:

```
* / %
```

14. The following single-character and two-character sequences shall be recognized as the token **REL_OP**:

```
== <= >= != < >
```

15. The following two-character sequences shall be recognized as the token **INCR_DECR**:

bc Utilities

7743 ++ --

16. The following single characters shall be recognized as tokens whose names are the character:

```
<newline> ( ) , + - ; [ ] ^ { }
```

17. The token **EOF** is returned when the end of input is reached.

Operations in bc

There are three kinds of identifiers: ordinary identifiers, array identifiers, and function identifiers. All three types consist of single lowercase letters. Array identifiers shall be followed by square brackets ("[]"). An array subscript is required except in an argument or auto list. Arrays are singly dimensioned and can contain up to {BC_DIM_MAX} elements. Indexing shall begin at zero so an array is indexed from 0 to {BC_DIM_MAX}–1. Subscripts shall be truncated to integers. The application shall ensure that function identifiers are followed by parentheses, possibly enclosing arguments. The three types of identifiers do not conflict.

The following table summarizes the rules for precedence and associativity of all operators. Operators on the same line shall have the same precedence; rows are in order of decreasing precedence.

Table 4-3 Operators in bc

Operator	Associativity
++,	N/A
unary -	N/A
^	Right to left
*, /, %	Left to right
+, binary -	Left to right
=, +=, -=, *=, /=, %=, ^=	Right to left
==, <=, >=, !=, <, >	None

Each expression or named expression has a *scale*, which is the number of decimal digits that shall be maintained as the fractional portion of the expression.

Named expressions are places where values are stored. Named expressions shall be valid on the left side of an assignment. The value of a named expression shall be the value stored in the place named. Simple identifiers and array elements are named expressions; they have an initial value of zero and an initial scale of zero.

The internal registers **scale**, **ibase**, and **obase** are all named expressions. The scale of an expression consisting of the name of one of these registers shall be zero; values assigned to any of these registers are truncated to integers. The **scale** register shall contain a global value used in computing the scale of expressions (as described below). The value of the register **scale** is limited to $0 \le \text{scale} \le \{\text{BC_SCALE_MAX}\}$ and shall have a default value of zero. The **ibase** and **obase** registers are the input and output number radix, respectively. The value of **ibase** shall be limited to:

```
2 \le ibase \le 16
```

The value of **obase** shall be limited to:

```
2 \le obase \le \{BC BASE MAX\}
```

When either **ibase** or **obase** is assigned a single **digit** value from the list in **Lexical Conventions** in **bc** (on page 196), the value shall be assumed in hexadecimal. (For example, **ibase**=A sets to

base ten, regardless of the current **ibase** value.) Otherwise, the behavior is undefined when digits greater than or equal to the value of **ibase** appear in the input. Both **ibase** and **obase** shall have initial values of 10.

Internal computations shall be conducted as if in decimal, regardless of the input and output bases, to the specified number of decimal digits. When an exact result is not achieved (for example, scale=0; 3.2/1), the result shall be truncated.

For all values of **obase** specified by this volume of IEEE Std 1003.1-2001, *bc* shall output numeric values by performing each of the following steps in order:

- 1. If the value is less than zero, a hyphen ('-') character shall be output.
- 2. One of the following is output, depending on the numerical value:
 - If the absolute value of the numerical value is greater than or equal to one, the integer portion of the value shall be output as a series of digits appropriate to **obase** (as described below), most significant digit first. The most significant non-zero digit shall be output next, followed by each successively less significant digit.
 - If the absolute value of the numerical value is less than one but greater than zero and the scale of the numerical value is greater than zero, it is unspecified whether the character 0 is output.
 - If the numerical value is zero, the character 0 shall be output.
- 3. If the scale of the value is greater than zero and the numeric value is not zero, a period character shall be output, followed by a series of digits appropriate to **obase** (as described below) representing the most significant portion of the fractional part of the value. If *s* represents the scale of the value being output, the number of digits output shall be *s* if **obase** is 10, less than or equal to *s* if **obase** is greater than 10, or greater than or equal to *s* if **obase** is less than 10. For **obase** values other than 10, this should be the number of digits needed to represent a precision of 10^s.

For **obase** values from 2 to 16, valid digits are the first **obase** of the single characters:

0 1 2 3 4 5 6 7 8 9 A B C D E F

7813 which represent the values zero to 15, inclusive, respectively.

For bases greater than 16, each digit shall be written as a separate multi-digit decimal number. Each digit except the most significant fractional digit shall be preceded by a single <space>. For bases from 17 to 100, *bc* shall write two-digit decimal numbers; for bases from 101 to 1000, three-digit decimal strings, and so on. For example, the decimal number 1024 in base 25 would be written as:

 $\Delta 01\Delta 15\Delta 24$

7820 and in base 125, as:

 Δ 008 Δ 024

Very large numbers shall be split across lines with 70 characters per line in the POSIX locale; other locales may split at different character boundaries. Lines that are continued shall end with a backslash ($' \setminus '$).

A function call shall consist of a function name followed by parentheses containing a commaseparated list of expressions, which are the function arguments. A whole array passed as an argument shall be specified by the array name followed by empty square brackets. All function arguments shall be passed by value. As a result, changes made to the formal parameters shall have no effect on the actual arguments. If the function terminates by executing a **return** **bc** Utilities

statement, the value of the function shall be the value of the expression in the parentheses of the return statement or shall be zero if no expression is provided or if there is no return statement.

The result of sart(expression) shall be the square root of the expression. The result shall be

The result of **sqrt**(*expression*) shall be the square root of the expression. The result shall be truncated in the least significant decimal place. The scale of the result shall be the scale of the expression or the value of **scale**, whichever is larger.

The result of **length**(*expression*) shall be the total number of significant decimal digits in the expression. The scale of the result shall be zero.

The result of **scale**(*expression*) shall be the scale of the expression. The scale of the result shall be zero.

A numeric constant shall be an expression. The scale shall be the number of digits that follow the radix point in the input representing the constant, or zero if no radix point appears.

The sequence (*expression*) shall be an expression with the same value and scale as *expression*. The parentheses can be used to alter the normal precedence.

The semantics of the unary and binary operators are as follows:

-expression

7833 7834

7836

7837 7838

7839

7840

7841

7842

7843 7844

7845

7846

7847

7848

7849

7850

7851 7852

7853

7854 7855

7857

7858 7859

7860

7861

7862

7863

7864

7865

7866

7867

7868

7869

7870

7871

The result shall be the negative of the *expression*. The scale of the result shall be the scale of *expression*.

The unary increment and decrement operators shall not modify the scale of the named expression upon which they operate. The scale of the result shall be the scale of that named expression.

++named-expression

The named expression shall be incremented by one. The result shall be the value of the named expression after incrementing.

-- named-expression

The named expression shall be decremented by one. The result shall be the value of the named expression after decrementing.

named-expression++

The named expression shall be incremented by one. The result shall be the value of the named expression before incrementing.

named-expression--

The named expression shall be decremented by one. The result shall be the value of the named expression before decrementing.

The exponentiation operator, circumflex $(' ^)$, shall bind right to left.

expression expression

The result shall be the first *expression* raised to the power of the second *expression*. If the second expression is not an integer, the behavior is undefined. If *a* is the scale of the left expression and *b* is the absolute value of the right expression, the scale of the result shall be:

```
if b >= 0 \min(a * b, \max(scale, a)) if b < 0 scale
```

The multiplicative operators ('*', '/', '%') shall bind left to right.

expression*expression

The result shall be the product of the two expressions. If *a* and *b* are the scales of the two expressions, then the scale of the result shall be:

```
7872
                   min(a+b, max(scale, a, b))
7873
              expression/expression
7874
                   The result shall be the quotient of the two expressions. The scale of the result shall be the
                   value of scale.
7875
7876
              expression%expression
                   For expressions a and b, a\%b shall be evaluated equivalent to the steps:
7877
                    1. Compute a/b to current scale.
7878
                    2. Use the result to compute:
7879
                        a - (a / b) * b
7880
                        to scale:
7881
                        max(scale + scale(b), scale(a))
7882
                   The scale of the result shall be:
7883
                   max(scale + scale(b), scale(a))
7884
                   When scale is zero, the '%' operator is the mathematical remainder operator.
7885
              The additive operators (' + ', ' - ') shall bind left to right.
7886
7887
              expression+expression
                   The result shall be the sum of the two expressions. The scale of the result shall be the
7888
                   maximum of the scales of the expressions.
7889
7890
              expression-expression
                   The result shall be the difference of the two expressions. The scale of the result shall be the
7891
                   maximum of the scales of the expressions.
7892
              The assignment operators ('=', "+=", "-=", "*=", "/=", "%=", "^=") shall bind right to left.
7893
7894
              named-expression=expression
                   This expression shall result in assigning the value of the expression on the right to the
7895
                   named expression on the left. The scale of both the named expression and the result shall be
                   the scale of expression.
7897
              The compound assignment forms:
7898
7899
              named-expression <operator>= expression
              shall be equivalent to:
7900
7901
              named-expression=named-expression <operator> expression
              except that the named-expression shall be evaluated only once.
7902
              Unlike all other operators, the relational operators (' < ', ' > ', " <= ", " >= ", " == ", " != ") shall be
7903
              only valid as the object of an if, while, or inside a for statement.
7904
              expression1<expression2
7905
7906
                   The relation shall be true if the value of expression1 is strictly less than the value of
7907
                   expression2.
              expression1>expression2
7908
                   The relation shall be true if the value of expression1 is strictly greater than the value of
7909
                   expression2.
7910
```

bc Utilities

```
7911 expression1<=expression2
```

 The relation shall be true if the value of *expression1* is less than or equal to the value of *expression2*.

expression1>=expression2

The relation shall be true if the value of *expression1* is greater than or equal to the value of *expression2*.

expression1 = expression2

The relation shall be true if the values of *expression1* and *expression2* are equal.

expression1!=expression2

The relation shall be true if the values of *expression1* and *expression2* are unequal.

There are only two storage classes in *bc*: global and automatic (local). Only identifiers that are local to a function need be declared with the **auto** command. The arguments to a function shall be local to the function. All other identifiers are assumed to be global and available to all functions. All identifiers, global and local, have initial values of zero. Identifiers declared as auto shall be allocated on entry to the function and released on returning from the function. They therefore do not retain values between function calls. Auto arrays shall be specified by the array name followed by empty square brackets. On entry to a function, the old values of the names that appear as parameters and as automatic variables shall be pushed onto a stack. Until the function returns, reference to these names shall refer only to the new values.

References to any of these names from other functions that are called from this function also refer to the new value until one of those functions uses the same name for a local variable.

When a statement is an expression, unless the main operator is an assignment, execution of the statement shall write the value of the expression followed by a <newline>.

When a statement is a string, execution of the statement shall write the value of the string.

Statements separated by semicolons or <newline>s shall be executed sequentially. In an interactive invocation of bc, each time a <newline> is read that satisfies the grammatical production:

```
input item : semicolon list NEWLINE
```

the sequential list of statements making up the **semicolon_list** shall be executed immediately and any output produced by that execution shall be written without any delay due to buffering.

In an **if** statement (**if**(relation) statement), the statement shall be executed if the relation is true.

The **while** statement (**while**(*relation*) *statement*) implements a loop in which the *relation* is tested; each time the *relation* is true, the *statement* shall be executed and the *relation* retested. When the *relation* is false, execution shall resume after *statement*.

A for statement(for(expression; relation; expression) statement) shall be the same as:

The application shall ensure that all three expressions are present.

The **break** statement shall cause termination of a **for** or **while** statement.

The **auto** statement (**auto** *identifier* [,*identifier*] ...) shall cause the values of the identifiers to be pushed down. The identifiers can be ordinary identifiers or array identifiers. Array identifiers

shall be specified by following the array name by empty square brackets. The application shall ensure that the **auto** statement is the first statement in a function definition.

A **define** statement:

7955

7956 7957

7962

7963 7964

7965

7966

7967

7968

7969

7970

7971

7972

7973

7974

7975

7976

7977

7988

7989

7990

7991

7992

7993 7994

7996

defines a function named **LETTER**. If a function named **LETTER** was previously defined, the **define** statement shall replace the previous definition. The expression:

```
LETTER ( opt argument list )
```

shall invoke the function named **LETTER**. The behavior is undefined if the number of arguments in the invocation does not match the number of parameters in the definition. Functions shall be defined before they are invoked. A function shall be considered to be defined within its own body, so recursive calls are valid. The values of numeric constants within a function shall be interpreted in the base specified by the value of the **ibase** register when the function is invoked.

The **return** statements (**return** and **return**(*expression*)) shall cause termination of a function, popping of its auto variables, and specification of the result of the function. The first form shall be equivalent to **return**(0). The value and scale of the result returned by the function shall be the value and scale of the expression returned.

The **quit** statement (**quit**) shall stop execution of a *bc* program at the point where the statement occurs in the input, even if it occurs in a function definition, or in an **if**, **for**, or **while** statement.

The following functions shall be defined when the –l option is specified:

```
7978
              s(expression)
                   Sine of argument in radians.
7979
7980
              c(expression)
                   Cosine of argument in radians.
7981
              a(expression)
7982
                   Arctangent of argument.
7983
              l(expression)
7984
                   Natural logarithm of argument.
7985
7986
              e(expression)
                   Exponential function of argument.
7987
```

Exponential function of argument.

j(expression, expression)

Bessel function of integer order.

The scale of the result returned by these functions shall be the value of the **scale** register at the time the function is invoked. The value of the **scale** register after these functions have completed their execution shall be the same value it had upon invocation. The behavior is undefined if any of these functions is invoked with an argument outside the domain of the mathematical function.

7995 EXIT STATUS

The following exit values shall be returned:

7997 0 All input files were processed successfully.

bc Utilities

unspecified An error occurred.

CONSEQUENCES OF ERRORS

If any *file* operand is specified and the named file cannot be accessed, *bc* shall write a diagnostic message to standard error and terminate without any further action.

In an interactive invocation of *bc*, the utility should print an error message and recover following any error in the input. In a non-interactive invocation of *bc*, invalid input causes undefined behavior.

APPLICATION USAGE

Automatic variables in bc do not work in exactly the same way as in either C or PL/1.

For historical reasons, the exit status from *bc* cannot be relied upon to indicate that an error has occurred. Returning zero after an error is possible. Therefore, *bc* should be used primarily by interactive users (who can react to error messages) or by application programs that can somehow validate the answers returned as not including error messages.

The *bc* utility always uses the period (' . ') character to represent a radix point, regardless of any decimal-point character specified as part of the current locale. In languages like C or *awk*, the period character is used in program source, so it can be portable and unambiguous, while the locale-specific character is used in input and output. Because there is no distinction between source and input in *bc*, this arrangement would not be possible. Using the locale-specific character in *bc*'s input would introduce ambiguities into the language; consider the following example in a locale with a comma as the decimal-point character:

Because of such ambiguities, the period character is used in input. Having input follow different conventions from output would be confusing in either pipeline usage or interactive usage, so the period is also used in output.

8026 EXAMPLES

In the shell, the following assigns an approximation of the first ten digits of ' π ' to the variable \mathbf{x} '.

```
x=\$(printf "\$s\n" 'scale = 10; 104348/33215' | bc)
```

The following *bc* program prints the same approximation of ' π ', with a label, to standard output:

```
scale = 10
"pi equals "
104348 / 33215
```

The following defines a function to compute an approximate value of the exponential function (note that such a function is predefined if the –l option is specified):

```
8037 scale = 20

8038 define e(x) {

8039 auto a, b, c, i, s

8040 a = 1

8041 b = 1

8042 s = 1
```

```
8043
                  for (i = 1; 1 == 1; i++)
8044
                        a = a*x
                       b = b*i
8045
                        c = a/b
8046
8047
                        if (c == 0) {
                              return(s)
2012
8049
8050
                        s = s + c
                  }
8051
             }
8052
```

The following prints approximate values of the exponential function of the first ten integers:

```
for (i = 1; i <= 10; ++i) {
    e(i)
}</pre>
```

RATIONALE

The bc utility is implemented historically as a front-end processor for dc; dc was not selected to be part of this volume of IEEE Std 1003.1-2001 because bc was thought to have a more intuitive programmatic interface. Current implementations that implement bc using dc are expected to be compliant.

The exit status for error conditions has been left unspecified for several reasons:

- The *bc* utility is used in both interactive and non-interactive situations. Different exit codes may be appropriate for the two uses.
- It is unclear when a non-zero exit should be given; divide-by-zero, undefined functions, and syntax errors are all possibilities.
- · It is not clear what utility the exit status has.
- In the 4.3 BSD, System V, and Ninth Edition implementations, *bc* works in conjunction with *dc*. The *dc* utility is the parent, *bc* is the child. This was done to cleanly terminate *bc* if *dc* aborted.

The decision to have bc exit upon encountering an inaccessible input file is based on the belief that bc file1 file2 is used most often when at least file1 contains data/function declarations/initializations. Having bc continue with prerequisite files missing is probably not useful. There is no implication in the CONSEQUENCES OF ERRORS section that bc must check all its files for accessibility before opening any of them.

There was considerable debate on the appropriateness of the language accepted by *bc*. Several reviewers preferred to see either a pure subset of the C language or some changes to make the language more compatible with C. While the *bc* language has some obvious similarities to C, it has never claimed to be compatible with any version of C. An interpreter for a subset of C might be a very worthwhile utility, and it could potentially make *bc* obsolete. However, no such utility is known in historical practice, and it was not within the scope of this volume of IEEE Std 1003.1-2001 to define such a language and utility. If and when they are defined, it may be appropriate to include them in a future version of IEEE Std 1003.1. This left the following alternatives:

1. Exclude any calculator language from this volume of IEEE Std 1003.1-2001.

The consensus of the standard developers was that a simple programmatic calculator language is very useful for both applications and interactive users. The only arguments for excluding any calculator were that it would become obsolete if and when a C-compatible

bc Utilities

 one emerged, or that the absence would encourage the development of such a C-compatible one. These arguments did not sufficiently address the needs of current application writers.

2. Standardize the historical *dc*, possibly with minor modifications.

 The consensus of the standard developers was that dc is a fundamentally less usable language and that that would be far too severe a penalty for avoiding the issue of being similar to but incompatible with C.

 3. Standardize the historical *bc*, possibly with minor modifications.

This was the approach taken. Most of the proponents of changing the language would not have been satisfied until most or all of the incompatibilities with C were resolved. Since most of the changes considered most desirable would break historical applications and require significant modification to historical implementations, almost no modifications were made. The one significant modification that was made was the replacement of the historical bc assignment operators "=+", and so on, with the more modern "+=", and so on. The older versions are considered to be fundamentally flawed because of the lexical ambiguity in uses like a=-1.

In order to permit implementations to deal with backwards-compatibility as they see fit, the behavior of this one ambiguous construct was made undefined. (At least three implementations have been known to support this change already, so the degree of change involved should not be great.)

The '%' operator is the mathematical remainder operator when **scale** is zero. The behavior of this operator for other values of **scale** is from historical implementations of bc, and has been maintained for the sake of historical applications despite its non-intuitive nature.

Historical implementations permit setting **ibase** and **obase** to a broader range of values. This includes values less than 2, which were not seen as sufficiently useful to standardize. These implementations do not interpret input properly for values of **ibase** that are greater than 16. This is because numeric constants are recognized syntactically, rather than lexically, as described in this volume of IEEE Std 1003.1-2001. They are built from lexical tokens of single hexadecimal digits and periods. Since <blank>s between tokens are not visible at the syntactic level, it is not possible to recognize the multi-digit "digits" used in the higher bases properly. The ability to recognize input in these bases was not considered useful enough to require modifying these implementations. Note that the recognition of numeric constants at the syntactic level is not a problem with conformance to this volume of IEEE Std 1003.1-2001, as it does not impact the behavior of conforming applications (and correct bc programs). Historical implementations also accept input with all of the digits '0'-'9' and 'A'-'F' regardless of the value of **ibase**; since digits with value greater than or equal to **ibase** are not really appropriate, the behavior when they appear is undefined, except for the common case of:

```
8126
8127
8128
8129
8130
```

```
ibase=8;
    /* Process in octal base. */
...
ibase=A
    /* Restore decimal base. */
```

In some historical implementations, if the expression to be written is an uninitialized array element, a leading <space> and/or up to four leading 0 characters may be output before the character zero. This behavior is considered a bug; it is unlikely that any currently conforming application relies on:

8135	echo 'b[3]' bc
8136	returning 00000 rather than 0.
8137 8138 8139 8140 8141	Exact calculation of the number of fractional digits to output for a given value in a base other than 10 can be computationally expensive. Historical implementations use a faster approximation, and this is permitted. Note that the requirements apply only to values of obase that this volume of IEEE Std 1003.1-2001 requires implementations to support (in particular, not to 1, 0, or negative bases, if an implementation supports them as an extension).
8142 8143 8144	Historical implementations of bc did not allow array parameters to be passed as the last parameter to a function. New implementations are encouraged to remove this restriction even though it is not required by the grammar.
	FUTURE DIRECTIONS
8146	None.
8147 S 8148	SEE ALSO Section 1.10 (on page 19), awk
8149	CHANGE HISTORY
8150	First released in Issue 4.
8151 I 8152	Issue 5 The FUTURE DIRECTIONS section is added.
8153] 8154 8155	Updated to align with the IEEE P1003.2b draft standard, which included resolution of several interpretations of the ISO POSIX-2: 1993 standard.
8156	The normative text is reworded to avoid use of the term "must" for application requirements.

bg **Utilities**

```
8157
     NAME
              bg — run jobs in the background
8158
8159
     SYNOPSIS
              bg [job id ...]
8160
     UP
8161
     DESCRIPTION
8162
              If job control is enabled (see the description of set - m), the bg utility shall resume suspended jobs
8163
              from the current environment (see Section 2.12 (on page 61)) by running them as background
8164
              jobs. If the job specified by job_id is already a running background job, the bg utility shall have no
8165
8166
              effect and shall exit successfully.
              Using bg to place a job into the background shall cause its process ID to become "known in the
8167
              current shell execution environment", as if it had been started as an asynchronous list; see
8168
8169
              Section 2.9.3.1 (on page 50).
     OPTIONS
8170
              None.
8171
     OPERANDS
8172
8173
              The following operand shall be supported:
                           Specify the job to be resumed as a background job. If no job_id operand is given,
8174
              job id
                           the most recently suspended job shall be used. The format of job_id is described in
8175
                           the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.203, Job Control Job
8176
8177
                           ID.
     STDIN
8178
              Not used.
8179
     INPUT FILES
8180
              None.
8181
     ENVIRONMENT VARIABLES
8182
              The following environment variables shall affect the execution of bg:
8183
              LANG
                           Provide a default value for the internationalization variables that are unset or null.
8184
                           (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
8185
                           Internationalization Variables for the precedence of internationalization variables
8186
                           used to determine the values of locale categories.)
8187
              LC_ALL
8188
                           If set to a non-empty string value, override the values of all the other
8189
                           internationalization variables.
              LC CTYPE
                           Determine the locale for the interpretation of sequences of bytes of text data as
8190
                           characters (for example, single-byte as opposed to multi-byte characters in
8191
8192
                           arguments).
              LC_MESSAGES
8193
                           Determine the locale that should be used to affect the format and contents of
8194
                           diagnostic messages written to standard error.
8195
              NLSPATH
                           Determine the location of message catalogs for the processing of LC_MESSAGES.
8196
     XSI
     ASYNCHRONOUS EVENTS
```

Default.

8197

8198

8199 **STDOUT** The output of *bg* shall consist of a line in the format: 8200 8201 "[%d] %sn", <job-number>, <command> 8202 where the fields are as follows: <job-number> A number that can be used to identify the job to the wait, fg, and kill utilities. Using 8203 these utilities, the job can be identified by prefixing the job number with '%'. 8204 The associated command that was given to the shell. 8205 <command> **STDERR** 8206 The standard error shall be used only for diagnostic messages. 8207 **OUTPUT FILES** 8208 None. 8209 **EXTENDED DESCRIPTION** 8210 8211 None. **EXIT STATUS** 8212 The following exit values shall be returned: 8213 Successful completion. 8214 8215 >0 An error occurred. 8216 **CONSEQUENCES OF ERRORS** If job control is disabled, the bg utility shall exit with an error and no job shall be placed in the 8217 background. 8218 APPLICATION USAGE 8219 A job is generally suspended by typing the SUSP character (<control>-Z on most systems); see 8220 the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface. At 8221 that point, bg can put the job into the background. This is most effective when the job is 8222 8223 expecting no terminal input and its output has been redirected to non-terminal files. A background job can be forced to stop when it has terminal output by issuing the command: 8224 stty tostop 8225 A background job can be stopped with the command: 8226 8227 kill -s stop job ID The bg utility does not work as expected when it is operating in its own utility execution 8228 environment because that environment has no suspended jobs. In the following examples: 8229 8230 ... | xargs bg 8231 (bg) each bg operates in a different environment and does not share its parent shell's understanding 8232 8233 of jobs. For this reason, bg is generally implemented as a shell regular built-in. **EXAMPLES** 8234 None. 8235 RATIONALE 8236 The extensions to the shell specified in this volume of IEEE Std 1003.1-2001 have mostly been 8237 based on features provided by the KornShell. The job control features provided by bg, fg, and jobs 8238 8239 are also based on the KornShell. The standard developers examined the characteristics of the C

8240

shell versions of these utilities and found that differences exist. Despite widespread use of the C

bgUtilities

8241 8242 8243	shell, the KornShell versions were selected for this volume of IEEE Std 1003.1-2001 to maintain a degree of uniformity with the rest of the KornShell features selected (such as the very popular command line editing features).
8244 8245	The <i>bg</i> utility is expected to wrap its output if the output exceeds the number of display columns.
8246 8247	FUTURE DIRECTIONS None.
8248 8249	SEE ALSO Section 2.9.3.1 (on page 50), fg, kill, jobs, wait
8250 8251	CHANGE HISTORY First released in Issue 4.
8252 8253	Issue 6 This utility is marked as part of the User Portability Utilities option.
8254 8255	The JC margin marker on the SYNOPSIS is removed since support for Job Control is mandatory in this issue. This is a FIPS requirement.

```
8256 NAME
```

c99 — compile standard C programs

8258 SYNOPSIS

```
8259 CD c99 [-c] [-D name[=value]]...[-E] [-g] [-I directory] ... [-L directory] 8260 ... [-o outfile] [-Ooptlevel] [-s] [-U name]... operand ...
```

8261 8262

8263

8264 8265

8266

8267

8268

8270 8271

8272

8273 8274

8275

8276

8277 8278

8280 8281

8282

8283

8284

8285 8286

8287

8290

8291

8292

8257

DESCRIPTION

The *c99* utility is an interface to the standard C compilation system; it shall accept source code conforming to the ISO C standard. The system conceptually consists of a compiler and link editor. The files referenced by *operands* shall be compiled and linked to produce an executable file. (It is unspecified whether the linking occurs entirely within the operation of *c99*; some implementations may produce objects that are not fully resolved until the file is executed.)

If the -c option is specified, for all pathname operands of the form *file.c*, the files:

```
$(basename pathname .c).o
```

shall be created as the result of successful compilation. If the -c option is not specified, it is unspecified whether such .o files are created or deleted for the *file* .c operands.

If there are no options that prevent link editing (such as -c or -E), and all operands compile and link without error, the resulting executable file shall be written according to the -o outfile option (if present) or to the file **a.out**.

The executable file shall be created as specified in Section 1.7.1.4 (on page 4), except that the file permission bits shall be set to:

```
S_IRWXO | S_IRWXG | S_IRWXU
```

and the bits specified by the *umask* of the process shall be cleared.

8279 OPTIONS

The *c99* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines, except that:

- The **–l** *library* operands have the format of options, but their position within a list of operands affects the order in which libraries are searched.
- The order of specifying the –I and –L options is significant.
- Conforming applications shall specify each option separately; that is, grouping option letters (for example, -cO) need not be recognized by all implementations.

The following options shall be supported:

- Suppress the link-edit phase of the compilation, and do not remove any object files that are produced.
 - Produce symbolic information in the object or executable files; the nature of this information is unspecified, and may be modified by implementation-defined interactions with other options.
- Produce object or executable files, or both, from which symbolic and other information not required for proper execution using the *exec* family defined in the System Interfaces volume of IEEE Std 1003.1-2001 has been removed (stripped). If both –g and –s options are present, the action taken is unspecified.
- Use the pathname *outfile*, instead of the default **a.out**, for the executable file produced. If the $-\mathbf{o}$ option is present with $-\mathbf{c}$ or $-\mathbf{E}$, the result is unspecified.

c99 Utilities

8299	− D name[=va	alue]
8300	-	Define <i>name</i> as if by a C-language #define directive. If no =value is given, a value of
8301		1 shall be used. The –D option has lower precedence than the –U option. That is, if
8302		name is used in both a –U and a –D option, name shall be undefined regardless of
8303		the order of the options. Additional implementation-defined <i>names</i> may be
8304 8305		provided by the compiler. Implementations shall support at least 2 048 bytes of –D definitions and 256 <i>names</i> .
8306	- E	Copy C-language source files to standard output, expanding all preprocessor
8307	2	directives; no compilation shall be performed. If any operand is not a text file, the
8308		effects are unspecified.
8309	-I directory	Change the algorithm for searching for headers whose names are not absolute
8310		pathnames to look in the directory named by the directory pathname before
8311		looking in the usual places. Thus, headers whose names are enclosed in double-
8312		quotes (" ") shall be searched for first in the directory of the file with the #include
8313		line, then in directories named in –I options, and last in the usual places. For
8314 8315		headers whose names are enclosed in angle brackets ("<>"), the header shall be searched for only in directories named in —I options and then in the usual places.
8316		Directories named in –I options shall be searched in the order specified.
8317		Implementations shall support at least ten instances of this option in a single <i>c99</i>
8318		command invocation.
8319	-L directory	Change the algorithm of searching for the libraries named in the $-\mathbf{l}$ objects to look
8320		in the directory named by the <i>directory</i> pathname before looking in the usual
8321		places. Directories named in –L options shall be searched in the order specified.
8322		Implementations shall support at least ten instances of this option in a single <i>c99</i> command invocation. If a directory specified by a –L option contains files named
8323 8324		libc.a, libm.a, libl.a, or liby.a, the results are unspecified.
8325	−O optlevel	Specify the level of code optimization. If the <i>optlevel</i> option-argument is the digit
8326 8327		'0', all special code optimizations shall be disabled. If it is the digit '1', the nature of the optimization is unspecified. If the –O option is omitted, the nature of
8328		the system's default optimization is unspecified. It is unspecified whether code
8329		generated in the presence of the –O 0 option is the same as that generated when
8330		−O is omitted. Other <i>optlevel</i> values may be supported.
8331	− U name	Remove any initial definition of <i>name</i> .
8332	Multiple inst	tances of the $-\mathbf{D}$, $-\mathbf{I}$, $-\mathbf{U}$, and $-\mathbf{L}$ options can be specified.
	PERANDS	
8334		is either in the form of a pathname or the form –l library. The application shall
8335		at least one operand of the pathname form is specified. The following operands shall
8336	be supported	
8337 8338	file.c	A C-language source file to be compiled and optionally linked. The application shall ensure that the operand is of this form if the –c option is used.
8339	file.a	A library of object files typically produced by the ar utility, and passed directly to
8340		the link editor. Implementations may recognize implementation-defined suffixes
8341		other than .a as denoting object file libraries.
8342	file.o	An object file produced by $c99$ –c and passed directly to the link editor.
8343		Implementations may recognize implementation-defined suffixes other than .o as
8344		denoting object files.

8345		The processi	ng of other files is implementation-defined.
8346		-l library	(The letter ell.) Search the library named:
8347			lib <i>library</i> .a
8348 8349 8350 8351			A library shall be searched when its name is encountered, so the placement of a -1 operand is significant. Several standard libraries can be specified in this manner, as described in the EXTENDED DESCRIPTION section. Implementations may recognize implementation-defined suffixes other than $.a$ as denoting libraries.
8352	STDIN		
8353		Not used.	
8354 8355 8356 8357 8358	INPUT 1	The input file an object file by archiving	the shall be one of the following: a text file containing a C-language source program, in the format produced by $c99$ –c, or a library of object files, in the format produced zero or more object files, using ar . Implementations may supply additional utilities in these formats. Additional input file formats are implementation-defined.
8359 8360			
8361 8362 8363 8364		LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
8365 8366		LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
8367 8368 8369		LC_CTYPE	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).
8370 8371 8372		LC_MESSAC	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
8373	XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
8374 8375 8376	XSI	TMPDIR	Provide a pathname that should override the default directory for temporary files, if any. On XSI-conforming systems, provide a pathname that shall override the default directory for temporary files, if any.
8377 8378	ASYNC	HRONOUS I Default.	EVENTS
8379 8380 8381	STDOU		n one $\it file$ operand ending in $\it .c$ (or possibly other unspecified suffixes) is given, for e:
8382		"%s:\n",	<file></file>
8383 8384 8385			ten. These messages, if written, shall precede the processing of each input file; they written to the standard output if they are written to the standard error, as described RR section.
8386			ion is specified, the standard output shall be a text file that represents the results of

subsequent compilation passes.

8387

8388

the preprocessing stage of the language; it may contain extra information appropriate for

c99 Utilities

STDERR 8389

8390 The standard error shall be used only for diagnostic messages. If more than one file operand ending in .c (or possibly other unspecified suffixes) is given, for each such file: 8391

"%s:\n", <file> 8392

8393 may be written to allow identification of the diagnostic and warning messages with the 8394 appropriate input file. These messages, if written, shall precede the processing of each input file; they shall not be written to the standard error if they are written to the standard output, as 8395 described in the STDOUT section. 8396

This utility may produce warning messages about certain conditions that do not warrant 8397 returning an error (non-zero) exit value. 8398

OUTPUT FILES 8399

8400

8402

Object files or executable files or both are produced in unspecified formats.

EXTENDED DESCRIPTION 8401

Standard Libraries

8403	The c99 utili	ty shall recognize the following – l operands for standard libraries:
8404 8405 8406 8407 8408 8409 8410 8411	−l c	This operand shall make visible all functions referenced in the System Interfaces volume of IEEE Std 1003.1-2001, with the possible exception of those functions listed as residing in <aio.h>, <arpa inet.h="">, <complex.h>, <fenv.h>, <math.h>, <mqueue.h>, <netdb.h>, <netinet in.h="">, <pthread_h>, <sched.h>, <semaphore.h>, <spawn.h>, <sys socket.h="">, pthread_kill(), and pthread_sigmask() in <signal.h>, <trace.h>, functions marked as extensions other than as part of the MF or MPR extensions in <sys mman.h="">, functions marked as ADV in <fcntl.h>, and functions marked as CS, CPT, and TMR in <time.h>. This operand shall not be required to be present to cause a search of this library.</time.h></fcntl.h></sys></trace.h></signal.h></sys></spawn.h></semaphore.h></sched.h></pthread_h></netinet></netdb.h></mqueue.h></math.h></fenv.h></complex.h></arpa></aio.h>
8413 8414	-l l	This operand shall make visible all functions required by the C-language output of lex that are not made available through the $-\mathbf{l}$ \mathbf{c} operand.
8415 8416 8417	-l pthread	This operand shall make visible all functions referenced in <pre>pthread.h></pre> and <pre>pthread_kill()</pre> and <pre>pthread_sigmask()</pre> referenced in <pre>signal.h></pre> . An implementation may search this library in the absence of this operand.
8418 8419 8420	−l m	This operand shall make visible all functions referenced in <math.h>, <complex.h>, and <fenv.h>. An implementation may search this library in the absence of this operand.</fenv.h></complex.h></math.h>
8421 8422 8423 8424 8425	−l rt	This operand shall make visible all functions referenced in <aio.h>, <mqueue.h>, <sched.h>, <semaphore.h>, and <spawn.h>, functions marked as extensions other than as part of the MF or MPR extensions in <sys mman.h="">, functions marked as ADV in <fcntl.h>, and functions marked as CS, CPT, and TMR in <time.h>. An implementation may search this library in the absence of this operand.</time.h></fcntl.h></sys></spawn.h></semaphore.h></sched.h></mqueue.h></aio.h>
8426 8427	-l trace	This operand shall make visible all functions referenced in <trace.h>. An implementation may search this library in the absence of this operand.</trace.h>
8428 8429 8430	-l xnet	This operand makes visible all functions referenced in <arpa inet.h="">, <netdb.h>, <netinet in.h="">, and <sys socket.h="">. An implementation may search this library in the absence of this operand.</sys></netinet></netdb.h></arpa>
8431 8432	-l y	This operand shall make visible all functions required by the C-language output of $yacc$ that are not made available through the $-\mathbf{l}$ \mathbf{c} operand.

Utilities c99

In the absence of options that inhibit invocation of the link editor, such as $-\mathbf{c}$ or $-\mathbf{E}$, the c99 utility shall cause the equivalent of a $-\mathbf{l}$ \mathbf{c} operand to be passed to the link editor as the last $-\mathbf{l}$ operand, causing it to be searched after all other object files and libraries are loaded.

It is unspecified whether the libraries **libc.a**, **libm.a**, **librt.a**, **libpthread.a**, **libl.a**, **liby.a**, or **libxnet.a** exist as regular files. The implementation may accept as $-\mathbf{l}$ operands names of objects that do not exist as regular files.

External Symbols

The C compiler and link editor shall support the significance of external symbols up to a length of at least 31 bytes; the action taken upon encountering symbols exceeding the implementation-defined maximum symbol length is unspecified.

The compiler and link editor shall support a minimum of 511 external symbols per source or object file, and a minimum of 4095 external symbols in total. A diagnostic message shall be written to the standard output if the implementation-defined limit is exceeded; other actions are unspecified.

Programming Environments

All implementations shall support one of the following programming environments as a default. Implementations may support more than one of the following programming environments. Applications can use <code>sysconf()</code> or <code>getconf</code> to determine which programming environments are supported.

Programming Environment getconf Name	Bits in int	Bits in long	Bits in pointer	Bits in off_t
_POSIX_V6_ILP32_OFF32	32	32	32	32
_POSIX_V6_ILP32_OFFBIG	32	32	32	≥64

≥32

≥64

≥64

≥64

Table 4-4 Programming Environments: Type Sizes

All implementations shall support one or more environments where the widths of the following types are no greater than the width of type **long**:

blksize_t, cc_t, mode_t, nfds_t, pid_t, ptrdiff_t, size_t, speed_t, ssize_t, suseconds_t, tcflag_t, useconds_t, wchar_t, wint_t

The executable files created when these environments are selected shall be in a proper format for execution by the *exec* family of functions. Each environment may be one of the ones in Table 4-4, or it may be another environment. The names for the environments that meet this requirement shall be output by a *getconf* command using the _POSIX_V6_WIDTH_RESTRICTED_ENVS argument. If more than one environment meets the requirement, the names of all such environments shall be output on separate lines. Any of these names can then be used in a subsequent *getconf* command to obtain the flags specific to that environment with the following suffixes added as appropriate:

- _CFLAGS To get the C compiler flags.
- LDFLAGS To get the linker/loader flags.
- LIBS To get the libraries.
- This requirement may be removed in a future version of IEEE Std 1003.1.

POSIX_V6_LP64_OFF64

POSIX_V6_LPBIG_OFFBIG

c99 Utilities

When this utility processes a file containing a function called main(), it shall be defined with a return type equivalent to int. Using return from the initial call to main() shall be equivalent (other than with respect to language scope issues) to calling exit() with the returned value. Reaching the end of the initial call to main() shall be equivalent to calling exit(0). The implementation shall not declare a prototype for this function.

Implementations provide configuration strings for C compiler flags, linker/loader flags, and libraries for each supported environment. When an application needs to use a specific programming environment rather than the implementation default programming environment while compiling, the application shall first verify that the implementation supports the desired environment. If the desired programming environment is supported, the application shall then invoke *c99* with the appropriate C compiler flags as the first options for the compile, the appropriate linker/loader flags after any other options but before any operands, and the appropriate libraries at the end of the operands.

Conforming applications shall not attempt to link together object files compiled for different programming models. Applications shall also be aware that binary data placed in shared memory or in files might not be recognized by applications built for other programming models.

Programming Environment c99 and cc Arguments getconf Name Use getconf Name _POSIX_V6_ILP32_OFF32 C Compiler Flags POSIX V6 ILP32 OFF32 CFLAGS Linker/Loader Flags POSIX_V6_ILP32_OFF32_LDFLAGS Libraries POSIX_V6_ILP32_OFF32_LIBS _POSIX_V6_ILP32_OFFBIG C Compiler Flags POSIX_V6_ILP32_OFFBIG_CFLAGS Linker/Loader Flags POSIX V6 ILP32 OFFBIG LDFLAGS POSIX_V6_ILP32_OFFBIG_LIBS Libraries C Compiler Flags _POSIX_V6_LP64_OFF64 POSIX V6 LP64 OFF64 CFLAGS Linker/Loader Flags POSIX_V6_LP64_OFF64_LDFLAGS Libraries POSIX_V6_LP64_OFF64_LIBS _POSIX_V6_LPBIG_OFFBIG C Compiler Flags POSIX V6 LPBIG OFFBIG CFLAGS Linker/Loader Flags POSIX V6 LPBIG OFFBIG LDFLAGS Libraries POSIX_V6_LPBIG_OFFBIG_LIBS

Table 4-5 Programming Environments: *c99* and *cc* Arguments

EXIT STATUS

8475

8476

8477

8478 8479

8480

8481 8482

8483

8484

8485

8486

8487

8488

8489

8490

8491 8492

8493 8494

8495

8496

8497

8498

8499

8500

8501 8502

8503

8505

8506

8507

8508 8509

8510

8511 8512

8513

8514 8515

8516

The following exit values shall be returned:

- 0 Successful compilation or link edit.
- >0 An error occurred.

CONSEQUENCES OF ERRORS

When *c99* encounters a compilation error that causes an object file not to be created, it shall write a diagnostic to standard error and continue to compile other source code operands, but it shall not perform the link phase and return a non-zero exit status. If the link edit is unsuccessful, a diagnostic message shall be written to standard error and *c99* exits with a non-zero status. A conforming application shall rely on the exit status of *c99*, rather than on the existence or mode of the executable file.

Utilities c99

APPLICATION USAGE

Since the *c99* utility usually creates files in the current directory during the compilation process, it is typically necessary to run the *c99* utility in a directory in which a file can be created.

On systems providing POSIX Conformance (see the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 2, Conformance), *c99* is required only with the C-Language Development option; XSI-conformant systems always provide *c99*.

Some historical implementations have created .o files when -c is not specified and more than one source file is given. Since this area is left unspecified, the application cannot rely on .o files being created, but it also must be prepared for any related .o files that already exist being deleted at the completion of the link edit.

Some historical implementations have permitted -L options to be interspersed with -l operands on the command line. For an application to compile consistently on systems that do not behave like this, it is necessary for a conforming application to supply all -L options before any of the -l options.

There is the possible implication that if a user supplies versions of the standard functions (before they would be encountered by an implicit $-\mathbf{l} \ \mathbf{c}$ or explicit $-\mathbf{l} \ \mathbf{m}$), that those versions would be used in place of the standard versions. There are various reasons this might not be true (functions defined as macros, manipulations for clean name space, and so on), so the existence of files named in the same manner as the standard libraries within the $-\mathbf{L}$ directories is explicitly stated to produce unspecified behavior.

All of the functions specified in the System Interfaces volume of IEEE Std 1003.1-2001 may be made visible by implementations when the Standard C Library is searched. Conforming applications must explicitly request searching the other standard libraries when functions made visible by those libraries are used.

EXAMPLES

1. The following usage example compiles **foo.c** and creates the executable file **foo**:

```
c99 -o foo foo.c
```

The following usage example compiles **foo.c** and creates the object file **foo.o**:

```
c99 -c foo.c
```

The following usage example compiles **foo.c** and creates the executable file **a.out**:

```
c99 foo.c
```

The following usage example compiles **foo.c**, links it with **bar.o**, and creates the executable file **a.out**. It may also create and leave **foo.o**:

```
c99 foo.c bar.o
```

2. The following example shows how an application using threads interfaces can test for support of and use a programming environment supporting 32-bit **int**, **long**, and **pointer** types and an **off_t** type using at least 64 bits:

```
if [ $(getconf _POSIX_V6_ILP32_OFFBIG) != "-1" ]

then

c99 $(getconf POSIX_V6_ILP32_OFFBIG_CFLAGS) -D_XOPEN_SOURCE=600 \

$(getconf POSIX_V6_ILP32_OFFBIG_LDFLAGS) foo.c -o foo \

$(getconf POSIX_V6_ILP32_OFFBIG_LIBS) -l pthread

$(getconf POSIX_V6_ILP32_OFFBIG_LIBS) -l pthread

clse

cho ILP32 OFFBIG programming environment not supported
```

c99 Utilities

```
8561 exit 1
8562 fi
```

8563

8564

8565

8566

8567

8568 8569

8570

8571

8572

8573

8574 8575

8576

8577 8578

8579

8580

8581

3. The following examples clarify the use and interactions of –L options and –l operands.

Consider the case in which module **a.c** calls function f() in library **libQ.a**, and module **b.c** calls function g() in library **libp.a**. Assume that both libraries reside in /a/b/c. The command line to compile and link in the desired way is:

```
c99 -L /a/b/c main.o a.c -l Q b.c -l p
```

In this case the -l Q operand need only precede the first -l p operand, since both libQ.a and libp.a reside in the same directory.

Multiple -L operands can be used when library name collisions occur. Building on the previous example, suppose that the user wants to use a new **libp.a**, in /a/a/a, but still wants f() from /a/b/c/libQ.a:

```
c99 -L /a/a/a -L /a/b/c main.o a.c -l Q b.c -l p
```

In this example, the linker searches the -L options in the order specified, and finds /a/a/a/libp.a before /a/b/c/libp.a when resolving references for b.c. The order of the -l operands is still important, however.

4. The following example shows how an application can use a programming environment where the widths of the following types:

blksize_t, cc_t, mode_t, nfds_t, pid_t, ptrdiff_t, size_t, speed_t, ssize_t, suseconds_t, tcflag_t, useconds_t, wchar_t, wint_t

are no greater than the width of type **long**:

```
8582
                # First choose one of the listed environments ...
                # ... if there are no additional constraints, the first one will do:
8583
                CENV=$(getconf POSIX V6 WIDTH RESTRICTED ENVS | head -n 1)
8584
                # ... or, if an environment that supports large files is preferred,
8585
8586
                # look for names that contain "OFF64" or "OFFBIG". (This chooses
                # the last one in the list if none match.)
8587
                for CENV in $(getconf _POSIX_V6_WIDTH_RESTRICTED_ENVS)
8588
8589
                do
8590
                    case $CENV in
8591
                    *OFF64* | *OFFBIG*) break ;;
8592
                    esac
                done
8593
                # The chosen environment name can now be used like this:
8594
                c99 $(getconf ${CENV} CFLAGS) -D POSIX C SOURCE=200112L \
8595
                (getconf \{CENV\}_LDFLAGS) foo.c -o foo \
8596
                $(getconf ${CENV} LIBS)
8597
```

RATIONALE

8598

8599

8600 8601

8602

The c99 utility is based on the c89 utility originally introduced in the ISO POSIX-2: 1993 standard.

Some of the changes from *c89* include the modification to the contents of the Standard Libraries section to account for new headers and options; for example, <**spawn.h**> added to the -**l rt** operand, and the -**l** trace operand added for the Tracing functions.

Utilities c99

8603 8604	FUTURE DIRECTIONS None.
8605 8606 8607 8608	SEE ALSO Section 1.7.1.4 (on page 4), ar, getconf, make, nm, strip, umask, the System Interfaces volume of IEEE Std 1003.1-2001, exec, sysconf(), the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 13, Headers
8609 8610	CHANGE HISTORY First released in Issue 6. Included for alignment with the ISO/IEC 9899: 1999 standard.
8611 8612 8613	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/12 is applied, correcting the EXTENDED DESCRIPTION of –l c and –l m . Previously, the text did not take into account the presence of the <i>c99</i> math headers.
8614 8615	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/13 is applied, changing the reference to the libxnet library to libxnet.a .

cal Utilities

8616	NAME		
8617		cal — print a	a calendar
8618	SYNOP		
8619 8620	XSI	cal [[mon	th] year]
8621	DESCR	IPTION	
8622	DESCR		y shall write a calendar to standard output using the Julian calendar for dates from
8623		January 1, 1	through September 2, 1752 and the Gregorian calendar for dates from September 14,
8624		_	th December 31, 9999 as though the Gregorian calendar had been adopted on
8625		September 1	4, 1/52.
8626 8627	OPTIO	NS None.	
8628	OPERA	NDS	
8629		The following	g operands shall be supported:
8630 8631		month	Specify the month to be displayed, represented as a decimal integer from 1 (January) to 12 (December). The default shall be the current month.
8632 8633		year	Specify the year for which the calendar is displayed, represented as a decimal integer from 1 to 9999. The default shall be the current year.
8634	STDIN		
8635		Not used.	
8636	INPUT		
8637		None.	
8638 8639	ENVIR	ONMENT VA The followin	ARIABLES ag environment variables shall affect the execution of <i>cal</i> :
8640 8641 8642 8643		LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
8644 8645		LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
8646 8647 8648		LC_CTYPE	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).
8649		LC_MESSAC	GES
8650			Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error, and informative messages written
8651 8652			to standard output.
8653		LC_TIME	Determine the format and contents of the calendar.
8654		NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .

8655

TZ

Determine the timezone used to calculate the value of the current month.

Utilities cal

ASYNCHRONOUS EVENTS 8656 8657 Default. **STDOUT** 8658 The standard output shall be used to display the calendar, in an unspecified format. 8659 **STDERR** 8660 The standard error shall be used only for diagnostic messages. 8661 **OUTPUT FILES** 8662 None. 8663 8664 EXTENDED DESCRIPTION None. 8665 **EXIT STATUS** 8666 The following exit values shall be returned: 8667 8668 Successful completion. >0 An error occurred. 8669 **CONSEQUENCES OF ERRORS** 8670 Default. 8671 APPLICATION USAGE Note that: 8673 8674 cal 83 refers to A.D. 83, not 1983. 8675 **EXAMPLES** 8676 None. 8677 **RATIONALE** 8678 8679 None. **FUTURE DIRECTIONS** 8680 A future version of IEEE Std 1003.1-2001 may support locale-specific recognition of the date of 8681 adoption of the Gregorian calendar. 8682 **SEE ALSO** 8683 None. 8684 **CHANGE HISTORY** 8685 First released in Issue 2. 8686 Issue 6 8687 The DESCRIPTION is updated to allow for traditional behavior for years before the adoption of

8688

8689

the Gregorian calendar.

cat **Utilities**

8690 NAME cat — concatenate and print files 8691 8692 **SYNOPSIS** cat [-u] [file ...] 8693 DESCRIPTION 8694 The cat utility shall read files in sequence and shall write their contents to the standard output in 8695 the same sequence. 8696 **OPTIONS** 8697 The cat utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, 8698 Utility Syntax Guidelines. 8699 The following option shall be supported: 8700 Write bytes from the input file to the standard output without delay as each is 8701 -u read. 8702 **OPERANDS** 8703 The following operand shall be supported: 8704 file A pathname of an input file. If no file operands are specified, the standard input 8705 shall be used. If a file is '-', the cat utility shall read from the standard input at 8706 that point in the sequence. The *cat* utility shall not close and reopen standard input 8707 when it is referenced in this way, but shall accept multiple occurrences of '-' as a 8708 8709 file operand. **STDIN** 8710 The standard input shall be used only if no *file* operands are specified, or if a *file* operand is '-'. 8711 See the INPUT FILES section. 8712 8713 **INPUT FILES** The input files can be any file type. 8714 8715 **ENVIRONMENT VARIABLES** The following environment variables shall affect the execution of *cat*: 8716 8717 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 8718 Internationalization Variables for the precedence of internationalization variables 8719 used to determine the values of locale categories.) 8720 LC ALL If set to a non-empty string value, override the values of all the other 8721 internationalization variables. 8722 Determine the locale for the interpretation of sequences of bytes of text data as 8723 LC_CTYPE characters (for example, single-byte as opposed to multi-byte characters in 8724 arguments). 8725 LC_MESSAGES 8726 Determine the locale that should be used to affect the format and contents of 8727 8728 diagnostic messages written to standard error. **NLSPATH** Determine the location of message catalogs for the processing of LC MESSAGES. 8729 ASYNCHRONOUS EVENTS

Default.

8730

8731

Utilities cat

```
8732
     STDOUT
8733
              The standard output shall contain the sequence of bytes read from the input files. Nothing else
8734
              shall be written to the standard output.
     STDERR
8735
8736
              The standard error shall be used only for diagnostic messages.
     OUTPUT FILES
8737
              None.
8738
     EXTENDED DESCRIPTION
8739
              None.
     EXIT STATUS
8741
              The following exit values shall be returned:
8742
                  All input files were output successfully.
8743
                  An error occurred.
8744
     CONSEQUENCES OF ERRORS
8745
8746
              Default.
     APPLICATION USAGE
              The –u option has value in prototyping non-blocking reads from FIFOs. The intent is to support
8748
8749
              the following sequence:
8750
              mkfifo foo
              cat -u foo > /dev/tty13 &
8751
8752
              cat -u > foo
8753
              It is unspecified whether standard output is or is not buffered in the default case. This is
              sometimes of interest when standard output is associated with a terminal, since buffering may
8754
              delay the output. The presence of the –u option guarantees that unbuffered I/O is available. It is
8755
              implementation-defined whether the cat utility buffers output if the -\mathbf{u} option is not specified.
8756
8757
              Traditionally, the -u option is implemented using the equivalent of the setvbuf() function
              defined in the System Interfaces volume of IEEE Std 1003.1-2001.
8758
8759
     EXAMPLES
              The following command:
8760
              cat myfile
8761
8762
              writes the contents of the file myfile to standard output.
              The following command:
8763
              cat doc1 doc2 > doc.all
8764
              concatenates the files doc1 and doc2 and writes the result to doc.all.
8765
8766
              Because of the shell language mechanism used to perform output redirection, a command such
              as this:
8767
8768
              cat doc doc.end > doc
              causes the original data in doc to be lost.
8769
8770
              The command:
```

cat start - middle - end > file

8771

cat Utilities

when standard input is a terminal, gets two arbitrary pieces of input from the terminal with a single invocation of *cat*. Note, however, that if standard input is a regular file, this would be equivalent to the command:

```
cat start - middle /dev/null end > file
```

because the entire contents of the file would be consumed by *cat* the first time '-' was used as a *file* operand and an end-of-file condition would be detected immediately when '-' was referenced the second time.

RATIONALE

8772

8773

8774

8775 8776

8777

8778

8779

8780

8781

8782

8783

8784 8785

8786

8787

8807

8808 8809

8810 8811

8812

8813

8816

Historical versions of the *cat* utility include the options $-\mathbf{e}$, $-\mathbf{t}$, and $-\mathbf{v}$, which permit the ends of lines, <tab>s, and invisible characters, respectively, to be rendered visible in the output. The standard developers omitted these options because they provide too fine a degree of control over what is made visible, and similar output can be obtained using a command such as:

```
sed -n -e 's/$/$/' -e l pathname
```

The -s option was omitted because it corresponds to different functions in BSD and System V-based systems. The BSD -s option to squeeze blank lines can be accomplished by the shell script shown in the following example:

```
sed -n '
8788
            # Write non-empty lines.
8789
8790
            /./
8791
                   р
                   d
8792
8793
            # Write a single empty line, then look for more empty lines.
8794
            /^$/
8795
            # Get next line, discard the held <newline> (empty line),
8796
            # and look for more empty lines.
8797
8798
            :Empty
8799
            /^$/
8800
                  Ν
8801
                   s/.//
8802
                   b Empty
8803
            # Write the non-empty line before going back to search
            # for the first in a set of empty lines.
8805
8806
```

The System V –**s** option to silence error messages can be accomplished by redirecting the standard error. Note that the BSD documentation for *cat* uses the term "blank line" to mean the same as the POSIX "empty line": a line consisting only of a <newline>.

The BSD $-\mathbf{n}$ option was omitted because similar functionality can be obtained from the $-\mathbf{n}$ option of the pr utility.

FUTURE DIRECTIONS

8814 None.

8815 SEE ALSO

more, the System Interfaces volume of IEEE Std 1003.1-2001, *setvbuf()*

Utilities cat

8817 CHANGE HISTORY

First released in Issue 2.

cd Utilities

```
      8819
      NAME

      8820
      cd — change the working directory

      8821
      SYNOPSIS

      8822
      cd [-L | -P] [directory]

      8823
      cd -
```

DESCRIPTION

 The *cd* utility shall change the working directory of the current shell execution environment (see Section 2.12 (on page 61)) by executing the following steps in sequence. (In the following steps, the symbol **curpath** represents an intermediate value used to simplify the description of the algorithm used by *cd*. There is no requirement that **curpath** be made visible to the application.)

- 1. If no *directory* operand is given and the *HOME* environment variable is empty or undefined, the default behavior is implementation-defined and no further steps shall be taken.
- 2. If no *directory* operand is given and the *HOME* environment variable is set to a non-empty value, the *cd* utility shall behave as if the directory named in the *HOME* environment variable was specified as the *directory* operand.
- 3. If the *directory* operand begins with a slash character, set **curpath** to the operand and proceed to step 7.
- 4. If the first component of the *directory* operand is dot or dot-dot, proceed to step 6.
- 5. Starting with the first pathname in the colon-separated pathnames of *CDPATH* (see the ENVIRONMENT VARIABLES section) if the pathname is non-null, test if the concatenation of that pathname, a slash character, and the *directory* operand names a directory. If the pathname is null, test if the concatenation of dot, a slash character, and the operand names a directory. In either case, if the resulting string names an existing directory, set **curpath** to that string and proceed to step 7. Otherwise, repeat this step with the next pathname in *CDPATH* until all pathnames have been tested.
- 6. Set **curpath** to the string formed by the concatenation of the value of *PWD*, a slash character, and the operand.
- 7. If the **-P** option is in effect, the *cd* utility shall perform actions equivalent to the *chdir()* function, called with **curpath** as the *path* argument. If these actions succeed, the *PWD* environment variable shall be set to an absolute pathname for the current working directory and shall not contain filename components that, in the context of pathname resolution, refer to a file of type symbolic link. If there is insufficient permission on the new directory, or on any parent of that directory, to determine the current working directory, the value of the *PWD* environment variable is unspecified. If the actions equivalent to *chdir()* fail for any reason, the *cd* utility shall display an appropriate error message and not alter the *PWD* environment variable. Whether the actions equivalent to *chdir()* succeed or fail, no further steps shall be taken.
- 8. The **curpath** value shall then be converted to canonical form as follows, considering each component from beginning to end, in sequence:
 - Dot components and any slashes that separate them from the next component shall be deleted.
 - b. For each dot-dot component, if there is a preceding component and it is neither root nor dot-dot, the preceding component, all slashes separating the preceding component from dot-dot, dot-dot and all slashes separating dot-dot from the following component shall be deleted.

Utilities cd

8865 c. An implementation may further simplify curpath by removing any trailing slash characters that are not also leading slashes, replacing multiple non-leading 8866 consecutive slashes with a single slash, and replacing three or more leading slashes with a single slash. If, as a result of this canonicalization, the **curpath** variable is null, 8868 8869 no further steps shall be taken. 9. The cd utility shall then perform actions equivalent to the chdir() function called with 8870 **curpath** as the *path* argument. If these actions failed for any reason, the *cd* utility shall 8871 8872 display an appropriate error message and no further steps shall be taken. The PWD environment variable shall be set to **curpath**. 8873 8874 If, during the execution of the above steps, the PWD environment variable is changed, the OLDPWD environment variable shall also be changed to the value of the old working directory 8875 (that is the current working directory immediately prior to the call to *cd*). 8876 **OPTIONS** 8877 The *cd* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, 8878 Utility Syntax Guidelines. 8879 The following options shall be supported by the implementation: 8880 $-\mathbf{L}$ Handle the operand dot-dot logically; symbolic link components shall not be 8881 resolved before dot-dot components are processed (see steps 8. and 9. in the 8882 DESCRIPTION). 8883 $-\mathbf{P}$ Handle the operand dot-dot physically; symbolic link components shall be 8884 resolved before dot-dot components are processed (see step 7. in the 8885 8886 DESCRIPTION). If both -L and -P options are specified, the last of these options shall be used and all others 8887 ignored. If neither **–L** nor **–P** is specified, the operand shall be handled dot-dot logically; see the 8888 DESCRIPTION. 8889 **OPERANDS** 8890 8891 The following operands shall be supported: 8892 directory An absolute or relative pathname of the directory that shall become the new working directory. The interpretation of a relative pathname by cd depends on the 8893 -L option and the CDPATH and PWD environment variables. If directory is an 8894 empty string, the results are unspecified. 8895 When a hyphen is used as the operand, this shall be equivalent to the command: 8896 8897 cd "\$OLDPWD" && pwd which changes to the previous working directory and then writes its name. 8898 **STDIN** 8899 Not used. 8900 INPUT FILES 8901 None. 8902 **ENVIRONMENT VARIABLES** 8903 The following environment variables shall affect the execution of *cd*: 8904

A colon-separated list of pathnames that refer to directories. The cd utility shall use

this list in its attempt to change the directory, as described in the DESCRIPTION.

An empty string in place of a directory pathname represents the current directory.

If CDPATH is not set, it shall be treated as if it were an empty string.

CDPATH

8905

8906 8907

8908

 \mathbf{cd} **Utilities**

8909	НОМ	ЛE '	The name of the directory, used when no directory operand is specified.		
8910 8911 8912 8913	LAN		Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)		
8914 8915	LC_A		If set to a non-empty string value, override the values of all the other internationalization variables.		
8916 8917 8918	LC_C		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).		
8919	LC_N	MESSAGI	ES		
8920 8921			Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.		
8922	XSI NLSI	PATH :	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .		
8923	OLD	PWD .	A pathname of the previous working directory, used by cd –.		
8924 8925	PWD		This variable shall be set as specified in the DESCRIPTION. If an application sets or unsets the value of <i>PWD</i> , the behavior of <i>cd</i> is unspecified.		
8926	ASYNCHRO	NOUS E	VENTS		
8927	Defa	ult.			
8928					
8929 8930	If a non-empty directory name from $CDPATH$ is used, or if cd – is used, an absolute pathname of the new working directory shall be written to the standard output as follows:				
8931	"%s\	n'', < n	ew directory>		
8932	Otherwise, there shall be no output.				
8933 8934					
8935	OUTPUT FILES				
8936	None.				
8937 8938	··				
8939 8940					
8941	0 7	The direc	ctory was successfully changed.		
8942	>0 An error occurred.				
8943	CONSEQUENCES OF ERRORS				
8944			directory shall remain unchanged.		

Utilities cd

8945 APPLICATION USAGE

Since *cd* affects the current shell execution environment, it is always provided as a shell regular built-in. If it is called in a subshell or separate utility execution environment, such as one of the following:

8949 (cd /tmp) 8950 nohup cd

8951 find . -exec cd {} \;

it does not affect the working directory of the caller's environment.

The user must have execute (search) permission in *directory* in order to change to it.

8954 EXAMPLES

8955 None.

8956 RATIONALE

8957

8958

8959

8960

8961

8962

8963 8964

8969

8970 8971

8972

8973

8979

8981

8982

8983

The use of the *CDPATH* was introduced in the System V shell. Its use is analogous to the use of the *PATH* variable in the shell. The BSD C shell used a shell parameter *cdpath* for this purpose.

A common extension when *HOME* is undefined is to get the login directory from the user database for the invoking user. This does not occur on System V implementations.

Some historical shells, such as the KornShell, took special actions when the directory name contained a dot-dot component, selecting the logical parent of the directory, rather than the actual parent directory; that is, it moved up one level toward the '/' in the pathname, remembering what the user typed, rather than performing the equivalent of:

8965 chdir("..");

In such a shell, the following commands would not necessarily produce equivalent output for all directories:

8968 cd .. && ls ls ..

This behavior is now the default. It is not consistent with the definition of dot-dot in most historical practice; that is, while this behavior has been optionally available in the KornShell, other shells have historically not supported this functionality. The logical pathname is stored in the *PWD* environment variable when the *cd* utility completes and this value is used to construct the next directory name if *cd* is invoked with the –L option.

8974 FUTURE DIRECTIONS

8975 None.

8976 SEE ALSO

8977 Section 2.12 (on page 61), pwd, the System Interfaces volume of IEEE Std 1003.1-2001, chdir()

8978 CHANGE HISTORY

First released in Issue 2.

8980 **Issue 6**

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

• The *cd* – operand, *PWD*, and *OLDPWD* are added.

The -L and -P options are added to align with the IEEE P1003.2b draft standard. This also includes the introduction of a new description to include the effect of these options.

cd Utilities

8986 8987 IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/14 is applied, changing the SYNOPSIS to make it clear that the **–L** and **–P** options are mutually-exclusive.

Utilities cflow

8988 NAME

8989

8995

8996 8997

8999

9000

9001

9002

9007

9008

9009

9010

9011

9012

9013

9014

9015

9017

9018

9019 9020

9021

9022

9023

9025

9027

9029

cflow — generate a C-language flowgraph (**DEVELOPMENT**)

8990 SYNOPSIS

```
8991 XSI cflow [-r] [-d num] [-D name[=def]] ... [-i incl] [-I dir] ...
8992 [-U dir] ... file ...
8993
```

8994 **DESCRIPTION**

The *cflow* utility shall analyze a collection of object files or assembler, C-language, *lex*, or *yacc* source files, and attempt to build a graph, written to standard output, charting the external references.

8998 OPTIONS

The *cflow* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines, except that the order of the $-\mathbf{D}$, $-\mathbf{I}$, and $-\mathbf{U}$ options (which are identical to their interpretation by c99) is significant.

The following options shall be supported:

9003 —**d** num Indicate the depth at which the flowgraph is cut off. The application shall ensure that the argument num is a decimal integer. By default this is a very large number (typically greater than 32 000). Attempts to set the cut-off depth to a non-positive integer shall be ignored.

–i *incl* Increase the number of included symbols. The *incl* option-argument is one of the following characters:

- x Include external and static data symbols. The default shall be to include only functions in the flowgraph.
- _ (Underscore) Include names that begin with an underscore. The default shall be to exclude these functions (and data if -i x is used).

Reverse the caller:callee relationship, producing an inverted listing showing the callers of each function. The listing shall also be sorted in lexicographical order by callee.

9016 **OPERANDS**

 $-\mathbf{r}$

file

The following operand is supported:

The pathname of a file for which a graph is to be generated. Filenames suffixed by .l shall shall be taken to be *lex* input, .y as *yacc* input, .c as *c99* input, and .i as the output of *c99* –E. Such files shall be processed as appropriate, determined by their suffix.

Files suffixed by **.s** (conventionally assembler source) may have more limited information extracted from them.

9024 **STDIN**

Not used.

9026 INPUT FILES

The input files shall be object files or assembler, C-language, *lex*, or *yacc* source files.

9028 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *cflow*:

9030 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,

cflow Utilities

9032 Internationalization Variables for the precedence of internationalization variables 9033 used to determine the values of locale categories.) 9034 LC ALL If set to a non-empty string value, override the values of all the other internationalization variables. 9035 9036 LC_COLLATE Determine the locale for the ordering of the output when the **-r** option is used. 9037 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 9038 characters (for example, single-byte as opposed to multi-byte characters in 9039 arguments and input files). 9040 LC_MESSAGES 9041 Determine the locale that should be used to affect the format and contents of 9042 diagnostic messages written to standard error. 9043 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 9044 ASYNCHRONOUS EVENTS 9045 Default. 9046 **STDOUT** 9047 The flowgraph written to standard output shall be formatted as follows: 9048 "%d %s:%s\n", <reference number>, <global>, <definition> 9049 9050 Each line of output begins with a reference (that is, line) number, followed by indentation of at least one column position per level. This is followed by the name of the global, a colon, and its 9051 definition. Normally globals are only functions not defined as an external or beginning with an 9052 underscore; see the OPTIONS section for the -i inclusion option. For information extracted from 9053 C-language source, the definition consists of an abstract type declaration (for example, char *) 9054 and, delimited by angle brackets, the name of the source file and the line number where the 9055 definition was found. Definitions extracted from object files indicate the filename and location 9056 counter under which the symbol appeared (for example, *text*). 9057 Once a definition of a name has been written, subsequent references to that name contain only 9058 9059 the reference number of the line where the definition can be found. For undefined references, only "< >" shall be written. 9060 **STDERR** 9061 The standard error shall be used only for diagnostic messages. 9062 **OUTPUT FILES** 9063 None 9064 EXTENDED DESCRIPTION 9065 None. 9066 **EXIT STATUS** 9067 9068 The following exit values shall be returned: Successful completion. 9069 9070 >0 An error occurred. CONSEQUENCES OF ERRORS 9071 Default. 9072

Utilities cflow

```
APPLICATION USAGE
9073
9074
             Files produced by lex and yacc cause the reordering of line number declarations, and this can
             confuse cflow. To obtain proper results, the input of yacc or lex must be directed to cflow.
9075
     EXAMPLES
9076
             Given the following in file.c:
9077
9078
              int i;
              int f();
9079
             int g();
9080
9081
             int h();
9082
             int
9083
             main()
9084
9085
                   f();
9086
                   g();
9087
                   f();
              }
9088
              int
9089
             f()
9090
9091
                   i = h();
9092
9093
             The command:
9094
             cflow -i x file.c
9095
             produces the output:
9096
             1 main: int(), <file.c 6>
9097
9098
             2
                    f: int(), <file.c 13>
             3
                         h: <>
9099
9100
             4
                         i: int, <file.c 1>
              5
9101
                    g: <>
     RATIONALE
9102
9103
             None.
     FUTURE DIRECTIONS
9104
             None.
9105
9106
     SEE ALSO
              c99, lex, yacc
9107
     CHANGE HISTORY
9108
9109
             First released in Issue 2.
9110
     Issue 6
```

The normative text is reworded to avoid use of the term "must" for application requirements.

9111

chgrp Utilities

```
9112 NAME
9113 chgrp — change the file group ownership
9114 SYNOPSIS
9115 chgrp [-hR] group file ...
9116 chgrp -R [-H | -L | -P ] group file ...
```

DESCRIPTION

 The *chgrp* utility shall set the group ID of the file named by each *file* operand to the group ID specified by the *group* operand.

For each *file* operand, or, if the $-\mathbf{R}$ option is used, each file encountered while walking the directory trees specified by the *file* operands, the *chgrp* utility shall perform actions equivalent to the *chown*() function defined in the System Interfaces volume of IEEE Std 1003.1-2001, called with the following arguments:

- The file operand shall be used as the path argument.
- The user ID of the file shall be used as the owner argument.
- The specified group ID shall be used as the group argument.

Unless *chgrp* is invoked by a process with appropriate privileges, the set-user-ID and set-group-ID bits of a regular file shall be cleared upon successful completion; the set-user-ID and set-group-ID bits of other file types may be cleared.

9130 OPTIONS

The *chgrp* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported by the implementation:

- 9134 —h If the system supports group IDs for symbolic links, for each *file* operand that
 9135 names a file of type symbolic link, *chgrp* shall attempt to set the group ID of the
 9136 symbolic link instead of the file referenced by the symbolic link. If the system does
 9137 not support group IDs for symbolic links, for each *file* operand that names a file of
 9138 type symbolic link, *chgrp* shall do nothing more with the current file and shall go
 9139 on to any remaining files.
 - -H If the -R option is specified and a symbolic link referencing a file of type directory is specified on the command line, *chgrp* shall change the group of the directory referenced by the symbolic link and all files in the file hierarchy below it.
 - -L If the -R option is specified and a symbolic link referencing a file of type directory is specified on the command line or encountered during the traversal of a file hierarchy, *chgrp* shall change the group of the directory referenced by the symbolic link and all files in the file hierarchy below it.
 - -P If the −R option is specified and a symbolic link is specified on the command line or encountered during the traversal of a file hierarchy, *chgrp* shall change the group ID of the symbolic link if the system supports this operation. The *chgrp* utility shall not follow the symbolic link to any other part of the file hierarchy.
- 9151 —**R** Recursively change file group IDs. For each *file* operand that names a directory, 9152 chgrp shall change the group of the directory and all files in the file hierarchy below it. Unless a -**H**, -**L**, or -**P** option is specified, it is unspecified which of these options will be used as the default.

Utilities chgrp

9155 9156			more than one of the mutually-exclusive options $-\mathbf{H}$, $-\mathbf{L}$, and $-\mathbf{P}$ shall not be an error. The last option specified shall determine the behavior of the utility.
9157	OPERA	NDS	
9158			ng operands shall be supported:
9159 9160 9161 9162		group	A group name from the group database or a numeric group ID. Either specifies a group ID to be given to each file named by one of the <i>file</i> operands. If a numeric <i>group</i> operand exists in the group database as a group name, the group ID number associated with that group name is used as the group ID.
9163		file	A pathname of a file whose group ID is to be modified.
9164 9165	STDIN	Not used.	
9166 9167	INPUT	FILES None.	
9168 9169	ENVIR	ONMENT VA The followin	ARIABLES ng environment variables shall affect the execution of <i>chgrp</i> :
9170 9171 9172 9173		LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
9174 9175		LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
9176 9177 9178		LC_CTYPE	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).
9179		LC_MESSAC	GES
9180 9181		_	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
9182	XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .
9183 9184	ASYNC	HRONOUS I Default.	EVENTS
9185	STDOU	T	
9186		Not used.	
9187 9188	STDER		d error shall be used only for diagnostic messages.
9189 9190	OUTPU	T FILES None.	
9191 9192	EXTEN	DED DESCR l None.	IPTION
9193 9194	EXIT ST		ng exit values shall be returned:
9195		0 The util	ity executed successfully and all requested changes were made.

>0 An error occurred.

9196

chgrp Utilities

9197 CONSEQUENCES OF ERRORS

9198 Default.

9199 APPLICATION USAGE

9200 Only the owner of a file or the user with appropriate privileges may change the owner or group 9201 of a file.

Some implementations restrict the use of *chgrp* to a user with appropriate privileges when the *group* specified is not the effective group ID or one of the supplementary group IDs of the calling process.

9205 EXAMPLES

9206

9208

9209

9210

9211 9212

9213

9214

9220

9222

9223

None.

9207 RATIONALE

The System V and BSD versions use different exit status codes. Some implementations used the exit status as a count of the number of errors that occurred; this practice is unworkable since it can overflow the range of valid exit status values. The standard developers chose to mask these by specifying only 0 and >0 as exit values.

The functionality of *chgrp* is described substantially through references to *chown*(). In this way, there is no duplication of effort required for describing the interactions of permissions, multiple groups, and so on.

9215 FUTURE DIRECTIONS

9216 None.

9217 SEE ALSO

9218 chmod, chown, the System Interfaces volume of IEEE Std 1003.1-2001, chown()

9219 CHANGE HISTORY

First released in Issue 2.

9221 Issue 6

New options **–H**, **–L**, and **–P** are added to align with the IEEE P1003.2b draft standard. These options affect the processing of symbolic links.

9224 IEEE PASC Interpretation 1003.2 #172 is applied, changing the CONSEQUENCES OF ERRORS 9225 section to "Default.".

9226 IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/15 is applied, changing the SYNOPSIS to make it clear that **-h** and **-R** are optional.

Utilities chmod

9228 9229	NAME chmod — cl	nange the file modes			
9230	SYNOPSIS				
9231	chmod [-R] mode file				
9232 9233 9234	DESCRIPTION The <i>chmod</i> utility shall change any or all of the file mode bits of the file named by each <i>file</i> operand in the way specified by the <i>mode</i> operand.				
9235 9236 9237	additional	It is implementation-defined whether and how the <i>chmod</i> utility affects any alternate or additional file access control mechanism (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 4.4, File Access Permissions) being used for the specified file.			
9238 9239		cess whose effective user ID matches the user ID of the file, or a process with the privileges, shall be permitted to change the file mode bits of a file.			
9240 9241 9242		utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section Syntax Guidelines.			
9243	The following	ng option shall be supported:			
9244 9245 9246	–R	Recursively change file mode bits. For each <i>file</i> operand that names a directory, <i>chmod</i> shall change the file mode bits of the directory and all files in the file hierarchy below it.			
9247 9248	OPERANDS The following	ng operands shall be supported:			
9249 9250	mode	Represents the change to be made to the file mode bits of each file named by one of the <i>file</i> operands; see the EXTENDED DESCRIPTION section.			
9251	file	A pathname of a file whose file mode bits shall be modified.			
9252 9253	STDIN Not used.				
9254	INPUT FILES				
9255	None.				
9256 9257	ENVIRONMENT VA	ARIABLES ng environment variables shall affect the execution of <i>chmod</i> :			
9258 9259 9260 9261	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)			
9262 9263	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.			
9264 9265 9266	LC_CTYPE	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).			
9267	7 LC_MESSAGES				
9268		Determine the locale that should be used to affect the format and contents of			

9269

diagnostic messages written to standard error.

chmod Utilities

9270 XSI NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

9271 ASYNCHRONOUS EVENTS

Default.

9273 STDOUT

9272

9274

9279

9280

9281

9282

9283 9284

9285

9286 9287

9288

9289

9290

9291 9292

9293 9294

9295

9296

9297

9298

9299 9300

9301

9302

9303

9304 9305

9306 9307

9308

9309

9310 9311 Not used.

9275 STDERR

9276 The standard error shall be used only for diagnostic messages.

9277 OUTPUT FILES

9278 None.

EXTENDED DESCRIPTION

The *mode* operand shall be either a *symbolic_mode* expression or a non-negative octal integer. The *symbolic_mode* form is described by the grammar later in this section.

Each **clause** shall specify an operation to be performed on the current file mode bits of each *file*. The operations shall be performed on each *file* in the order in which the **clause**s are specified.

The **who** symbols **u**, **g**, and **o** shall specify the *user*, *group*, and *other* parts of the file mode bits, respectively. A **who** consisting of the symbol **a** shall be equivalent to **ugo**.

The **perm** symbols **r**, **w**, and **x** represent the *read*, *write*, and *execute/search* portions of file mode bits, respectively. The **perm** symbol **s** shall represent the *set-user-ID-on-execution* (when **who** contains or implies **u**) and *set-group-ID-on-execution* (when **who** contains or implies **g**) bits.

The **perm** symbol **X** shall represent the execute/search portion of the file mode bits if the file is a directory or if the current (unmodified) file mode bits have at least one of the execute bits (S_IXUSR, S_IXGRP, or S_IXOTH) set. It shall be ignored if the file is not a directory and none of the execute bits are set in the current file mode bits.

The **permcopy** symbols **u**, **g**, and **o** shall represent the current permissions associated with the user, group, and other parts of the file mode bits, respectively. For the remainder of this section, **perm** refers to the non-terminals **perm** and **permcopy** in the grammar.

If multiple **actionlists** are grouped with a single **wholist** in the grammar, each **actionlist** shall be applied in the order specified with that **wholist**. The *op* symbols shall represent the operation performed, as follows:

+ If **perm** is not specified, the '+' operation shall not change the file mode bits.

If **who** is not specified, the file mode bits represented by **perm** for the owner, group, and other permissions, except for those with corresponding bits in the file mode creation mask of the invoking process, shall be set.

Otherwise, the file mode bits represented by the specified **who** and **perm** values shall be set.

If perm is not specified, the '−' operation shall not change the file mode bits.

If **who** is not specified, the file mode bits represented by **perm** for the owner, group, and other permissions, except for those with corresponding bits in the file mode creation mask of the invoking process, shall be cleared.

Otherwise, the file mode bits represented by the specified **who** and **perm** values shall be cleared.

= Clear the file mode bits specified by the **who** value, or, if no **who** value is specified, all of the file mode bits specified in this volume of IEEE Std 1003.1-2001.

Utilities chmod

If **perm** is not specified, the '=' operation shall make no further modifications to the file mode bits.

If **who** is not specified, the file mode bits represented by **perm** for the owner, group, and other permissions, except for those with corresponding bits in the file mode creation mask of the invoking process, shall be set.

Otherwise, the file mode bits represented by the specified **who** and **perm** values shall be set.

When using the symbolic mode form on a regular file, it is implementation-defined whether or not:

- Requests to set the set-user-ID-on-execution or set-group-ID-on-execution bit when all execute bits are currently clear and none are being set are ignored.
- Requests to clear all execute bits also clear the set-user-ID-on-execution and set-group-ID-on-execution bits.
- Requests to clear the set-user-ID-on-execution or set-group-ID-on-execution bits when all execute bits are currently clear are ignored. However, if the command *ls* –*l file* writes an *s* in the position indicating that the set-user-ID-on-execution or set-group-ID-on-execution is set, the commands *chmod* **u**–**s** *file* or *chmod* **g**–**s** *file*, respectively, shall not be ignored.

When using the symbolic mode form on other file types, it is implementation-defined whether or not requests to set or clear the set-user-ID-on-execution or set-group-ID-on-execution bits are honored.

If the **who** symbol \mathbf{o} is used in conjunction with the **perm** symbol \mathbf{s} with no other **who** symbols being specified, the set-user-ID-on-execution and set-group-ID-on-execution bits shall not be modified. It shall not be an error to specify the **who** symbol \mathbf{o} in conjunction with the **perm** symbol \mathbf{s} .

The **perm** symbol **t** shall specify the S_ISVTX bit. When used with a file of type directory, it can be used with the **who** symbol **a**, or with no **who** symbol. It shall not be an error to specify a **who** symbol of **u**, **g**, or **o** in conjunction with the **perm** symbol **t**, but the meaning of these combinations is unspecified. The effect when using the **perm** symbol **t** with any file type other than directory is unspecified.

For an octal integer *mode* operand, the file mode bits shall be set absolutely.

For each bit set in the octal number, the corresponding file permission bit shown in the following table shall be set; all other file permission bits shall be cleared. For regular files, for each bit set in the octal number corresponding to the set-user-ID-on-execution or the set-group-ID-on-execution, bits shown in the following table shall be set; if these bits are not set in the octal number, they are cleared. For other file types, it is implementation-defined whether or not requests to set or clear the set-user-ID-on-execution or set-group-ID-on-execution bits are honored.

Octal	Mode Bit	Octal	Mode Bit	Octal	Mode Bit	Octal	Mode Bit
4000	S_ISUID	0400	S_IRUSR	0040	S_IRGRP	0004	S_IROTH
2000	S_ISGID	0200	S_IWUSR	0020	S_IWGRP	0002	S_IWOTH
1000	S_ISVTX	0100	S_IXUSR	0010	S_IXGRP	0001	S_IXOTH

XSI

 XSI

When bits are set in the octal number other than those listed in the table above, the behavior is unspecified.

chmod Utilities

Grammar for chmod

9354

9355 9356

9357

9358

9359

9360

9361

The grammar and lexical conventions in this section describe the syntax for the *symbolic_mode* operand. The general conventions for this style of grammar are described in Section 1.10 (on page 19). A valid *symbolic_mode* can be represented as the non-terminal symbol *symbolic_mode* in the grammar. This formal syntax shall take precedence over the preceding text syntax description.

The lexical processing is based entirely on single characters. Implementations need not allow

blank>s within the single argument being processed.

```
9362
                         symbolic mode
             응응
9363
9364
             symbolic_mode
                                  : clause
9365
                                    symbolic mode ',' clause
9366
9367
             clause
                                   actionlist
                                    wholist actionlist
9368
9369
             wholist
9370
                                   who
9371
                                    wholist who
9372
                                  ;
9373
             who
                                           'q' | 'o' | 'a'
9374
9375
             actionlist
                                  : action
9376
                                    actionlist action
9377
             action
9378
                                  :
                                    op
9379
                                    op permlist
9380
                                    op permcopy
9381
9382
             permcopy
9383
9384
             op
9385
9386
             permlist
                                    perm
9387
                                    perm permlist
9388
                                           'w' | 'x' | 'X' | 's' | 't'
9389
    XSI
                                    'r'
             perm
9390
```

9391 EXIT STATUS

9392

9393

The following exit values shall be returned:

- 0 The utility executed successfully and all requested changes were made.
- 9394 >0 An error occurred.

Utilities chmod

CONSEQUENCES OF ERRORS

Default.

9397 APPLICATION USAGE

Some implementations of the *chmod* utility change the mode of a directory before the files in the directory when performing a recursive (**-R** option) change; others change the directory mode after the files in the directory. If an application tries to remove read or search permission for a file hierarchy, the removal attempt fails if the directory is changed first; on the other hand, trying to re-enable permissions to a restricted hierarchy fails if directories are changed last. Users should not try to make a hierarchy inaccessible to themselves.

Some implementations of *chmod* never used the process' *umask* when changing modes; systems conformant with this volume of IEEE Std 1003.1-2001 do so when **who** is not specified. Note the difference between:

chmod a-w file

which removes all write permissions, and:

chmod -- -w file

which removes write permissions that would be allowed if **file** was created with the same *umask*.

Conforming applications should never assume that they know how the set-user-ID and setgroup-ID bits on directories are interpreted.

9414 EXAMPLES

Mode	Results
a+=	Equivalent to $a+,a=$; clears all file mode bits.
<i>go</i> +-w	Equivalent to $go+,go-w$; clears group and other write bits.
g=o-w	Equivalent to $g=o,g-w$; sets group bit to match other bits and then clears group write bit.
g-r+w	Equivalent to $g-r,g+w$; clears group read bit and sets group write bit.
uo=g	Sets owner bits to match group bits and sets other bits to match group bits.

RATIONALE

The functionality of *chmod* is described substantially through references to concepts defined in the System Interfaces volume of IEEE Std 1003.1-2001. In this way, there is less duplication of effort required for describing the interactions of permissions. However, the behavior of this utility is not described in terms of the *chmod()* function from the System Interfaces volume of IEEE Std 1003.1-2001 because that specification requires certain side effects upon alternate file access control mechanisms that might not be appropriate, depending on the implementation.

Implementations that support mandatory file and record locking as specified by the 1984 /usr/group standard historically used the combination of set-group-ID bit set and group execute bit clear to indicate mandatory locking. This condition is usually set or cleared with the symbolic mode **perm** symbol **l** instead of the **perm** symbols **s** and **x** so that the mandatory locking mode is not changed without explicit indication that that was what the user intended. Therefore, the details on how the implementation treats these conditions must be defined in the documentation. This volume of IEEE Std 1003.1-2001 does not require mandatory locking (nor does the System Interfaces volume of IEEE Std 1003.1-2001), but does allow it as an extension. However, this volume of IEEE Std 1003.1-2001 does require that the *ls* and *chmod* utilities work

chmod Utilities

consistently in this area. If ls –l file indicates that the set-group-ID bit is set, chmod g–s file must clear it (assuming appropriate privileges exist to change modes).

The System V and BSD versions use different exit status codes. Some implementations used the exit status as a count of the number of errors that occurred; this practice is unworkable since it can overflow the range of valid exit status values. This problem is avoided here by specifying only 0 and >0 as exit values.

The System Interfaces volume of IEEE Std 1003.1-2001 indicates that implementation-defined restrictions may cause the S_ISUID and S_ISGID bits to be ignored. This volume of IEEE Std 1003.1-2001 allows the *chmod* utility to choose to modify these bits before calling *chmod()* (or some function providing equivalent capabilities) for non-regular files. Among other things, this allows implementations that use the set-user-ID and set-group-ID bits on directories to enable extended features to handle these extensions in an intelligent manner.

The **X perm** symbol was adopted from BSD-based systems because it provides commonly desired functionality when doing recursive ($-\mathbf{R}$ option) modifications. Similar functionality is not provided by the *find* utility. Historical BSD versions of *chmod*, however, only supported **X** with op+; it has been extended in this volume of IEEE Std 1003.1-2001 because it is also useful with op=. (It has also been added for op- even though it duplicates **x**, in this case, because it is intuitive and easier to explain.)

The grammar was extended with the *permcopy* non-terminal to allow historical-practice forms of symbolic modes like $\mathbf{o}=\mathbf{u}-\mathbf{g}$ (that is, set the "other" permissions to the permissions of "owner" minus the permissions of "group").

9462 FUTURE DIRECTIONS

None.

9464 SEE ALSO

ls, *umask*, the System Interfaces volume of IEEE Std 1003.1-2001, *chmod*()

9466 CHANGE HISTORY

First released in Issue 2.

Issue 6

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

• Octal modes have been kept and made mandatory despite being marked obsolescent in the ISO POSIX-2:1993 standard.

IEEE PASC Interpretation 1003.2 #172 is applied, changing the CONSEQUENCES OF ERRORS section to "Default.".

The Open Group Base Resolution bwg2001-010 is applied, adding the description of the S_ISVTX bit and the **t perm** symbol as an XSI extension.

IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/16 is applied, changing the XSI shaded text in the EXTENDED DESCRIPTION from:

"The **perm** symbol **t** shall specify the S_ISVTX bit and shall apply to directories only. The effect when using it with any other file type is unspecified. It can be used with the **who** symbols **o**, **a**, or with no **who** symbol. It shall not be an error to specify a **who** symbol of **u** or **g** in conjunction with the **perm** symbol **t**; it shall be ignored for **u** and **g**."

9483 to:

Utilities chmod

9484	"The perm symbol t shall specify the S_ISVTX bit. When used with a file of type directory, it	I
9485	can be used with the who symbol a, or with no who symbol. It shall not be an error to specify	İ
9486	a who symbol of u, g, or o in conjunction with the perm symbol t, but the meaning of these	Ì
9487	combinations is unspecified. The effect when using the perm symbol t with any file type	Ĺ
9488	other than directory is unspecified.''	
9489	This change is to permit historical behavior.	

chown Utilities

```
9490 NAME
9491 chown — change the file ownership
9492 SYNOPSIS
9493 chown [-hR] owner[:group] file ...
9494 chown —R [-H | -L | -P ] owner[:group] file ...
```

DESCRIPTION

The *chown* utility shall set the user ID of the file named by each *file* operand to the user ID specified by the *owner* operand.

For each *file* operand, or, if the $-\mathbf{R}$ option is used, each file encountered while walking the directory trees specified by the *file* operands, the *chown* utility shall perform actions equivalent to the *chown*() function defined in the System Interfaces volume of IEEE Std 1003.1-2001, called with the following arguments:

- 1. The *file* operand shall be used as the *path* argument.
- 2. The user ID indicated by the *owner* portion of the first operand shall be used as the *owner* argument.
- 3. If the *group* portion of the first operand is given, the group ID indicated by it shall be used as the *group* argument; otherwise, the group ownership shall not be changed.

Unless *chown* is invoked by a process with appropriate privileges, the set-user-ID and set-group-ID bits of a regular file shall be cleared upon successful completion; the set-user-ID and set-group-ID bits of other file types may be cleared.

OPTIONS

The *chown* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported by the implementation:

- If the system supports user IDs for symbolic links, for each *file* operand that names a file of type symbolic link, *chown* shall attempt to set the user ID of the symbolic link. If the system supports group IDs for symbolic links, and a group ID was specified, for each *file* operand that names a file of type symbolic link, *chown* shall attempt to set the group ID of the symbolic link. If the system does not support user or group IDs for symbolic links, for each *file* operand that names a file of type symbolic link, *chown* shall do nothing more with the current file and shall go on to any remaining files.
- -H If the -R option is specified and a symbolic link referencing a file of type directory is specified on the command line, *chown* shall change the user ID (and group ID, if specified) of the directory referenced by the symbolic link and all files in the file hierarchy below it.
- -L If the -R option is specified and a symbolic link referencing a file of type directory is specified on the command line or encountered during the traversal of a file hierarchy, *chown* shall change the user ID (and group ID, if specified) of the directory referenced by the symbolic link and all files in the file hierarchy below it.
- -P If the -R option is specified and a symbolic link is specified on the command line or encountered during the traversal of a file hierarchy, *chown* shall change the owner ID (and group ID, if specified) of the symbolic link if the system supports this operation. The *chown* utility shall not follow the symbolic link to any other part of the file hierarchy.

Utilities chown

9535 $-\mathbf{R}$ Recursively change file user and group IDs. For each file operand that names a directory, chown shall change the user ID (and group ID, if specified) of the 9536 directory and all files in the file hierarchy below it. Unless a -H, -L, or -P option is 9537 specified, it is unspecified which of these options will be used as the default. 9538 Specifying more than one of the mutually-exclusive options -H, -L, and -P shall not be 9539 considered an error. The last option specified shall determine the behavior of the utility. 9540 **OPERANDS** 9541 The following operands shall be supported: 9542 owner[:group] A user ID and optional group ID to be assigned to file. The owner portion of this 9543 operand shall be a user name from the user database or a numeric user ID. Either 9544 specifies a user ID which shall be given to each file named by one of the file 9545 operands. If a numeric owner operand exists in the user database as a user name, 9546 9547 the user ID number associated with that user name shall be used as the user ID. Similarly, if the group portion of this operand is present, it shall be a group name 9548 from the group database or a numeric group ID. Either specifies a group ID which 9549 shall be given to each file. If a numeric group operand exists in the group database 9550 as a group name, the group ID number associated with that group name shall be 9551 used as the group ID. 9552 file A pathname of a file whose user ID is to be modified. 9553 STDIN 9554 9555 Not used. **INPUT FILES** 9556 None. 9557 ENVIRONMENT VARIABLES 9558 The following environment variables shall affect the execution of *chown*: 9559 LANG Provide a default value for the internationalization variables that are unset or null. 9560 9561 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables 9562 used to determine the values of locale categories.) 9563 LC ALL If set to a non-empty string value, override the values of all the other 9564 internationalization variables. 9565 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 9566 characters (for example, single-byte as opposed to multi-byte characters in 9567 arguments). 9568 LC_MESSAGES 9569 Determine the locale that should be used to affect the format and contents of 9570 diagnostic messages written to standard error. 9571 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 9572 XSI ASYNCHRONOUS EVENTS 9573 Default. 9574 **STDOUT** 9575 Not used. 9576

chown Utilities

9577 STDERR

9582

9584 9585

9586

9590 9591

9592

9594

9595 9596

9597

9598

9599

9600

9601

9602 9603

9604

9605

9606 9607

9608

9609

9610

9611

9612

9613

9614

9615 9616

9617

9618 9619

9620

9578 The standard error shall be used only for diagnostic messages.

9579 **OUTPUT FILES**

9580 None.

9581 EXTENDED DESCRIPTION

None.

9583 EXIT STATUS

The following exit values shall be returned:

- 0 The utility executed successfully and all requested changes were made.
- >0 An error occurred.

9587 CONSEQUENCES OF ERRORS

9588 Default.

9589 APPLICATION USAGE

Only the owner of a file or the user with appropriate privileges may change the owner or group of a file.

Some implementations restrict the use of *chown* to a user with appropriate privileges.

9593 EXAMPLES

None.

RATIONALE

The System V and BSD versions use different exit status codes. Some implementations used the exit status as a count of the number of errors that occurred; this practice is unworkable since it can overflow the range of valid exit status values. These are masked by specifying only 0 and >0 as exit values.

The functionality of *chown* is described substantially through references to functions in the System Interfaces volume of IEEE Std 1003.1-2001. In this way, there is no duplication of effort required for describing the interactions of permissions, multiple groups, and so on.

The 4.3 BSD method of specifying both owner and group was included in this volume of IEEE Std 1003.1-2001 because:

- There are cases where the desired end condition could not be achieved using the *chgrp* and *chown* (that only changed the user ID) utilities. (If the current owner is not a member of the desired group and the desired owner is not a member of the current group, the *chown*() function could fail unless both owner and group are changed at the same time.)
- Even if they could be changed independently, in cases where both are being changed, there is a 100% performance penalty caused by being forced to invoke both utilities.

The BSD syntax *user*[.*group*] was changed to *user*[:*group*] in this volume of IEEE Std 1003.1-2001 because the period is a valid character in login names (as specified by the Base Definitions volume of IEEE Std 1003.1-2001, login names consist of characters in the portable filename character set). The colon character was chosen as the replacement for the period character because it would never be allowed as a character in a user name or group name on historical implementations.

The $-\mathbf{R}$ option is considered by some observers as an undesirable departure from the historical UNIX system tools approach; since a tool, *find*, already exists to recurse over directories, there seemed to be no good reason to require other tools to have to duplicate that functionality. However, the $-\mathbf{R}$ option was deemed an important user convenience, is far more efficient than

Utilities chown

9621 9622	forking a separate process for each element of the directory hierarchy, and is in widespread historical use.
9623 9624	FUTURE DIRECTIONS None.
9625 9626	SEE ALSO chmod, chgrp, the System Interfaces volume of IEEE Std 1003.1-2001, chown()
9627 9628	CHANGE HISTORY First released in Issue 2.
9629 9630 9631	Issue 6 New options -h, -H, -L, and -P are added to align with the IEEE P1003.2b draft standard. These options affect the processing of symbolic links.
9632	The normative text is reworded to avoid use of the term "must" for application requirements.
9633 9634	IEEE PASC Interpretation 1003.2 #172 is applied, changing the CONSEQUENCES OF ERRORS section to "Default.".
9635 9636	The "otherwise," text in item 3. of the DESCRIPTION is changed to "otherwise, the group ownership shall not be changed".
9637 9638	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/17 is applied, changing the SYNOPSIS to make it clear that $-\mathbf{h}$ and $-\mathbf{R}$ are optional.

cksum Utilities

9639 NAME

9644

9645

9646

9647

9648

9650

9651

9652

9653

9654

9655

9656

9657

9658

9659

9660

9661

9662

9666

9670

9671

9673

9675

9640 cksum — write file checksums and sizes

9641 SYNOPSIS

9642 cksum [file ...]

9643 **DESCRIPTION**

The *cksum* utility shall calculate and write to standard output a cyclic redundancy check (CRC) for each input file, and also write to standard output the number of octets in each file. The CRC used is based on the polynomial used for CRC error checking in the ISO/IEC 8802-3: 1996 standard (Ethernet).

The encoding for the CRC checksum is defined by the generating polynomial:

9649
$$G(x) = x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^{8} + x^{7} + x^{5} + x^{4} + x^{2} + x + 1$$

Mathematically, the CRC value corresponding to a given file shall be defined by the following procedure:

- 1. The n bits to be evaluated are considered to be the coefficients of a mod 2 polynomial M(x) of degree n-1. These n bits are the bits from the file, with the most significant bit being the most significant bit of the first octet of the file and the last bit being the least significant bit of the last octet, padded with zero bits (if necessary) to achieve an integral number of octets, followed by one or more octets representing the length of the file as a binary value, least significant octet first. The smallest number of octets capable of representing this integer shall be used.
- 2. M(x) is multiplied by x^{32} (that is, shifted left 32 bits) and divided by G(x) using mod 2 division, producing a remainder R(x) of degree ≤ 31 .
- 3. The coefficients of R(x) are considered to be a 32-bit sequence.
- 4. The bit sequence is complemented and the result is the CRC.

9663 OPTIONS

9664 None.

9665 **OPERANDS**

The following operand shall be supported:

9667 *file* A pathname of a file to be checked. If no *file* operands are specified, the standard input shall be used.

9669 **STDIN**

The standard input shall be used only if no *file* operands are specified. See the INPUT FILES section.

9672 INPUT FILES

The input files can be any file type.

9674 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *cksum*:

9676 LANG Provide a default value for the internationalization variables that are unset or null.
9677 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
9678 Internationalization Variables for the precedence of internationalization variables
9679 used to determine the values of locale categories.)

9680 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

Utilities cksum

9682 LC_CTYPE 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 9683 arguments). 9684 LC MESSAGES 9685 Determine the locale that should be used to affect the format and contents of 9686 diagnostic messages written to standard error. 9687 **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 9688 XSI ASYNCHRONOUS EVENTS 9689 Default. 9690 **STDOUT** 9691 For each file processed successfully, the *cksum* utility shall write in the following format: 9692 "%u %d %s\n", <checksum>, <# of octets>, <pathname> 9693 If no *file* operand was specified, the pathname and its leading <space> shall be omitted. 9694 **STDERR** 9695 9696 The standard error shall be used only for diagnostic messages. **OUTPUT FILES** 9697 None. 9698 EXTENDED DESCRIPTION 9699 9700 None.

EXIT STATUS 9701

9702 9703

9706

9708 9709

9710

9711 9712

9713

9714

9715

9716

9717

9718 9719

9720

9721

The following exit values shall be returned:

All files were processed successfully.

>0 An error occurred.

CONSEQUENCES OF ERRORS 9705

Default.

APPLICATION USAGE 9707

The *cksum* utility is typically used to quickly compare a suspect file against a trusted version of the same, such as to ensure that files transmitted over noisy media arrive intact. However, this comparison cannot be considered cryptographically secure. The chances of a damaged file producing the same CRC as the original are small; deliberate deception is difficult, but probably not impossible.

Although input files to *cksum* can be any type, the results need not be what would be expected on character special device files or on file types not described by the System Interfaces volume of IEEE Std 1003.1-2001. Since this volume of IEEE Std 1003.1-2001 does not specify the block size used when doing input, checksums of character special files need not process all of the data in those files.

The algorithm is expressed in terms of a bitstream divided into octets. If a file is transmitted between two systems and undergoes any data transformation (such as changing little-endian byte ordering to big-endian), identical CRC values cannot be expected. Implementations performing such transformations may extend *cksum* to handle such situations.

cksum Utilities

```
9722 EXAMPLES
9723 None.
9724 RATIONALE
9725 The fo
```

9727

9728

The following C-language program can be used as a model to describe the algorithm. It assumes that a **char** is one octet. It also assumes that the entire file is available for one pass through the function. This was done for simplicity in demonstrating the algorithm, rather than as an implementation model.

```
static unsigned long crctab[] = {
9729
9730
           0x00000000,
9731
           0x04c11db7, 0x09823b6e, 0x0d4326d9, 0x130476dc, 0x17c56b6b,
           0x1a864db2, 0x1e475005, 0x2608edb8, 0x22c9f00f, 0x2f8ad6d6,
9732
           0x2b4bcb61, 0x350c9b64, 0x31cd86d3, 0x3c8ea00a, 0x384fbdbd,
9733
           0x4c11db70, 0x48d0c6c7, 0x4593e01e, 0x4152fda9, 0x5f15adac,
9734
           0x5bd4b01b, 0x569796c2, 0x52568b75, 0x6a1936c8, 0x6ed82b7f,
9735
           0x639b0da6, 0x675a1011, 0x791d4014, 0x7ddc5da3, 0x709f7b7a,
9736
           0x745e66cd, 0x9823b6e0, 0x9ce2ab57, 0x91a18d8e, 0x95609039,
9737
           0x8b27c03c, 0x8fe6dd8b, 0x82a5fb52, 0x8664e6e5, 0xbe2b5b58,
9738
           Oxbaea46ef, Oxb7a96036, Oxb3687d81, Oxad2f2d84, Oxa9ee3033,
9739
           0xa4ad16ea, 0xa06c0b5d, 0xd4326d90, 0xd0f37027, 0xddb056fe,
9740
           0xd9714b49, 0xc7361b4c, 0xc3f706fb, 0xceb42022, 0xca753d95,
9741
9742
           0xf23a8028, 0xf6fb9d9f, 0xfbb8bb46, 0xff79a6f1, 0xe13ef6f4,
           0xe5ffeb43, 0xe8bccd9a, 0xec7dd02d, 0x34867077, 0x30476dc0,
9743
           0x3d044b19, 0x39c556ae, 0x278206ab, 0x23431b1c, 0x2e003dc5,
9744
           0x2ac12072, 0x128e9dcf, 0x164f8078, 0x1b0ca6a1, 0x1fcdbb16,
9745
           0x018aeb13, 0x054bf6a4, 0x0808d07d, 0x0cc9cdca, 0x7897ab07,
9746
           0x7c56b6b0, 0x71159069, 0x75d48dde, 0x6b93dddb, 0x6f52c06c,
9747
           0x6211e6b5, 0x66d0fb02, 0x5e9f46bf, 0x5a5e5b08, 0x571d7dd1,
9748
           0x53dc6066, 0x4d9b3063, 0x495a2dd4, 0x44190b0d, 0x40d816ba,
9749
           0xaca5c697, 0xa864db20, 0xa527fdf9, 0xa1e6e04e, 0xbfa1b04b,
9750
9751
           0xbb60adfc, 0xb6238b25, 0xb2e29692, 0x8aad2b2f, 0x8e6c3698,
           0x832f1041, 0x87ee0df6, 0x99a95df3, 0x9d684044, 0x902b669d,
9752
9753
           0x94ea7b2a, 0xe0b41de7, 0xe4750050, 0xe9362689, 0xedf73b3e,
           0xf3b06b3b, 0xf771768c, 0xfa325055, 0xfef34de2, 0xc6bcf05f,
9754
           0xc27dede8, 0xcf3ecb31, 0xcbffd686, 0xd5b88683, 0xd1799b34,
9755
           0xdc3abded, 0xd8fba05a, 0x690ce0ee, 0x6dcdfd59, 0x608edb80,
9756
           0x644fc637, 0x7a089632, 0x7ec98b85, 0x738aad5c, 0x774bb0eb,
9757
           0x4f040d56, 0x4bc510e1, 0x46863638, 0x42472b8f, 0x5c007b8a,
9758
           0x58c1663d, 0x558240e4, 0x51435d53, 0x251d3b9e, 0x21dc2629,
9759
           0x2c9f00f0, 0x285e1d47, 0x36194d42, 0x32d850f5, 0x3f9b762c,
9760
           0x3b5a6b9b, 0x0315d626, 0x07d4cb91, 0x0a97ed48, 0x0e56f0ff,
9761
           0x1011a0fa, 0x14d0bd4d, 0x19939b94, 0x1d528623, 0xf12f560e,
9762
           0xf5ee4bb9, 0xf8ad6d60, 0xfc6c70d7, 0xe22b20d2, 0xe6ea3d65,
9763
9764
           0xeba91bbc, 0xef68060b, 0xd727bbb6, 0xd3e6a601, 0xdea580d8,
9765
           0xda649d6f, 0xc423cd6a, 0xc0e2d0dd, 0xcda1f604, 0xc960ebb3,
9766
           0xbd3e8d7e, 0xb9ff90c9, 0xb4bcb610, 0xb07daba7, 0xae3afba2,
           0xaafbe615, 0xa7b8c0cc, 0xa379dd7b, 0x9b3660c6, 0x9ff77d71,
9767
           0x92b45ba8, 0x9675461f, 0x8832161a, 0x8cf30bad, 0x81b02d74,
9768
           0x857130c3, 0x5d8a9099, 0x594b8d2e, 0x5408abf7, 0x50c9b640,
9769
           0x4e8ee645, 0x4a4ffbf2, 0x470cdd2b, 0x43cdc09c, 0x7b827d21,
9770
9771
           0x7f436096, 0x7200464f, 0x76c15bf8, 0x68860bfd, 0x6c47164a,
           0x61043093, 0x65c52d24, 0x119b4be9, 0x155a565e, 0x18197087,
9772
```

Utilities **cksum**

```
9773
           0x1cd86d30, 0x029f3d35, 0x065e2082, 0x0b1d065b, 0x0fdc1bec,
9774
           0x3793a651, 0x3352bbe6, 0x3e119d3f, 0x3ad08088, 0x2497d08d,
           0x2056cd3a, 0x2d15ebe3, 0x29d4f654, 0xc5a92679, 0xc1683bce,
9775
           0xcc2b1d17, 0xc8ea00a0, 0xd6ad50a5, 0xd26c4d12, 0xdf2f6bcb,
9776
9777
           0xdbee767c, 0xe3a1cbc1, 0xe760d676, 0xea23f0af, 0xeee2ed18,
           0xf0a5bd1d, 0xf464a0aa, 0xf9278673, 0xfde69bc4, 0x89b8fd09,
9778
           0x8d79e0be, 0x803ac667, 0x84fbdbd0, 0x9abc8bd5, 0x9e7d9662,
9779
           0x933eb0bb, 0x97ffad0c, 0xafb010b1, 0xab710d06, 0xa6322bdf,
9780
           0xa2f33668, 0xbcb4666d, 0xb8757bda, 0xb5365d03, 0xb1f740b4
9781
9782
           unsigned long memcrc(const unsigned char *b, size t n)
9783
9784
           /*
9785
                Input arguments:
9786
                const char*
                               b == byte sequence to checksum
             *
                               n == length of sequence
9787
                size t
9788
                register unsigned
9789
                                     i, c, s = 0;
9790
                for (i = n; i > 0; --i) {
                    c = (unsigned)(*b++);
9791
                    s = (s << 8) ^ crctab[(s >> 24) ^ c];
9792
9793
                /* Extend with the length of the string. */
9794
9795
                while (n != 0) {
                    c = n \& 0377;
9796
                    n >>= 8;
9797
                    s = (s << 8) ^ crctab[(s >> 24) ^ c];
9798
9799
9800
                return ~s;
9801
```

9802

9803

9804

9805 9806

9807

9808

9809

9810

9811

9812

9813 9814

9815

9816

9817 9818

9819

9820

The historical practice of writing the number of "blocks" has been changed to writing the number of octets, since the latter is not only more useful, but also since historical implementations have not been consistent in defining what a "block" meant. Octets are used instead of bytes because bytes can differ in size between systems.

The algorithm used was selected to increase the operational robustness of *cksum*. Neither the System V nor BSD *sum* algorithm was selected. Since each of these was different and each was the default behavior on those systems, no realistic compromise was available if either were selected—some set of historical applications would break. Therefore, the name was changed to *cksum*. Although the historical *sum* commands will probably continue to be provided for many years, programs designed for portability across systems should use the new name.

The algorithm selected is based on that used by the ISO/IEC 8802-3: 1996 standard (Ethernet) for the frame check sequence field. The algorithm used does not match the technical definition of a *checksum*; the term is used for historical reasons. The length of the file is included in the CRC calculation because this parallels inclusion of a length field by Ethernet in its CRC, but also because it guards against inadvertent collisions between files that begin with different series of zero octets. The chance that two different files produce identical CRCs is much greater when their lengths are not considered. Keeping the length and the checksum of the file itself separate would yield a slightly more robust algorithm, but historical usage has always been that a single number (the checksum as printed) represents the signature of the file. It was decided that

cksum Utilities

historical usage was the more important consideration.

Early proposals contained modifications to the Ethernet algorithm that involved extracting table values whenever an intermediate result became zero. This was demonstrated to be less robust than the current method and mathematically difficult to describe or justify.

The calculation used is identical to that given in pseudo-code in the referenced Sarwate article. The pseudo-code rendition is:

```
9827
             X < -0; Y < -0;
             for i <- m -1 step -1 until 0 do
9828
9829
                  begin
                  T <- X(1) ^ A[i];
9830
                  X(1) \leftarrow X(0); X(0) \leftarrow Y(1); Y(1) \leftarrow Y(0); Y(0) \leftarrow 0;
9831
                  comment: f[T] and f'[T] denote the T-th words in the
9832
                        table f and f';
9833
                  X \leftarrow X ^f[T]; Y \leftarrow Y ^f'[T];
9834
9835
                  end
```

The pseudo-code is reproduced exactly as given; however, note that in the case of *cksum*, **A[i]** represents a byte of the file, the words **X** and **Y** are treated as a single 32-bit value, and the tables **f** and **f**' are a single table containing 32-bit values.

The referenced Sarwate article also discusses generating the table.

9840 FUTURE DIRECTIONS

9841 None.

9842 SEE ALSO

9822

9823

9824 9825

9826

9836

9837

9838 9839

9845

9843 None.

9844 CHANGE HISTORY

First released in Issue 4.

Utilities cmp

9846 **NAME** cmp — compare two files 9847 9848 **SYNOPSIS** cmp [-l | -s] file1 file2 9849 DESCRIPTION 9850 The *cmp* utility shall compare two files. The *cmp* utility shall write no output if the files are the 9851 same. Under default options, if they differ, it shall write to standard output the byte and line 9852 number at which the first difference occurred. Bytes and lines shall be numbered beginning with 9853 9854 **OPTIONS** 9855 The *cmp* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 9856 12.2, Utility Syntax Guidelines. 9857 The following options shall be supported: 9858 -1(Lowercase ell.) Write the byte number (decimal) and the differing bytes (octal) for 9859 9860 each difference. Write nothing for differing files; return exit status only. 9861 -5 **OPERANDS** 9862 The following operands shall be supported: 9863 file1 A pathname of the first file to be compared. If *file1* is '-', the standard input shall 9864 be used. 9865 file2 A pathname of the second file to be compared. If file2 is '-', the standard input 9866 shall be used. 9867 If both file1 and file2 refer to standard input or refer to the same FIFO special, block special, or 9868 character special file, the results are undefined. 9869 **STDIN** 9870 The standard input shall be used only if the *file1* or *file2* operand refers to standard input. See the 9871 9872 INPUT FILES section. **INPUT FILES** 9873 The input files can be any file type. 9874 **ENVIRONMENT VARIABLES** 9875 9876 The following environment variables shall affect the execution of *cmp*: LANG Provide a default value for the internationalization variables that are unset or null. 9877 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 9878 Internationalization Variables for the precedence of internationalization variables 9879 used to determine the values of locale categories.) 9880 LC ALL If set to a non-empty string value, override the values of all the other 9881 internationalization variables. 9882 9883 LC_CTYPE 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). 9885 LC_MESSAGES 9886

standard output.

9887

9888 9889 Determine the locale that should be used to affect the format and contents of

diagnostic messages written to standard error and informative messages written to

cmp Utilities

9890 XSI NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

9891 ASYNCHRONOUS EVENTS

Default.

9893 STDOUT

9892

9894

9895

9901

9902

9903

9905

9907

9909

9910

9916

9917

9918

9919

9920

9924

9925 9926

9927

9928

9929

9930

In the POSIX locale, results of the comparison shall be written to standard output. When no options are used, the format shall be:

```
9896 "%s %s differ: char %d, line %d\n", file1, file2, 
9897 <br/>
obyte number>, <line number>
```

9898 When the **-l** option is used, the format shall be:

```
9899 "%d %o %o\n", <byte number>, <differing byte>,
9900 <differing byte>
```

for each byte that differs. The first *<differing byte>* number is from *file1* while the second is from *file2*. In both cases, *<byte number>* shall be relative to the beginning of the file, beginning with 1.

No output shall be written to standard output when the $-\mathbf{s}$ option is used.

9904 STDERR

The standard error shall be used only for diagnostic messages. If *file1* and *file2* are identical for the entire length of the shorter file, in the POSIX locale the following diagnostic message shall be written, unless the -s option is specified:

```
9908 "cmp: EOF on %s%s\n", <name of shorter file>, <additional info>
```

The *<additional info>* field shall either be null or a string that starts with a *<*blank*>* and contains no *<*newline*>*s. Some implementations report on the number of lines in this case.

9911 OUTPUT FILES

9912 None.

9913 EXTENDED DESCRIPTION

9914 None.

9915 EXIT STATUS

The following exit values shall be returned:

- 0 The files are identical.
- 1 The files are different; this includes the case where one file is identical to the first part of the other.
- >1 An error occurred.

9921 CONSEQUENCES OF ERRORS

9922 Default.

9923 APPLICATION USAGE

Although input files to *cmp* can be any type, the results might not be what would be expected on character special device files or on file types not described by the System Interfaces volume of IEEE Std 1003.1-2001. Since this volume of IEEE Std 1003.1-2001 does not specify the block size used when doing input, comparisons of character special files need not compare all of the data in those files.

For files which are not text files, line numbers simply reflect the presence of a <newline>, without any implication that the file is organized into lines.

Utilities cmp

9931 **EXAMPLES** 9932 None. 9933 **RATIONALE** The global language in Section 1.11 (on page 20) indicates that using two mutually-exclusive 9934 options together produces unspecified results. Some System V implementations consider the 9935 9936 option usage: cmp -1 -s ... 9937 to be an error. They also treat: 9938 9939 $cmp -s -l \dots$ as if no options were specified. Both of these behaviors are considered bugs, but are allowed. 9940 9941 The word char in the standard output format comes from historical usage, even though it is actually a byte number. When cmp is supported in other locales, implementations are 9942 encouraged to use the word byte or its equivalent in another language. Users should not 9943 interpret this difference to indicate that the functionality of the utility changed between locales. 9944 9945 Some implementations report on the number of lines in the identical-but-shorter file case. This is 9946 allowed by the inclusion of the *<additional info>* fields in the output format. The restriction on having a leading
blank> and no <newline>s is to make parsing for the filename easier. It is 9947 recognized that some filenames containing white-space characters make parsing difficult 9948 anyway, but the restriction does aid programs used on systems where the names are 9949 9950 predominantly well behaved. **FUTURE DIRECTIONS** 9951 9952 None. **SEE ALSO** 9953 9954 comm, diff **CHANGE HISTORY** 9955 9956 First released in Issue 2.

comm Utilities

9957 NAME comm — select or reject lines common to two files 9958 9959 **SYNOPSIS** comm [-123] file1 file2 9960 9961 DESCRIPTION The *comm* utility shall read *file1* and *file2*, which should be ordered in the current collating 9962 sequence, and produce three text columns as output: lines only in file1, lines only in file2, and 9963 lines in both files. 9964 If the lines in both files are not ordered according to the collating sequence of the current locale, 9965 the results are unspecified. 9966 **OPTIONS** 9967 The *comm* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 9968 12.2, Utility Syntax Guidelines. 9969 9970 The following options shall be supported: -1Suppress the output column of lines unique to *file1*. 9971 **-2** 9972 Suppress the output column of lines unique to *file2*. -3Suppress the output column of lines duplicated in *file1* and *file2*. 9973 **OPERANDS** 9974 9975 The following operands shall be supported: file1 A pathname of the first file to be compared. If *file1* is '-', the standard input shall 9976 be used. 9977 file2 A pathname of the second file to be compared. If *file2* is '-', the standard input 9978 shall be used. 9979 If both file1 and file2 refer to standard input or to the same FIFO special, block special, or 9980 9981 character special file, the results are undefined. **STDIN** 9982 The standard input shall be used only if one of the file1 or file2 operands refers to standard input. 9983 See the INPUT FILES section. 9984 **INPUT FILES** 9985 The input files shall be text files. 9986 9987 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *comm*: 9988 LANG Provide a default value for the internationalization variables that are unset or null. 9989 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 9990 Internationalization Variables for the precedence of internationalization variables 9991 used to determine the values of locale categories.) 9992 LC ALL If set to a non-empty string value, override the values of all the other 9993 internationalization variables. 9994 LC_COLLATE 9995 Determine the locale for the collating sequence *comm* expects to have been used 9996 when the input files were sorted. 9997 Determine the locale for the interpretation of sequences of bytes of text data as LC_CTYPE 9998 9999 characters (for example, single-byte as opposed to multi-byte characters in

Utilities comm

10000 arguments and input files). LC_MESSAGES 10001 10002 Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. 10003 10004 XSI **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 10005 ASYNCHRONOUS EVENTS 10006 Default. 10007 STDOUT 10008 The *comm* utility shall produce output depending on the options selected. If the -1, -2, and -3options are all selected, *comm* shall write nothing to standard output. 10009 If the -1 option is not selected, lines contained only in *file1* shall be written using the format: 10010 "%s\n", <line in file1> 10011 If the **-2** option is not selected, lines contained only in *file2* are written using the format: 10012 "%s%sn", <lead>, <line in file2> 10013 where the string < *lead*> is as follows: 10014 <tab> The -1 option is not selected. 10015 null string 10016 The -1 option is selected. If the -3 option is not selected, lines contained in both files shall be written using the format: 10017 "%s%s\n", <lead>, <line in both> 10018 where the string < *lead*> is as follows: 10019 10020 <tab><tab> Neither the -1 nor the -2 option is selected. <tab> Exactly one of the -1 and -2 options is selected. 10021 null string Both the -1 and -2 options are selected. 10022 If the input files were ordered according to the collating sequence of the current locale, the lines 10023 written shall be in the collating sequence of the original lines. 10024 10025 STDERR The standard error shall be used only for diagnostic messages. 10026 10027 OUTPUT FILES None. 10028 10029 EXTENDED DESCRIPTION None. 10030 10031 EXIT STATUS The following exit values shall be returned: 10032 All input files were successfully output as specified. 10033 >0 An error occurred. 10034 10035 CONSEQUENCES OF ERRORS

Default.

comm Utilities

10037 APPLICATION USAGE 10038 If the input files are not properly presorted, the output of *comm* might not be useful. 10039 EXAMPLES If a file named xcu contains a sorted list of the utilities in this volume of IEEE Std 1003.1-2001, a 10040 file named xpg3 contains a sorted list of the utilities specified in the X/Open Portability Guide, 10041 Issue 3, and a file named svid89 contains a sorted list of the utilities in the System V Interface 10042 **Definition Third Edition:** 10043 comm -23 xcu xpq3 | comm -23 - svid89 10044 would print a list of utilities in this volume of IEEE Std 1003.1-2001 not specified by either of the 10045 other documents: 10046 10047 comm -12 xcu xpg3 | comm -12 - svid89 would print a list of utilities specified by all three documents, and: 10048 10049 comm -12 xpg3 svid89 | comm -23 - xcu would print a list of utilities specified by both XPG3 and the SVID, but not specified in this 10050 volume of IEEE Std 1003.1-2001. 10051 10052 RATIONALE 10053 None. 10054 FUTURE DIRECTIONS 10055 None. 10056 SEE ALSO 10057 cmp, diff, sort, uniq 10058 CHANGE HISTORY First released in Issue 2. 10059 10060 Issue 6

The normative text is reworded to avoid use of the term "must" for application requirements.

command **Utilities**

10062 **NAME** 10063 command — execute a simple command 10064 SYNOPSIS 10065 command [-p] command name [argument ...] 10066 UP command [-v | -V] command name 10067

10068 DESCRIPTION

The *command* utility shall cause the shell to treat the arguments as a simple command, suppressing the shell function lookup that is described in Section 2.9.1.1 (on page 48), item 1b.

If the command_name is the same as the name of one of the special built-in utilities, the special 10071 properties in the enumerated list at the beginning of Section 2.14 (on page 64) shall not occur. In 10072 every other respect, if command_name is not the name of a function, the effect of command (with 10073 no options) shall be the same as omitting *command*. 10074

> On systems supporting the User Portability Utilities option, the *command* utility also shall provide information concerning how a command name is interpreted by the shell; see -v and -V.

10078 OPTIONS

The command utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported:

Perform the command search using a default value for PATH that is guaranteed to 10082 -p find all of the standard utilities. 10083

> (On systems supporting the User Portability Utilities option.) Write a string to standard output that indicates the pathname or command that will be used by the shell, in the current shell execution environment (see Section 2.12 (on page 61)), to invoke command name, but do not invoke command name.

- Utilities, regular built-in utilities, command_names including a slash character, and any implementation-defined functions that are found using the PATH variable (as described in Section 2.9.1.1 (on page 48)), shall be written as absolute pathnames.
- Shell functions, special built-in utilities, regular built-in utilities not associated with a PATH search, and shell reserved words shall be written as just their
- An alias shall be written as a command line that represents its alias definition.
- Otherwise, no output shall be written and the exit status shall reflect that the name was not found.

(On systems supporting the User Portability Utilities option.) Write a string to standard output that indicates how the name given in the *command_name* operand will be interpreted by the shell, in the current shell execution environment (see Section 2.12 (on page 61)), but do not invoke *command_name*. Although the format of this string is unspecified, it shall indicate in which of the following categories *command_name* falls and shall include the information stated:

• Utilities, regular built-in utilities, and any implementation-defined functions that are found using the PATH variable (as described in Section 2.9.1.1 (on page 48)), shall be identified as such and include the absolute pathname in the string.

10069

10070

10075

10076

10077

10079

10081

10084 $-\mathbf{v}$

10085 10086 10087

10092 10093 10094

10095 10096

10097 10098 10099

10100

10101

10102

10103

 $-\mathbf{V}$

command Utilities

10107		 Other shell functions shall be identified as functions.
10108		 Aliases shall be identified as aliases and their definitions included in the string.
10109		 Special built-in utilities shall be identified as special built-in utilities.
10110 10111		• Regular built-in utilities not associated with a <i>PATH</i> search shall be identified as regular built-in utilities. (The term "regular" need not be used.)
10112		 Shell reserved words shall be identified as reserved words.
10113 OPERA	NDS	
10114	The followin	g operands shall be supported:
10115	argument	One of the strings treated as an argument to <i>command_name</i> .
10116 10117	command_na	<i>me</i> The name of a utility or a special built-in utility.
10118 STDIN	NI. a	
10119	Not used.	
10120 INPUT 10121	FILES None.	
10122 ENVIR (ARIARLES
10123		g environment variables shall affect the execution of <i>command</i> :
10124 10125 10126 10127	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
10128 10129	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
10130 10131 10132	LC_CTYPE	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).
10133	LC_MESSAC	
10134 10135 10136		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.
10137 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .
10138 10139	PATH	Determine the search path used during the command search described in Section 2.9.1.1 (on page 48), except as described under the $-\mathbf{p}$ option.
10140 ASYNC	HRONOUS I	EVENTS
10141	Default.	
10142 STDOU 10143		option is specified, standard output shall be formatted as:
10144	"%s\n", <	pathname or command>
10145	When the -V	option is specified, standard output shall be formatted as:
10146	"%s\n", <	unspecified>

Utilities command

10147 STDERR

10148 The standard error shall be used only for diagnostic messages.

10149 OUTPUT FILES

None. 10150

10151 EXTENDED DESCRIPTION

None. 10152

10153 EXIT STATUS

When the $-\mathbf{v}$ or $-\mathbf{V}$ options are specified, the following exit values shall be returned: 10154

- 10155 Successful completion.
- >0 The *command_name* could not be found or an error occurred. 10156
- 10157 Otherwise, the following exit values shall be returned:
- 126 The utility specified by *command_name* was found but could not be invoked. 10158
- 127 An error occurred in the *command* utility or the utility specified by *command_name* could not 10159 be found. 10160
- Otherwise, the exit status of command shall be that of the simple command specified by the 10161 10162 arguments to *command*.

10163 CONSEQUENCES OF ERRORS

Default. 10164

10177

10178

10179

10180

10181

10182

10183

10184

10185

10165 APPLICATION USAGE

The order for command search allows functions to override regular built-ins and path searches. 10166 This utility is necessary to allow functions that have the same name as a utility to call the utility 10167 (instead of a recursive call to the function). 10168

The system default path is available using *getconf*; however, since *getconf* may need to have the 10169 *PATH* set up before it can be called itself, the following can be used: 10170

10171 command -p getconf CS PATH

There are some advantages to suppressing the special characteristics of special built-ins on 10172 occasion. For example: 10173

10174 command exec > unwritable-file

does not cause a non-interactive script to abort, so that the output status can be checked by the 10175 10176

The command, env, nohup, time, 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. 10186

10187 Since the -v and -V options of *command* produce output in relation to the current shell execution 10188 environment, command is generally provided as a shell regular built-in. If it is called in a subshell or separate utility execution environment, such as one of the following: 10189

command Utilities

```
10190 (PATH=foo command -v)
10191 nohup command -v
```

it does not necessarily produce correct results. For example, when called with *nohup* or an *exec* function, in a separate utility execution environment, most implementations are not able to identify aliases, functions, or special built-ins.

Two types of regular built-ins could be encountered on a system and these are described separately by *command*. The description of command search in Section 2.9.1.1 (on page 48) allows for a standard utility to be implemented as a regular built-in as long as it is found in the appropriate place in a *PATH* search. So, for example, *command* –v *true* might yield /bin/true or some similar pathname. Other implementation-defined utilities that are not defined by this volume of IEEE Std 1003.1-2001 might exist only as built-ins and have no pathname associated with them. These produce output identified as (regular) built-ins. Applications encountering these are not able to count on *exec*ing them, using them with *nohup*, overriding them with a different *PATH*, and so on.

10204 EXAMPLES

1. Make a version of *cd* that always prints out the new working directory exactly once:

```
cd() {
    command cd "$@" >/dev/null
    pwd
}
```

2. Start off a "secure shell script" in which the script avoids being spoofed by its parent:

```
IFS='
10211
10212
10213
                      The preceding value should be <space><tab><newline>.
10214
                #
                      Set IFS to its default value.
10215
                \unalias -a
10216
                      Unset all possible aliases.
                #
                      Note that unalias is escaped to prevent an alias
10217
10218
                      being used for unalias.
                unset -f command
10219
                      Ensure command is not a user function.
10220
                PATH="$(command -p getconf CS PATH): $PATH"
10221
10222
                      Put on a reliable PATH prefix.
10223
```

At this point, given correct permissions on the directories called by *PATH*, the script has the ability to ensure that any utility it calls is the intended one. It is being very cautious because it assumes that implementation extensions may be present that would allow user functions to exist when it is invoked; this capability is not specified by this volume of IEEE Std 1003.1-2001, but it is not prohibited as an extension. For example, the *ENV* variable precedes the invocation of the script with a user start-up script. Such a script could define functions to spoof the application.

10231 RATIONALE

Since *command* is a regular built-in utility it is always found prior to the *PATH* search.

There is nothing in the description of *command* that implies the command line is parsed any differently from that of any other simple command. For example:

Utilities command

10235 command a | b ; c

is not parsed in any special way that causes $' \mid '$ or '; ' to be treated other than a pipe operator or semicolon or that prevents function lookup on $\bf b$ or $\bf c$.

The *command* utility is somewhat similar to the Eighth Edition shell *builtin* command, but since *command* also goes to the file system to search for utilities, the name *builtin* would not be intuitive.

The *command* utility is most likely to be provided as a regular built-in. It is not listed as a special built-in for the following reasons:

- The removal of exportable functions made the special precedence of a special built-in unnecessary.
- A special built-in has special properties (see Section 2.14 (on page 64)) that were inappropriate for invoking other utilities. For example, two commands such as:

```
10247 date > unwritable-file
```

command date > unwritable-file

would have entirely different results; in a non-interactive script, the former would continue to execute the next command, the latter would abort. Introducing this semantic difference along with suppressing functions was seen to be non-intuitive.

The –**p** option is present because it is useful to be able to ensure a safe path search that finds all the standard utilities. This search might not be identical to the one that occurs through one of the *exec* functions (as defined in the System Interfaces volume of IEEE Std 1003.1-2001) when *PATH* is unset. At the very least, this feature is required to allow the script to access the correct version of *getconf* so that the value of the default path can be accurately retrieved.

The *command* –**v** and –**v** options were added to satisfy requirements from users that are currently accomplished by three different historical utilities: *type* in the System V shell, *whence* in the KornShell, and *which* in the C shell. Since there is no historical agreement on how and what to accomplish here, the POSIX *command* utility was enhanced and the historical utilities were left unmodified. The C shell *which* merely conducts a path search. The KornShell *whence* is more elaborate—in addition to the categories required by POSIX, it also reports on tracked aliases, exported aliases, and undefined functions.

The output format of -V was left mostly unspecified because human users are its only audience. Applications should not be written to care about this information; they can use the output of -v to differentiate between various types of commands, but the additional information that may be emitted by the more verbose -V is not needed and should not be arbitrarily constrained in its verbosity or localization for application parsing reasons.

10269 FUTURE DIRECTIONS

10270 None.

SEE ALSO

Section 2.9.1.1 (on page 48), Section 2.12 (on page 61), Section 2.14 (on page 64), *sh*, *type*, the System Interfaces volume of IEEE Std 1003.1-2001, *exec*

10274 CHANGE HISTORY

First released in Issue 4.

compress Utilities

10276 NAME 10277	compress –	– compress data		
10278 SYNOPSIS				
10279 XSI		[-fv] [-b bits] [file]		
10280 10281	compress	[-cfv] [-b bits] [file]		
10282 DESCR	IPTION			
10283 10284		ess utility shall attempt to reduce the size of the named files by using adaptive v coding algorithm.		
10285 10286		Lempel-Ziv is US Patent 4464650, issued to William Eastman, Abraham Lempel, Jacob Ziv, Martin Cohn on August 7th, 1984, and assigned to Sperry Corporation.		
10287 10288		Lempel-Ziv-Welch compression is covered by US Patent 4558302, issued to Terry A. Welch on December 10th, 1985, and assigned to Sperry Corporation.		
10289 10290 10291 10292 10293 10294 10295	changed an standard or has approp original file {NAME_M	s not supporting adaptive Lempel-Ziv coding algorithm, the input files shall not be an error value greater than two shall be returned. Except when the output is to the utput, each file shall be replaced by one with the extension .Z. If the invoking process priate privileges, the ownership, modes, access time, and modification time of the error preserved. If appending the .Z to the filename would make the name exceed AX} bytes, the command shall fail. If no files are specified, the standard input shall be d to the standard output.		
10296 OPTIO	NS			
10297 10298	The $\it compress$ utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.			
10299	The followi	The following options shall be supported:		
10300 10301	− b bits	Specify the maximum number of bits to use in a code. For a conforming application, the <i>bits</i> argument shall be:		
10302		9 <= bits <= 14		
10303 10304		The implementation may allow <i>bits</i> values of greater than 14. The default is 14, 15, or 16.		
10305 10306	- c	Cause <i>compress</i> to write to the standard output; the input file is not changed, and no . Z files are created.		
10307 10308 10309 10310	-f	Force compression of <i>file</i> , even if it does not actually reduce the size of the file, or if the corresponding <i>file</i> . Z file already exists. If the – f option is not given, and the process is not running in the background, the user is prompted as to whether an existing <i>file</i> . Z file should be overwritten.		
10311	- v	Write the percentage reduction of each file to standard error.		
10312 OPERA	NDS			
10313	The following	ing operand shall be supported:		
10314	file	A pathname of a file to be compressed.		

10315 **STDIN**

10316

The standard input shall be used only if no *file* operands are specified, or if a *file* operand is '-'.

Utilities compress

10317 INPUT FILES

10318 If *file* operands are specified, the input files contain the data to be compressed.

10319 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *compress*: 10320

10321 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 10322 Internationalization Variables for the precedence of internationalization variables 10323 used to determine the values of locale categories.) 10324

LC_ALL 10325 If set to a non-empty string value, override the values of all the other 10326

internationalization variables.

Determine the locale for the interpretation of sequences of bytes of text data as 10327 LC_CTYPE characters (for example, single-byte as opposed to multi-byte characters in 10328 10329

arguments).

LC_MESSAGES 10330

Determine the locale that should be used to affect the format and contents of 10331 10332 diagnostic messages written to standard error.

NLSPATH Determine the location of message catalogs for the processing of *LC MESSAGES*. 10333

10334 ASYNCHRONOUS EVENTS

Default. 10335

10336 STDOUT

If no file operands are specified, or if a file operand is '-', or if the -c option is specified, the 10337 10338 standard output contains the compressed output.

10339 STDERR

10340 The standard error shall be used only for diagnostic and prompt messages and the output from 10341

10342 OUTPUT FILES

The output files shall contain the compressed output. The format of compressed files is 10343 unspecified and interchange of such files between implementations (including access via 10344 unspecified file sharing mechanisms) is not required by IEEE Std 1003.1-2001. 10345

10346 EXTENDED DESCRIPTION

None. 10347

10348 EXIT STATUS

The following exit values shall be returned: 10349

Successful completion. 10350

An error occurred. 10351

One or more files were not compressed because they would have increased in size (and the 10352 -**f** option was not specified). 10353

10354 >2 An error occurred.

10355 CONSEQUENCES OF ERRORS

10356 The input file shall remain unmodified. **compress**Utilities

10357 APPLICATION USAGE

The amount of compression obtained depends on the size of the input, the number of *bits* per code, and the distribution of common substrings. Typically, text such as source code or English is reduced by 50-60%. Compression is generally much better than that achieved by Huffman coding or adaptive Huffman coding (*compact*), and takes less time to compute.

Although *compress* strictly follows the default actions upon receipt of a signal or when an error occurs, some unexpected results may occur. In some implementations it is likely that a partially compressed file is left in place, alongside its uncompressed input file. Since the general operation of *compress* is to delete the uncompressed file only after the .Z file has been successfully filled, an application should always carefully check the exit status of *compress* before arbitrarily deleting files that have like-named neighbors with .Z suffixes.

The limit of 14 on the *bits* option-argument is to achieve portability to all systems (within the restrictions imposed by the lack of an explicit published file format). Some implementations based on 16-bit architectures cannot support 15 or 16-bit uncompression.

10371 EXAMPLES

10368 10369

10370

10372 None.

10373 RATIONALE

10374 None.

10375 FUTURE DIRECTIONS

10376 None.

10377 SEE ALSO

10378 uncompress, zcat

10379 CHANGE HISTORY

10380 First released in Issue 4.

10381 Issue 6

The normative text is reworded to avoid use of the term "must" for application requirements.

An error case is added for systems not supporting adaptive Lempel-Ziv coding.

Utilities cp

10391 DESCRIPTION

The first synopsis form is denoted by two operands, neither of which are existing files of type directory. The *cp* utility shall copy the contents of *source_file* (or, if *source_file* is a file of type symbolic link, the contents of the file referenced by *source_file*) to the destination path named by *target_file*.

The second synopsis form is denoted by two or more operands where the $-\mathbf{R}$ or $-\mathbf{r}$ options are not specified and the first synopsis form is not applicable. It shall be an error if any *source_file* is a file of type directory, if *target* does not exist, or if *target* is a file of a type defined by the System Interfaces volume of IEEE Std 1003.1-2001, but is not a file of type directory. The *cp* utility shall copy the contents of each *source_file* (or, if *source_file* is a file of type symbolic link, the contents of the file referenced by *source_file*) to the destination path named by the concatenation of *target*, a slash character, and the last component of *source_file*.

The third and fourth synopsis forms are denoted by two or more operands where the $-\mathbf{R}$ or $-\mathbf{r}$ options are specified. The *cp* utility shall copy each file in the file hierarchy rooted in each *source_file* to a destination path named as follows:

- If *target* exists and is a file of type directory, the name of the corresponding destination path for each file in the file hierarchy shall be the concatenation of *target*, a slash character, and the pathname of the file relative to the directory containing *source_file*.
- If *target* does not exist and two operands are specified, the name of the corresponding destination path for *source_file* shall be *target*; the name of the corresponding destination path for all other files in the file hierarchy shall be the concatenation of *target*, a slash character, and the pathname of the file relative to *source_file*.

It shall be an error if *target* does not exist and more than two operands are specified, or if *target* exists and is a file of a type defined by the System Interfaces volume of IEEE Std 1003.1-2001, but is not a file of type directory.

In the following description, the term *dest_file* refers to the file named by the destination path. The term *source_file* refers to the file that is being copied, whether specified as an operand or a file in a file hierarchy rooted in a *source_file* operand. If *source_file* is a file of type symbolic link:

- If neither the $-\mathbf{R}$ nor $-\mathbf{r}$ options were specified, cp shall take actions based on the type and contents of the file referenced by the symbolic link, and not by the symbolic link itself.
- If the –R option was specified:
- If none of the options –H, –L, nor –P were specified, it is unspecified which of –H, –L, or –P will be used as a default.
- If the –H option was specified, *cp* shall take actions based on the type and contents of the file referenced by any symbolic link specified as a *source_file* operand.
- If the -L option was specified, cp shall take actions based on the type and contents of the file referenced by any symbolic link specified as a source_file operand or any symbolic

CP Utilities

links encountered during traversal of a file hierarchy.

 If the -P option was specified, cp shall copy any symbolic link specified as a source_file
operand and any symbolic links encountered during traversal of a file hierarchy, and shall
not follow any symbolic links.

• If the $-\mathbf{r}$ option was specified, the behavior is implementation-defined.

For each *source_file*, the following steps shall be taken:

 1. If *source_file* references the same file as *dest_file*, *cp* may write a diagnostic message to standard error; it shall do nothing more with *source_file* and shall go on to any remaining files.

2. If *source_file* is of type directory, the following steps shall be taken:

 a. If neither the $-\mathbf{R}$ or $-\mathbf{r}$ options were specified, cp shall write a diagnostic message to standard error, do nothing more with $source_file$, and go on to any remaining files.

 If source_file was not specified as an operand and source_file is dot or dot-dot, cp shall do nothing more with source_file and go on to any remaining files.

 c. If *dest_file* exists and it is a file type not specified by the System Interfaces volume of IEEE Std 1003.1-2001, the behavior is implementation-defined.

 d. If *dest_file* exists and it is not of type directory, *cp* shall write a diagnostic message to standard error, do nothing more with *source_file* or any files below *source_file* in the file hierarchy, and go on to any remaining files.

e. If the directory *dest_file* does not exist, it shall be created with file permission bits set to the same value as those of *source_file*, modified by the file creation mask of the user if the -**p** option was not specified, and then bitwise-inclusively OR'ed with S_IRWXU. If *dest_file* cannot be created, *cp* shall write a diagnostic message to standard error, do nothing more with *source_file*, and go on to any remaining files. It is unspecified if *cp* attempts to copy files in the file hierarchy rooted in *source_file*.

f. The files in the directory *source_file* shall be copied to the directory *dest_file*, taking the four steps (1 to 4) listed here with the files as *source_files*.

g. If *dest_file* was created, its file permission bits shall be changed (if necessary) to be the same as those of *source_file*, modified by the file creation mask of the user if the **-p** option was not specified.

h. The cp utility shall do nothing more with source_file and go on to any remaining files.

3. If *source_file* is of type regular file, the following steps shall be taken:

 a. If *dest_file* exists, the following steps shall be taken:

 i. If the —i option is in effect, the *cp* utility shall write a prompt to the standard error and read a line from the standard input. If the response is not affirmative, *cp* shall do nothing more with *source_file* and go on to any remaining files.

ii. A file descriptor for dest_file shall be obtained by performing actions equivalent to the open() function defined in the System Interfaces volume of IEEE Std 1003.1-2001 called using dest_file as the path argument, and the bitwise-inclusive OR of O_WRONLY and O_TRUNC as the oflag argument.

 iii. If the attempt to obtain a file descriptor fails and the **–f** option is in effect, *cp* shall attempt to remove the file by performing actions equivalent to the *unlink()* function defined in the System Interfaces volume of

Utilities cp

10471 IEEE Std 1003.1-2001 called using dest_file as the path argument. If this attempt 10472 succeeds, *cp* shall continue with step 3b. 10473 b. If dest file does not exist, a file descriptor shall be obtained by performing actions equivalent to the open() function defined in the System Interfaces volume of 10474 IEEE Std 1003.1-2001 called using dest_file as the path argument, and the bitwise-10475 inclusive OR of O_WRONLY and O_CREAT as the oflag argument. The file 10476 permission bits of *source_file* shall be the *mode* argument. 10477 c. If the attempt to obtain a file descriptor fails, cp shall write a diagnostic message to 10478 10479 standard error, do nothing more with *source_file*, and go on to any remaining files. d. The contents of *source_file* shall be written to the file descriptor. Any write errors 10480 shall cause *cp* to write a diagnostic message to standard error and continue to step 3e. 10481 The file descriptor shall be closed. 10482 The *cp* utility shall do nothing more with *source_file*. If a write error occurred in step 10483 3d, it is unspecified if cp continues with any remaining files. If no write error 10484 occurred in step 3d, *cp* shall go on to any remaining files. 10485 4. Otherwise, the following steps shall be taken: 10486 a. If the **-r** option was specified, the behavior is implementation-defined. 10487 10488 b. If the $-\mathbf{R}$ option was specified, the following steps shall be taken: 10489 i. The *dest_file* shall be created with the same file type as *source_file*. If source_file is a file of type FIFO, the file permission bits shall be the same as 10490 those of *source_file*, modified by the file creation mask of the user if the -p 10491 10492 option was not specified. Otherwise, the permissions, owner ID, and group ID of *dest_file* are implementation-defined. 10493 If this creation fails for any reason, cp shall write a diagnostic message to 10494 10495 standard error, do nothing more with *source_file*, and go on to any remaining files. 10496 If source_file is a file of type symbolic link, the pathname contained in dest_file 10497 shall be the same as the pathname contained in *source_file*. 10498 If this fails for any reason, cp shall write a diagnostic message to standard error, 10499 do nothing more with *source_file*, and go on to any remaining files. 10500 If the implementation provides additional or alternate access control mechanisms (see the Base 10501 Definitions volume of IEEE Std 1003.1-2001, Section 4.4, File Access Permissions), their effect on 10502 copies of files is implementation-defined. 10503 10504 OPTIONS The *cp* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, 10505 Utility Syntax Guidelines. 10506 The following options shall be supported: 10507 $-\mathbf{f}$ If a file descriptor for a destination file cannot be obtained, as described in step 10508 3.a.ii., attempt to unlink the destination file and proceed. 10509 -H10510 Take actions based on the type and contents of the file referenced by any symbolic link specified as a *source_file* operand. 10511 10512 $-\mathbf{i}$ Write a prompt to standard error before copying to any existing destination file. If 10513 the response from the standard input is affirmative, the copy shall be attempted;

cp Utilities

		all and the first terms of the second
10514		otherwise, it shall not.
10515	–L	Take actions based on the type and contents of the file referenced by any symbolic
10516		link specified as a <i>source_file</i> operand or any symbolic links encountered during traversal of a file hierarchy.
10517		v
10518	−P	Take actions on any symbolic link specified as a <i>source_file</i> operand or any
10519		symbolic link encountered during traversal of a file hierarchy.
10520 10521	-p	Duplicate the following characteristics of each source file in the corresponding destination file:
10522 10523		1. The time of last data modification and time of last access. If this duplication fails for any reason, <i>cp</i> shall write a diagnostic message to standard error.
10524 10525		2. The user ID and group ID. If this duplication fails for any reason, it is unspecified whether <i>cp</i> writes a diagnostic message to standard error.
10526 10527 10528		3. The file permission bits and the S_ISUID and S_ISGID bits. Other, implementation-defined, bits may be duplicated as well. If this duplication fails for any reason, <i>cp</i> shall write a diagnostic message to standard error.
10529 10530 10531		If the user ID or the group ID cannot be duplicated, the file permission bits S_ISUID and S_ISGID shall be cleared. If these bits are present in the source file but are not duplicated in the destination file, it is unspecified whether <i>cp</i> writes a diagnostic message to standard error.
10532		diagnostic message to standard error.
10533 10534		The order in which the preceding characteristics are duplicated is unspecified. The <i>dest_file</i> shall not be deleted if these characteristics cannot be preserved.
10535	$-\mathbf{R}$	Copy file hierarchies.
		copy the incruremes.
10536 ОВ	-r	Copy file hierarchies. The treatment of special files is implementation-defined.
	- r Specifying	• •
10536 ОВ 10537	- r Specifying considered	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be
10536 OB 10537 10538	-r Specifying considered	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be
10536 ОВ 10537 10538 10539 ОРЕК А	-r Specifying considered	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options – H , – L , and – P shall not be an error. The last option specified shall determine the behavior of the utility.
10536 OB 10537 10538 10539 OPER A 10540	-r Specifying considered ANDS The followi	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be an error. The last option specified shall determine the behavior of the utility. ng operands shall be supported:
10536 OB 10537 10538 10539 OPER A 10540 10541 10542 10543	-r Specifying considered ANDS The followi source_file target_file	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be an error. The last option specified shall determine the behavior of the utility. ng operands shall be supported: A pathname of a file to be copied. A pathname of an existing or nonexistent file, used for the output when a single file is copied.
10536 OB 10537 10538 10539 OPER A 10540 10541 10542 10543 10544	-r Specifying considered ANDS The followi source_file target_file	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be an error. The last option specified shall determine the behavior of the utility. ng operands shall be supported: A pathname of a file to be copied. A pathname of an existing or nonexistent file, used for the output when a single
10536 OB 10537 10538 10539 OPER 10540 10541 10542 10543 10544 10545 STDIN	-r Specifying considered ANDS The followi source_file target_file	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be an error. The last option specified shall determine the behavior of the utility. ng operands shall be supported: A pathname of a file to be copied. A pathname of an existing or nonexistent file, used for the output when a single file is copied. A pathname of a directory to contain the copied files.
10536 OB 10537 10538 10539 OPER A 10540 10541 10542 10543 10544	-r Specifying considered ANDS The followi source_file target_file target The standar	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be an error. The last option specified shall determine the behavior of the utility. ng operands shall be supported: A pathname of a file to be copied. A pathname of an existing or nonexistent file, used for the output when a single file is copied.
10536 OB 10537 10538 10539 OPER A 10540 10541 10542 10543 10544 10545 STDIN 10546 10547	-r Specifying considered ANDS The following source_file target_file target The standarthe STDERI	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be an error. The last option specified shall determine the behavior of the utility. ng operands shall be supported: A pathname of a file to be copied. A pathname of an existing or nonexistent file, used for the output when a single file is copied. A pathname of a directory to contain the copied files. rd input shall be used to read an input line in response to each prompt specified in
10536 OB 10537 10538 10539 OPER A 10540 10541 10542 10543 10544 10545 STDIN 10546	Specifying considered ANDS The followi source_file target_file target The standar the STDERI	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be an error. The last option specified shall determine the behavior of the utility. ng operands shall be supported: A pathname of a file to be copied. A pathname of an existing or nonexistent file, used for the output when a single file is copied. A pathname of a directory to contain the copied files. rd input shall be used to read an input line in response to each prompt specified in R section. Otherwise, the standard input shall not be used.
10536 OB 10537 10538 10539 OPER A 10540 10541 10542 10543 10544 10545 STDIN 10546 10547 10548 INPUT 10549	Specifying considered ANDS The following source_file target_file target The standare the STDERICATION The input file standare file target.	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be an error. The last option specified shall determine the behavior of the utility. ng operands shall be supported: A pathname of a file to be copied. A pathname of an existing or nonexistent file, used for the output when a single file is copied. A pathname of a directory to contain the copied files. rd input shall be used to read an input line in response to each prompt specified in R section. Otherwise, the standard input shall not be used. iles specified as operands may be of any file type.
10536 OB 10537 10538 10539 OPER A 10540 10541 10542 10543 10544 10545 STDIN 10546 10547 10548 INPUT 10549	Specifying considered ANDS The following source_file target_file target The standar the STDERIFILES The input file.	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be an error. The last option specified shall determine the behavior of the utility. ng operands shall be supported: A pathname of a file to be copied. A pathname of an existing or nonexistent file, used for the output when a single file is copied. A pathname of a directory to contain the copied files. rd input shall be used to read an input line in response to each prompt specified in R section. Otherwise, the standard input shall not be used. iles specified as operands may be of any file type.
10536 OB 10537 10538 10539 OPER 10540 10541 10542 10543 10544 10545 STDIN 10546 10547 10548 INPUT 10549 10550 ENVIR	Specifying considered ANDS The following source_file target_file target The standar the STDERIFILES The input file.	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be an error. The last option specified shall determine the behavior of the utility. ng operands shall be supported: A pathname of a file to be copied. A pathname of an existing or nonexistent file, used for the output when a single file is copied. A pathname of a directory to contain the copied files. rd input shall be used to read an input line in response to each prompt specified in a section. Otherwise, the standard input shall not be used. section. Otherwise, the standard input shall not be used. ARIABLES ng environment variables shall affect the execution of cp: Provide a default value for the internationalization variables that are unset or null.
10536 OB 10537 10538 10539 OPER 10540 10541 10542 10543 10544 10545 STDIN 10546 10547 10548 INPUT 10549 10550 ENVIR 10551 10552 10553	Specifying considered ANDS The following source_file target_file target The standar the STDERIFILES The input file The following specific target the standar the STDERIFILES The following specific target target the standar the STDERIFILES The following specific target target the standar th	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be an error. The last option specified shall determine the behavior of the utility. ng operands shall be supported: A pathname of a file to be copied. A pathname of an existing or nonexistent file, used for the output when a single file is copied. A pathname of a directory to contain the copied files. rd input shall be used to read an input line in response to each prompt specified in a section. Otherwise, the standard input shall not be used. slies specified as operands may be of any file type. ARIABLES ng environment variables shall affect the execution of cp: Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
10536 OB 10537 10538 10539 OPER 10540 10541 10542 10543 10544 10545 STDIN 10546 10547 10548 INPUT 10549 10550 ENVIR 10551	Specifying considered ANDS The following source_file target_file target The standar the STDERIFILES The input file The following specific target the standar the STDERIFILES The following specific target target the standar the STDERIFILES The following specific target target the standar th	Copy file hierarchies. The treatment of special files is implementation-defined. more than one of the mutually-exclusive options –H, –L, and –P shall not be an error. The last option specified shall determine the behavior of the utility. ng operands shall be supported: A pathname of a file to be copied. A pathname of an existing or nonexistent file, used for the output when a single file is copied. A pathname of a directory to contain the copied files. rd input shall be used to read an input line in response to each prompt specified in a section. Otherwise, the standard input shall not be used. section. Otherwise, the standard input shall not be used. ARIABLES ng environment variables shall affect the execution of cp: Provide a default value for the internationalization variables that are unset or null.

Utilities cp

10556 LC ALL If set to a non-empty string value, override the values of all the other 10557 internationalization variables. 10558 LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multi-10559 character collating elements used in the extended regular expression defined for 10560 the **yesexpr** locale keyword in the *LC_MESSAGES* category. 10561 Determine the locale for the interpretation of sequences of bytes of text data as LC_CTYPE 10562 characters (for example, single-byte as opposed to multi-byte characters in 10563 arguments and input files) and the behavior of character classes used in the 10564 10565 extended regular expression defined for the yesexpr locale keyword in the LC_MESSAGES category. 10566 LC_MESSAGES 10567 Determine the locale for the processing of affirmative responses that should be 10568 10569

used to affect the format and contents of diagnostic messages written to standard error.

NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 10571 XSI

10572 ASYNCHRONOUS EVENTS

Default. 10573

10574 STDOUT

10570

Not used. 10575

10576 STDERR

A prompt shall be written to standard error under the conditions specified in the DESCRIPTION 10577 section. The prompt shall contain the destination pathname, but its format is otherwise 10578 unspecified. Otherwise, the standard error shall be used only for diagnostic messages. 10579

10580 OUTPUT FILES

The output files may be of any type. 10581

10582 EXTENDED DESCRIPTION

None. 10583

10584 EXIT STATUS

The following exit values shall be returned: 10585

All files were copied successfully. 10586

10587 >0 An error occurred.

10588 CONSEQUENCES OF ERRORS

If cp is prematurely terminated by a signal or error, files or file hierarchies may be only partially 10589 copied and files and directories may have incorrect permissions or access and modification 10590 times. 10591

cp Utilities

10592 APPLICATION USAGE

The difference between $-\mathbf{R}$ and $-\mathbf{r}$ is in the treatment by cp of file types other than regular and directory. The original $-\mathbf{r}$ flag, for historic reasons, does not handle special files any differently from regular files, but always reads the file and copies its contents. This has obvious problems in the presence of special file types; for example, character devices, FIFOs, and sockets. The $-\mathbf{R}$ option is intended to recreate the file hierarchy and the $-\mathbf{r}$ option supports historical practice. It was anticipated that a future version of this volume of IEEE Std 1003.1-2001 would deprecate the $-\mathbf{r}$ option, and for that reason, there has been no attempt to fix its behavior with respect to FIFOs or other file types where copying the file is clearly wrong. However, some implementations support $-\mathbf{r}$ with the same abilities as the $-\mathbf{R}$ defined in this volume of IEEE Std 1003.1-2001. To accommodate them as well as systems that do not, the differences between $-\mathbf{r}$ and $-\mathbf{R}$ are implementation-defined. Implementations may make them identical. The $-\mathbf{r}$ option is marked obsolescent.

The set-user-ID and set-group-ID bits are explicitly cleared when files are created. This is to prevent users from creating programs that are set-user-ID or set-group-ID to them when copying files or to make set-user-ID or set-group-ID files accessible to new groups of users. For example, if a file is set-user-ID and the copy has a different group ID than the source, a new group of users has execute permission to a set-user-ID program than did previously. In particular, this is a problem for superusers copying users' trees.

10611 EXAMPLES

10612 None.

10613 RATIONALE

The -i option exists on BSD systems, giving applications and users a way to avoid accidentally removing files when copying. Although the 4.3 BSD version does not prompt if the standard input is not a terminal, the standard developers decided that use of -i is a request for interaction, so when the destination path exists, the utility takes instructions from whatever responds on standard input.

The exact format of the interactive prompts is unspecified. Only the general nature of the contents of prompts are specified because implementations may desire more descriptive prompts than those used on historical implementations. Therefore, an application using the $-\mathbf{i}$ option relies on the system to provide the most suitable dialog directly with the user, based on the behavior specified.

The -p option is historical practice on BSD systems, duplicating the time of last data modification and time of last access. This volume of IEEE Std 1003.1-2001 extends it to preserve the user and group IDs, as well as the file permissions. This requirement has obvious problems in that the directories are almost certainly modified after being copied. This volume of IEEE Std 1003.1-2001 requires that the modification times be preserved. The statement that the order in which the characteristics are duplicated is unspecified is to permit implementations to provide the maximum amount of security for the user. Implementations should take into account the obvious security issues involved in setting the owner, group, and mode in the wrong order or creating files with an owner, group, or mode different from the final value.

It is unspecified whether cp writes diagnostic messages when the user and group IDs cannot be set due to the widespread practice of users using $-\mathbf{p}$ to duplicate some portion of the file characteristics, indifferent to the duplication of others. Historic implementations only write diagnostic messages on errors other than [EPERM].

The $-\mathbf{r}$ option is historical practice on BSD and BSD-derived systems, copying file hierarchies as opposed to single files. This functionality is used heavily in historical applications, and its loss would significantly decrease consensus. The $-\mathbf{R}$ option was added as a close synonym to the $-\mathbf{r}$ option, selected for consistency with all other options in this volume of IEEE Std 1003.1-2001 that

Utilities cp

do recursive directory descent.

 When a failure occurs during the copying of a file hierarchy, *cp* is required to attempt to copy files that are on the same level in the hierarchy or above the file where the failure occurred. It is unspecified if *cp* shall attempt to copy files below the file where the failure occurred (which cannot succeed in any case).

Permissions, owners, and groups of created special file types have been deliberately left as implementation-defined. This is to allow systems to satisfy special requirements (for example, allowing users to create character special devices, but requiring them to be owned by a certain group). In general, it is strongly suggested that the permissions, owner, and group be the same as if the user had run the historical *mknod*, *ln*, or other utility to create the file. It is also probable that additional privileges are required to create block, character, or other implementation-defined special file types.

Additionally, the $-\mathbf{p}$ option explicitly requires that all set-user-ID and set-group-ID permissions be discarded if any of the owner or group IDs cannot be set. This is to keep users from unintentionally giving away special privilege when copying programs.

When creating regular files, historical versions of *cp* use the mode of the source file as modified by the file mode creation mask. Other choices would have been to use the mode of the source file unmodified by the creation mask or to use the same mode as would be given to a new file created by the user (plus the execution bits of the source file) and then modify it by the file mode creation mask. In the absence of any strong reason to change historic practice, it was in large part retained.

When creating directories, historical versions of *cp* use the mode of the source directory, plus read, write, and search bits for the owner, as modified by the file mode creation mask. This is done so that *cp* can copy trees where the user has read permission, but the owner does not. A side effect is that if the file creation mask denies the owner permissions, *cp* fails. Also, once the copy is done, historical versions of *cp* set the permissions on the created directory to be the same as the source directory, unmodified by the file creation mask.

This behavior has been modified so that *cp* is always able to create the contents of the directory, regardless of the file creation mask. After the copy is done, the permissions are set to be the same as the source directory, as modified by the file creation mask. This latter change from historical behavior is to prevent users from accidentally creating directories with permissions beyond those they would normally set and for consistency with the behavior of *cp* in creating files.

It is not a requirement that *cp* detect attempts to copy a file to itself; however, implementations are strongly encouraged to do so. Historical implementations have detected the attempt in most cases.

There are two methods of copying subtrees in this volume of IEEE Std 1003.1-2001. The other method is described as part of the *pax* utility (see *pax*). Both methods are historical practice. The *cp* utility provides a simpler, more intuitive interface, while *pax* offers a finer granularity of control. Each provides additional functionality to the other; in particular, *pax* maintains the hard-link structure of the hierarchy, while *cp* does not. It is the intention of the standard developers that the results be similar (using appropriate option combinations in both utilities). The results are not required to be identical; there seemed insufficient gain to applications to balance the difficulty of implementations having to guarantee that the results would be exactly identical.

The wording allowing *cp* to copy a directory to implementation-defined file types not specified by the System Interfaces volume of IEEE Std 1003.1-2001 is provided so that implementations supporting symbolic links are not required to prohibit copying directories to symbolic links. Other extensions to the System Interfaces volume of IEEE Std 1003.1-2001 file types may need to

cp Utilities

10689	use this loophole as well.	
10690 FUTUR 10691	E DIRECTIONS The –r option may be removed; use –R instead.	
10692 SEE AL 10693	SO mv, find, ln, pax, the System Interfaces volume of IEEE Std 1003.1-2001, open(), unlink()	ı
10694 CHAN (10695	GE HISTORY First released in Issue 2.	
10696 Issue 6 10697	The $-\mathbf{r}$ option is marked obsolescent.	
10698 10699	The new options $-\mathbf{H}$, $-\mathbf{L}$, and $-\mathbf{P}$ are added to align with the IEEE P1003.2b draft standard. These options affect the processing of symbolic links.	
10700	IEEE PASC Interpretation 1003.2 #194 is applied, adding a description of the $-\mathbf{P}$ option.	
10701 10702	IEEE Std 1003.1-2001/Cor 1-2002, item $XCU/TC1/D6/18$ is applied, correcting an error in the SEE ALSO section.	

Utilities crontab

314365				
10703 NAME 10704				
	10705 SYNOPSIS			
10703 311101		crontab [file]		
10707	crontab [[-e -l -r]		
10708				
10709 DESCR	RIPTION			
10710	The crontab utility shall create, replace, or edit a user's crontab entry; a crontab entry is a list of			
10711 10712		and the times at which they shall be executed. The new crontab entry can be input by ile or input from standard input if no file operand is specified, or by using an editor, if		
10713	-e is specific			
10714	Upon execu	tion of a command from a crontab entry, the implementation shall supply a default		
10715		nt, defining at least the following environment variables:		
10716	HOME	A pathname of the user's home directory.		
10717	LOGNAME	The user's login name.		
10718	PATH	A string representing a search path guaranteed to find all of the standard utilities.		
10719 10720	SHELL	A pathname of the command interpreter. When <i>crontab</i> is invoked as specified by this volume of IEEE Std 1003.1-2001, the value shall be a pathname for <i>sh</i> .		
10721		of these variables when crontab is invoked as specified by this volume of		
10722		03.1-2001 shall not affect the default values provided when the scheduled command		
10723	is run.			
10724 10725	If standard output and standard error are not redirected by commands executed from the crontab entry, any generated output or errors shall be mailed, via an implementation-defined			
10726		method, to the user.		
10727 XSI	Users shall	be permitted to use <i>crontab</i> if their names appear in the file /usr/lib/cron/cron.allow.		
10728		If that file does not exist, the file /usr/lib/cron/cron.deny shall be checked to determine whether		
10729 10730		the user shall be denied access to <i>crontab</i> . If neither file exists, only a process with appropriate		
10731		privileges shall be allowed to submit a job. If only cron.deny exists and is empty, global usage shall be permitted. The cron.allow and cron.deny files shall consist of one user name per line.		
10732 OPTIO	NS			
10733		utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section		
10734	12.2, Utility	Syntax Guidelines.		
10735	The following	ng options shall be supported:		
10736	−e	Edit a copy of the invoking user's crontab entry, or create an empty entry to edit if		
10737 10738		the crontab entry does not exist. When editing is complete, the entry shall be installed as the user's crontab entry.		
10739	- l	(The letter ell.) List the invoking user's crontab entry.		
10740	- r	Remove the invoking user's crontab entry.		
10741 OPER		· · · · · · · · · · · · · · · · · · ·		
10742		ng operand shall be supported:		
10743	file	The pathname of a file that contains specifications, in the format defined in the		
10744		INDITE ELECTRICION for grantal entries		

10744

INPUT FILES section, for crontab entries.

crontab Utilities

10745 **STDIN**

10751

10755

10756

10757

10758 10759

10760

10761

10762

10763 10764

10765

10766 10767

10768

10769

10770 10771

10777

See the INPUT FILES section.

10747 INPUT FILES

In the POSIX locale, the user or application shall ensure that a crontab entry is a text file consisting of lines of six fields each. The fields shall be separated by
blank>s. The first five fields shall be integer patterns that specify the following:

- 1. Minute [0,59]
- 10752 2. Hour [0,23]
- 3. Day of the month [1,31]
- 10754 4. Month of the year [1,12]
 - 5. Day of the week ([0,6] with 0=Sunday)

Each of these patterns can be either an asterisk (meaning all valid values), an element, or a list of elements separated by commas. An element shall be either a number or two numbers separated by a hyphen (meaning an inclusive range). The specification of days can be made by two fields (day of the month and day of the week). If month, day of month, and day of week are all asterisks, every day shall be matched. If either the month or day of month is specified as an element or list, but the day of week is an asterisk, the month and day of month fields shall specify the days that match. If both month and day of month are specified as an asterisk, but day of week is an element or list, then only the specified days of the week match. Finally, if either the month or day of month is specified as an element or list, and the day of week is also specified as an element or list, then any day matching either the month and day of month, or the day of week, shall be matched.

The sixth field of a line in a crontab entry is a string that shall be executed by *sh* at the specified times. A percent sign character in this field shall be translated to a <newline>. Any character preceded by a backslash (including the '%') shall cause that character to be treated literally. Only the first line (up to a '%' or end-of-line) of the command field shall be executed by the command interpreter. The other lines shall be made available to the command as standard input.

Blank lines and those whose first non-
blank> is '#' shall be ignored.

The text files /usr/lib/cron/cron.allow and /usr/lib/cron/cron.deny shall contain zero or more user names, one per line, of users who are, respectively, authorized or denied access to the service underlying the *crontab* utility.

10776 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *crontab*:

10778 10779	EDITOR	Determine the editor to be invoked when the $-\mathbf{e}$ option is specified. The default editor shall be vi .
10780 10781 10782 10783	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
10784 10785	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
10786 10787 10788	LC_CTYPE	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).

Utilities crontab

10789 LC_MESSAGES 10790 Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. 10791 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 10792 XSI 10793 ASYNCHRONOUS EVENTS Default. 10794 10795 STDOUT If the -l option is specified, the crontab entry shall be written to the standard output. 10796 10797 STDERR The standard error shall be used only for diagnostic messages. 10798 10799 OUTPUT FILES None. 10800 10801 EXTENDED DESCRIPTION None. 10802 10803 EXIT STATUS The following exit values shall be returned: 10804 Successful completion. 10805 >0 An error occurred. 10806 10807 CONSEQUENCES OF ERRORS 10808 The user's crontab entry is not submitted, removed, edited, or listed. 10809 APPLICATION USAGE The format of the crontab entry shown here is guaranteed only for the POSIX locale. Other 10810 10811 cultures may be supported with substantially different interfaces, although implementations are 10812 encouraged to provide comparable levels of functionality. 10813 The default settings of the HOME, LOGNAME, PATH, and SHELL variables that are given to the scheduled job are not affected by the settings of those variables when *crontab* is run; as stated, 10814 they are defaults. The text about "invoked as specified by this volume of IEEE Std 1003.1-2001" 10815 means that the implementation may provide extensions that allow these variables to be affected 10816 at runtime, but that the user has to take explicit action in order to access the extension, such as 10817 10818 give a new option flag or modify the format of the crontab entry. A typical user error is to type only *crontab*; this causes the system to wait for the new crontab 10819 10820 entry on standard input. If end-of-file is typed (generally <control>-D), the crontab entry is replaced by an empty file. In this case, the user should type the interrupt character, which 10821 prevents the crontab entry from being replaced. 10822 10823 EXAMPLES 10824 1. Clean up **core** files every weekday morning at 3:15 am: 15 3 * * 1-5 find \$HOME -name core 2>/dev/null | xargs rm -f 10825 2. Mail a birthday greeting: 10826 0 12 14 2 * mailx john%Happy Birthday!%Time for lunch. 10827 3. As an example of specifying the two types of days: 10828

0 0 1,15 * 1

crontab Utilities

10830 would run a command on the first and fifteenth of each month, as well as on every 10831 Monday. To specify days by only one field, the other field should be set to '*'; for example: 10832 0 0 * * 1 10833 10834 would run a command only on Mondays. 10835 RATIONALE 10836 All references to a cron daemon and to cron files have been omitted. Although historical implementations have used this arrangement, there is no reason to limit future implementations. 10837 10838 This description of *crontab* is designed to support only users with normal privileges. The format of the input is based on the System V crontab; however, there is no requirement here that the 10839 actual system database used by the cron daemon (or a similar mechanism) use this format 10840 internally. For example, systems derived from BSD are likely to have an additional field 10841 appended that indicates the user identity to be used when the job is submitted. 10842 The –e option was adopted from the SVID as a user convenience, although it does not exist in all 10843 historical implementations. 10844 10845 **FUTURE DIRECTIONS** None. 10846 10847 SEE ALSO 10848 at 10849 CHANGE HISTORY First released in Issue 2. 10850 10851 Issue 6

This utility is marked as part of the User Portability Utilities option.

The normative text is reworded to avoid use of the term "must" for application requirements.

Utilities csplit

```
10854 NAME
10855
              csplit — split files based on context
10856 SYNOPSIS
               csplit [-ks] [-f prefix] [-n number] file arg1 ...argn
10857 UP
10858
10859 DESCRIPTION
              The csplit utility shall read the file named by the file operand, write all or part of that file into
10860
              other files as directed by the arg operands, and write the sizes of the files.
10861
10862 OPTIONS
              The csplit utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
10863
              12.2, Utility Syntax Guidelines.
10864
              The following options shall be supported:
10865
                             Name the created files prefix 00, prefix 01, ..., prefix n. The default is xx 00 ... xx n. If
              -f prefix
10866
                             the prefix argument would create a filename exceeding {NAME MAX} bytes, an
10867
                             error shall result, csplit shall exit with a diagnostic message, and no files shall be
10868
10869
                             created.
              -\mathbf{k}
                             Leave previously created files intact. By default, csplit shall remove created files if
10870
                             an error occurs.
10871
              -n number
                             Use number decimal digits to form filenames for the file pieces. The default shall be
10872
10873
10874
              -\mathbf{s}
                             Suppress the output of file size messages.
10875 OPERANDS
              The following operands shall be supported:
10876
              file
                             The pathname of a text file to be split. If file is '-', the standard input shall be
10877
                             used.
10878
              The operands arg1 . . . argn can be a combination of the following:
10879
              /rexp/[offset]
10880
                             A file shall be created using the content of the lines from the current line up to, but
10881
                             not including, the line that results from the evaluation of the regular expression
10882
                             with offset, if any, applied. The regular expression rexp shall follow the rules for
10883
                             basic regular expressions described in the Base Definitions volume of
10884
                             IEEE Std 1003.1-2001, Section 9.3, Basic Regular Expressions. The application shall
10885
                             use the sequence "\/" to specify a slash character within the rexp. The optional
10886
                             offset shall be a positive or negative integer value representing a number of lines.
10887
                             A positive integer value can be preceded by '+'. If the selection of lines from an
10888
                             offset expression of this type would create a file with zero lines, or one with greater
10889
                             than the number of lines left in the input file, the results are unspecified. After the
10890
                             section is created, the current line shall be set to the line that results from the
10891
                             evaluation of the regular expression with any offset applied. If the current line is
10892
                             the first line in the file and a regular expression operation has not yet been
10893
                             performed, the pattern match of rexp shall be applied from the current line to the
10894
```

10895

10896

Equivalent to \(\text{rexp/[offset]} \), except that no file shall be created for the selected section of the input file. The application shall use the sequence "\%" to specify a

following the current line to the end of the file.

end of the file. Otherwise, the pattern match of rexp shall be applied from the line

csplit Utilities

10900		percent-sign character within the rexp.	
10901 10902 10903	line_no	Create a file from the current line up to (but not including) the line number <code>line_no</code> . Lines in the file shall be numbered starting at one. The current line becomes <code>line_no</code> .	
10904 10905 10906 10907	{num}	Repeat operand. This operand can follow any of the operands described previously. If it follows a <i>rexp</i> type operand, that operand shall be applied <i>num</i> more times. If it follows a <i>line_no</i> operand, the file shall be split every <i>line_no</i> lines, <i>num</i> times, from that point.	
10908 10909	An error sha and the end	all be reported if an operand does not reference a line between the current position of the file.	
10910 STDIN			
10911	See the INPU	UT FILES section.	
10912 INPUT 10913		e shall be a text file.	
10914 ENVIR			
10915	The following	ng environment variables shall affect the execution of <i>csplit</i> :	
10916 10917 10918 10919	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)	
10920 10921	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.	
10922	LC_COLLAT	\mathcal{E}	
10923 10924		Determine the locale for the behavior of ranges, equivalence classes, and multi- character collating elements within regular expressions.	
10925 10926 10927 10928	LC_CTYPE	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) and the behavior of character classes within regular expressions.	
10929 10930 10931	LC_MESSA(GES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.	
10932 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .	
10933 ASYNCHRONOUS EVENTS 10934 If the -k option is specified, created files shall be retained. Otherwise, the default action occurs.			
10935 STDOU	-	•	
10936 10937		${f s}$ option is used, the standard output shall consist of one line per file created, with a llows:	
10938	"%d\n", <	file size in bytes>	
10939 STDER	R		

10939 STDERR

The standard error shall be used only for diagnostic messages.

Utilities csplit

10941 **OUTPUT FILES**

The output files shall contain portions of the original input file; otherwise, unchanged.

10943 EXTENDED DESCRIPTION

10944 None.

10945 EXIT STATUS

10946 The following exit values shall be returned:

10947 0 Successful completion.

10948 >0 An error occurred.

10949 CONSEQUENCES OF ERRORS

By default, created files shall be removed if an error occurs. When the **-k** option is specified, created files shall not be removed if an error occurs.

10952 APPLICATION USAGE

10953 None.

10954 EXAMPLES

10960

10961 10962

10963

10964

10965 10966

10969

1. This example creates four files, **cobol00** ... **cobol03**:

```
10956 csplit -f cobol file '/procedure division/' /par5./ /par16./
```

10957 After editing the split files, they can be recombined as follows:

10958 cat cobol0[0-3] > file

Note that this example overwrites the original file.

2. This example would split the file after the first 99 lines, and every 100 lines thereafter, up to 9 999 lines; this is because lines in the file are numbered from 1 rather than zero, for historical reasons:

```
csplit -k file 100 {99}
```

3. Assuming that **prog.c** follows the C-language coding convention of ending routines with a '}' at the beginning of the line, this example creates a file containing each separate C routine (up to 21) in **prog.c**:

10967 csplit -k prog.c '%main(%' '/^}/+1' {20}

10968 RATIONALE

The $-\mathbf{n}$ option was added to extend the range of filenames that could be handled.

Consideration was given to adding a **–a** flag to use the alphabetic filename generation used by the historical *split* utility, but the functionality added by the **–n** option was deemed to make alphabetic naming unnecessary.

10973 FUTURE DIRECTIONS

10974 None.

10975 SEE ALSO

sed, split

10977 CHANGE HISTORY

First released in Issue 2.

csplit Utilities

10979 **Issue 5**

10980 The FUTURE DIRECTIONS section is added.

10981 **Issue 6**

This utility is marked as part of the User Portability Utilities option.

10983 The APPLICATION USAGE section is added.

The description of regular expression operands is changed to align with the IEEE P1003.2b draft

standard.

The normative text is reworded to avoid use of the term "must" for application requirements.

Utilities ctags

10987 **NAME** 10988 ctags — create a tags file (**DEVELOPMENT**, **FORTRAN**) 10989 SYNOPSIS ctags [-a][-f tagsfile] pathname 10990 UP 10991 ctags -x pathname ... 10992 10993 **DESCRIPTION** The ctags utility shall be provided on systems that support the User Portability Utilities option, 10994 the Software Development Utilities option, and either or both of the C-Language Development 10995 Utilities option and FORTRAN Development Utilities option. On other systems, it is optional. 10996 The ctags utility shall write a tagsfile or an index of objects from C-language or FORTRAN source 10997 files specified by the pathname operands. The tagsfile shall list the locators of language-specific 10998 objects within the source files. A locator consists of a name, pathname, and either a search 10999 pattern or a line number that can be used in searching for the object definition. The objects that 11000 shall be recognized are specified in the EXTENDED DESCRIPTION section. 11001 11002 OPTIONS The ctags utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 11003 11004 12.2, Utility Syntax Guidelines. The following options shall be supported: 11005 11006 -a Append to tagsfile. −**f** tagsfile Write the object locator lists into tagsfile instead of the default file named tags in 11007 11008 the current directory. Produce a list of object names, the line number, and filename in which each is 11009 $-\mathbf{x}$ 11010 defined, as well as the text of that line, and write this to the standard output. A 11011 *tagsfile* shall not be created when –**x** is specified. 11012 OPERANDS The following *pathname* operands are supported: 11013 11014 file.c Files with basenames ending with the .c suffix shall be treated as C-language source code. Such files that are not valid input to *c99* produce unspecified results. 11015 file.h 11016 Files with basenames ending with the .h suffix shall be treated as C-language source code. Such files that are not valid input to c99 produce unspecified results. 11017

The handling of other files is implementation-defined.

unspecified results.

11022 **STDIN**

11018

11019 11020

See the INPUT FILES section.

file.f

11024 INPUT FILES

The input files shall be text files containing source code in the language indicated by the operand filename suffixes.

Files with basenames ending with the .f suffix shall be treated as FORTRAN-language source code. Such files that are not valid input to fort77 produce

ctags Utilities

11027 ENVIR	ONMENT V	ARIABLES	
11028	The following environment variables shall affect the execution of <i>ctags</i> :		
11029 11030 11031 11032	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)	
11033 11034	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.	
11035	LC_COLLA	TE	
11036 11037		Determine the order in which output is sorted for the $-\mathbf{x}$ option. The POSIX locale determines the order in which the <i>tagsfile</i> is written.	
11038 11039 11040 11041 11042	LC_CTYPE	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). When processing C-language source code, if the locale is not compatible with the C locale described by the ISO C standard, the results are unspecified.	
11043	LC_MESSA	GES	
11044		Determine the locale that should be used to affect the format and contents of	
11045		diagnostic messages written to standard error.	
11046 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .	
11047 ASYNCHRONOUS EVENTS 11048 Default.			
11049 STDO	JT		
11050 11051	The list of object name information produced by the $-\mathbf{x}$ option shall be written to standard output in the following format:		
11052	"%s %d %s %s", <object-name>, <line-number>, <filename>, <text></text></filename></line-number></object-name>		
11053	where < text	> is the text of line < line-number> of file < filename>.	
11054 STDFR	11054 STDERR		
11054 51521			
11056 OUTPUT FILES			
11057		x option is not specified, the format of the output file shall be:	
11058	"%s\t%s\t	:/%s/\n", <identifier>, <filename>, <pattern></pattern></filename></identifier>	
11059 11060 11061	where <i><pattern></pattern></i> is a search pattern that could be used by an editor to find the defining instance of <i><identifier></identifier></i> in <i><filename></filename></i> (where <i>defining instance</i> is indicated by the declarations listed in the EXTENDED DESCRIPTION).		
11062 11063 11064 11065	An optional circumflex ('^') can be added as a prefix to <i><pattern></pattern></i> , and an optional dollar sign can be appended to <i><pattern></pattern></i> to indicate that the pattern is anchored to the beginning (end) of a line of text. Any slash or backslash characters in <i><pattern></pattern></i> shall be preceded by a backslash character. The anchoring circumflex, dollar sign, and escaping backslash characters shall not be		

considered literal characters.

11066

11067

considered part of the search pattern. All other characters in the search pattern shall be

Utilities ctags

11068

An alternative format is:

"%s\t%s\t?%s?\n", <identifier>, <filename>, <pattern> 11069 which is identical to the first format except that slashes in *pattern>* shall not be preceded by 11070 11071 escaping backslash characters, and question mark characters in *<pattern>* shall be preceded by 11072 backslash characters. A second alternative format is: 11073 "%s\t%s\t%d\n", <identifier>, <filename>, <lineno> 11074 where *< lineno>* is a decimal line number that could be used by an editor to find *< identifier>* in 11075 <filename>. 11076 11077 Neither alternative format shall be produced by ctags when it is used as described by 11078 IEEE Std 1003.1-2001, but the standard utilities that process tags files shall be able to process those formats as well as the first format. 11079 11080 In any of these formats, the file shall be sorted by identifier, based on the collation sequence in the POSIX locale. 11081 11082 EXTENDED DESCRIPTION If the operand identifies C-language source, the ctags utility shall attempt to produce an output 11083 line for each of the following objects: 11084 Function definitions 11085 Type definitions 11086 11087 Macros with arguments It may also produce output for any of the following objects: 11088 11089 Function prototypes 11090 Structures Unions 11091 Global variable definitions 11092 Enumeration types 11093 11094 Macros without arguments 11095 #define statements #line statements 11096 Any #if and #ifdef statements shall produce no output. The tag main is treated specially in C 11097 11098 programs. The tag formed shall be created by prefixing M to the name of the file, with the trailing .c, and leading pathname components (if any) removed. 11099 11100 On systems that do not support the C-Language Development Utilities option, ctags produces 11101 unspecified results for C-language source code files. It should write to standard error a message 11102 identifying this condition and cause a non-zero exit status to be produced. If the operand identifies FORTRAN source, the ctags utility shall produce an output line for each 11103 function definition. It may also produce output for any of the following objects: 11104 • Subroutine definitions 11105

COMMON statements

ctags Utilities

- PARAMETER statements
- DATA and BLOCK DATA statements
- Statement numbers

On systems that do not support the FORTRAN Development Utilities option, *ctags* produces unspecified results for FORTRAN source code files. It should write to standard error a message identifying this condition and cause a non-zero exit status to be produced.

It is implementation-defined what other objects (including duplicate identifiers) produce output.

11114 EXIT STATUS

11113

11124

11125 11126

11127

11136 11137

11138

11139

11140

11141

11142 11143

11144

11145

11146

11147 11148

11149

11115 The following exit values shall be returned:

11116 0 Successful completion.

11117 >0 An error occurred.

11118 CONSEQUENCES OF ERRORS

11119 Default.

11120 APPLICATION USAGE

The output with $-\mathbf{x}$ is meant to be a simple index that can be written out as an off-line readable function index. If the input files to *ctags* (such as $.\mathbf{c}$ files) were not created using the same locale as that in effect when $ctags - \mathbf{x}$ is run, results might not be as expected.

The description of C-language processing says "attempts to" because the C language can be greatly confused, especially through the use of **#defines**, and this utility would be of no use if the real C preprocessor were run to identify them. The output from *ctags* may be fooled and incorrect for various constructs.

11128 EXAMPLES

11129 None.

11130 RATIONALE

The option list was significantly reduced from that provided by historical implementations. The **F** option was omitted as redundant, since it is the default. The **B** option was omitted as being

of very limited usefulness. The **-t** option was omitted since the recognition of **typedefs** is now

required for C source files. The **-u** option was omitted because the update function was judged

to be not only inefficient, but also rarely needed.

An early proposal included a $-\mathbf{w}$ option to suppress warning diagnostics. Since the types of such diagnostics could not be described, the option was omitted as being not useful.

The text for *LC_CTYPE* about compatibility with the C locale acknowledges that the ISO C standard imposes requirements on the locale used to process C source. This could easily be a superset of that known as "the C locale" by way of implementation extensions, or one of a few alternative locales for systems supporting different codesets. No statement is made for FORTRAN because the ANSI X3.9-1978 standard (FORTRAN 77) does not (yet) define a similar locale concept. However, a general rule in this volume of IEEE Std 1003.1-2001 is that any time that locales do not match (preparing a file for one locale and processing it in another), the results are suspect.

The collation sequence of the tags file is not affected by *LC_COLLATE* because it is typically not used by human readers, but only by programs such as *vi* to locate the tag within the source files. Using the POSIX locale eliminates some of the problems of coordinating locales between the *ctags* file creator and the *vi* file reader.

Utilities ctags

11150 Historically, the tags file has been used only by ex and vi. However, the format of the tags file 11151 has been published to encourage other programs to use the tags in new ways. The format allows either patterns or line numbers to find the identifiers because the historical *vi* recognizes either. 11152 The ctags utility does not produce the format using line numbers because it is not useful 11153 11154 following any source file changes that add or delete lines. The documented search patterns 11155 match historical practice. It should be noted that literal leading circumflex or trailing dollar-sign characters in the search pattern will only behave correctly if anchored to the beginning of the 11156 line or end of the line by an additional circumflex or dollar-sign character. 11157 Historical implementations also understand the objects used by the languages Pascal and 11158 sometimes LISP, and they understand the C source output by lex and yacc. The ctags utility is 11159 not required to accommodate these languages, although implementors are encouraged to do so. 11160 11161 The following historical option was not specified, as vgrind is not included in this volume of IEEE Std 1003.1-2001: 11162 If the -v flag is given, an index of the form expected by vgrind is produced on the 11163 $-\mathbf{v}$ standard output. This listing contains the function name, filename, and page 11164 number (assuming 64-line pages). Since the output is sorted into lexicographic 11165 11166 order, it may be desired to run the output through *sort* –**f**. Sample use: ctags -v files | sort -f > index vgrind -x index 11167 The special treatment of the tag main makes the use of ctags practical in directories with more 11168 than one program. 11169 11170 FUTURE DIRECTIONS 11171 None. 11172 SEE ALSO 11173 c99, fort77, vi 11174 CHANGE HISTORY First released in Issue 4. 11175 11176 **Issue 5** The FUTURE DIRECTIONS section is added. 11177 11178 Issue 6 11179 This utility is marked as part of the User Portability Utilities option. The OUTPUT FILES section is changed to align with the IEEE P1003.2b draft standard. 11180 11181 The normative text is reworded to avoid use of the term "must" for application requirements. IEEE PASC Interpretation 1003.2 #168 is applied, changing "create" to "write" in the 11182

11183

DESCRIPTION.

cut Utilities

```
11184 NAME
              cut — cut out selected fields of each line of a file
11185
11186 SYNOPSIS
              cut -b list [-n] [file ...]
11187
              cut -c list [file ...]
11188
              cut -f list [-d delim] [-s] [file ...]
11189
11190 DESCRIPTION
              The cut utility shall cut out bytes (-b option), characters (-c option), or character-delimited fields
11191
11192
              (-f option) from each line in one or more files, concatenate them, and write them to standard
11193
11194 OPTIONS
              The cut utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
11195
              12.2, Utility Syntax Guidelines.
11196
              The application shall ensure that the option-argument list (see options -\mathbf{b}, -\mathbf{c}, and -\mathbf{f} below) is a
11197
              comma-separated list or <br/> separated list of positive numbers and ranges. Ranges can be
11198
              in three forms. The first is two positive numbers separated by a hyphen (low-high), which
11199
              represents all fields from the first number to the second number. The second is a positive
11200
              number preceded by a hyphen (-high), which represents all fields from field number 1 to that
11201
              number. The third is a positive number followed by a hyphen (low-), which represents that
11202
              number to the last field, inclusive. The elements in list can be repeated, can overlap, and can be
11203
              specified in any order, but the bytes, characters, or fields selected shall be written in the order of
11204
              the input data. If an element appears in the selection list more than once, it shall be written
11205
11206
              exactly once.
```

The following options shall be supported:

1120.	1110 10110 11	8 observe sum as authorisen.
11208 11209	− b list	Cut based on a <i>list</i> of bytes. Each selected byte shall be output unless the $-\mathbf{n}$ option is also specified. It shall not be an error to select bytes not present in the input line.
11210 11211	−c list	Cut based on a <i>list</i> of characters. Each selected character shall be output. It shall not be an error to select characters not present in the input line.
11212	$-\mathbf{d}$ delim	Set the field delimiter to the character <i>delim</i> . The default is the <tab>.</tab>
11213 11214 11215 11216 11217	−f list	Cut based on a <i>list</i> of fields, assumed to be separated in the file by a delimiter character (see $-\mathbf{d}$). Each selected field shall be output. Output fields shall be separated by a single occurrence of the field delimiter character. Lines with no field delimiters shall be passed through intact, unless $-\mathbf{s}$ is specified. It shall not be an error to select fields not present in the input line.
11218 11219	-n	Do not split characters. When specified with the -b option, each element in <i>list</i> of the form <i>low-high</i> (hyphen-separated numbers) shall be modified as follows:

• If the byte selected by *low* is not the first byte of a character, *low* shall be decremented to select the first byte of the character originally selected by *low*. If the byte selected by *high* is not the last byte of a character, *high* shall be decremented to select the last byte of the character prior to the character originally selected by *high*, or zero if there is no prior character. If the resulting range element has *high* equal to zero or *low* greater than *high*, the list element shall be dropped from *list* for that input line without causing an error.

Each element in *list* of the form *low*– shall be treated as above with *high* set to the number of bytes in the current line, not including the terminating <newline>. Each

11207

11220

11221

11222

11223

11224

11225

11226 11227

11228

Utilities cut

11229 11230 11231		element in <i>list</i> of the form <i>-high</i> shall be treated as above with <i>low</i> set to 1. Each element in <i>list</i> of the form <i>num</i> (a single number) shall be treated as above with <i>low</i> set to <i>num</i> and <i>high</i> set to <i>num</i> .
11232 – 11233	- s	Suppress lines with no delimiter characters, when used with the $-\mathbf{f}$ option. Unless specified, lines with no delimiters shall be passed through untouched.
11234 OPERAN		g an around shall be arranged a
		g operand shall be supported:
11236 f. 11237	île	A pathname of an input file. If no <i>file</i> operands are specified, or if a <i>file</i> operand is $'-'$, the standard input shall be used.
11238 STDIN	-1	
		d input shall be used only if no <i>file</i> operands are specified, or if a <i>file</i> operand is $'-'$. UT FILES section.
11241 INPUT FI		
	•	es shall be text files, except that line lengths shall be unlimited.
11243 ENVIRO I 11244 T		ARIABLES ag environment variables shall affect the execution of <i>cut</i> :
11245 I	LANG	Provide a default value for the internationalization variables that are unset or null.
11246		(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables
11247 11248		used to determine the values of locale categories.)
11249 I	LC_ALL	If set to a non-empty string value, override the values of all the other
11250	30_7 122	internationalization variables.
11251 <i>I</i> 11252	LC_CTYPE	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
11253		arguments and input files).
11254 I	LC_MESSAC	GES
11255 11256		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
11257 XSI - 1	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .
11258 ASYNCH 11259 I	RONOUS I Default.	EVENTS
11260 STDOUT	ı	
	The <i>cut</i> utilit he following	ty output shall be a concatenation of the selected bytes, characters, or fields (one of g):
11263	'%s\n", <	concatenation of bytes>
11264 "	'%s\n", <	concatenation of characters>
11265 "	'%s\n", <	concatenation of fields and field delimiters>
11266 STDERR		
	Γhe standard	d error shall be used only for diagnostic messages.
11268 OUTPUT 11269 N	FILES None.	

cut Utilities

11270 EXTENDED DESCRIPTION

11271 None.

11272 EXIT STATUS

The following exit values shall be returned:

11274 0 All input files were output successfully.

>0 An error occurred.

11276 CONSEQUENCES OF ERRORS

11277 Default.

11278 APPLICATION USAGE

Earlier versions of the *cut* utility worked in an environment where bytes and characters were considered equivalent (modulo <backspace> and <tab> processing in some implementations). In the extended world of multi-byte characters, the new -b option has been added. The -n option (used with -b) allows it to be used to act on bytes rounded to character boundaries. The algorithm specified for -n guarantees that:

```
11284 cut -b 1-500 -n file > file1
11285 cut -b 501- -n file > file2
```

ends up with all the characters in **file** appearing exactly once in **file1** or **file2**. (There is, however, a <newline> in both **file1** and **file2** for each <newline> in **file**.)

11288 EXAMPLES

11289 Examples of the option qualifier list:

11290 1,4,7 Select the first, fourth, and seventh bytes, characters, or fields and field delimiters.

11291 1–3,8 Equivalent to 1,2,3,8.

−5,10 Equivalent to 1,2,3,4,5,10.

11293 3– Equivalent to third to last, inclusive.

The *low-high* forms are not always equivalent when used with $-\mathbf{b}$ and $-\mathbf{n}$ and multi-byte characters; see the description of $-\mathbf{n}$.

The following command:

11297 cut -d : -f 1,6 /etc/passwd

reads the System V password file (user database) and produces lines of the form:

11299 <user ID>:<home directory>

Most utilities in this volume of IEEE Std 1003.1-2001 work on text files. The *cut* utility can be used to turn files with arbitrary line lengths into a set of text files containing the same data. The *paste* utility can be used to create (or recreate) files with arbitrary line lengths. For example, if **file** contains long lines:

```
11304 cut -b 1-500 -n file > file1
11305 cut -b 501- -n file > file2
```

creates **file1** (a text file) with lines no longer than 500 bytes (plus the <newline>) and **file2** that contains the remainder of the data from **file**. (Note that **file2** is not a text file if there are lines in **file** that are longer than 500 + {LINE_MAX} bytes.) The original file can be recreated from **file1** and **file2** using the command:

paste -d "\0" file1 file2 > file

Utilities cut

11311 RATIONALE 11312 Some historical implementations do not count

backspace>s in determining character counts 11313 with the -c option. This may be useful for using *cut* for processing *nroff* output. It was deliberately decided not to have the -c option treat either
backspace>s or <tab>s in any special 11314 11315 fashion. The *fold* utility does treat these characters specially. Unlike other utilities, some historical implementations of *cut* exit after not finding an input file, 11316 11317 rather than continuing to process the remaining file operands. This behavior is prohibited by this 11318 volume of IEEE Std 1003.1-2001, where only the exit status is affected by this problem. The behavior of *cut* when provided with either mutually-exclusive options or options that do 11319 11320 not work logically together has been deliberately left unspecified in favor of global wording in Section 1.11 (on page 20). 11321 The OPTIONS section was changed in response to IEEE PASC Interpretation 1003.2 #149. The 11322 change represents historical practice on all known systems. The original standard was 11323 ambiguous on the nature of the output. 11324 The *list* option-arguments are historically used to select the portions of the line to be written, but 11325 do not affect the order of the data. For example: 11326 11327 echo abcdefghi | cut -c6,2,4-7,1 yields "abdefg". 11328 A proposal to enhance *cut* with the following option: 11329 **−o** Preserve the selected field order. When this option is specified, each byte, character, or field 11330 11331 (or ranges of such) shall be written in the order specified by the *list* option-argument, even if 11332 this requires multiple outputs of the same bytes, characters, or fields. 11333 was rejected because this type of enhancement is outside the scope of the IEEE P1003.2b draft 11334 standard. 11335 FUTURE DIRECTIONS 11336 None. 11337 **SEE ALSO** 11338 grep, paste, Section 2.5 (on page 33) 11339 CHANGE HISTORY 11340 First released in Issue 2. 11341 **Issue 6** 11342 The OPTIONS section is changed to align with the IEEE P1003.2b draft standard.

The normative text is reworded to avoid use of the term "must" for application requirements.

11343

cxref Utilities

```
11344 NAME
11345
              cxref — generate a C-language program cross-reference table (DEVELOPMENT)
11346 SYNOPSIS
              cxref [-cs][-o file][-w num] [-D name[=def]]...[-I dir]...
11347 XSI
11348
                    [-U name]... file ...
11349
11350 DESCRIPTION
              The cxref utility shall analyze a collection of C-language files and attempt to build a cross-
11351
              reference table. Information from #define lines shall be included in the symbol table. A sorted
11352
11353
              listing shall be written to standard output of all symbols (auto, static, and global) in each file
              separately, or with the -c option, in combination. Each symbol shall contain an asterisk before
11354
              the declaring reference.
11355
11356 OPTIONS
              The cxref utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
11357
              12.2, Utility Syntax Guidelines, except that the order of the -D, -I, and -U options (which are
11358
              identical to their interpretation by c99) is significant. The following options shall be supported:
11359
                            Write a combined cross-reference of all input files.
11360
              -\mathbf{c}
                            Operate silently; do not print input filenames.
              -s
11361
              −o file
11362
                            Direct output to named file.
11363
              -w num
                            Format output no wider than num (decimal) columns. This option defaults to 80 if
                            num is not specified or is less than 51.
11364
              -\mathbf{D}
                            Equivalent to c99.
11365
              -\mathbf{I}
11366
                            Equivalent to c99.
              -\mathbf{U}
                            Equivalent to c99.
11367
11368 OPERANDS
              The following operand shall be supported:
11369
              file
11370
                            A pathname of a C-language source file.
11371 STDIN
              Not used.
11379
11373 INPUT FILES
11374
              The input files are C-language source files.
11375 ENVIRONMENT VARIABLES
11376
              The following environment variables shall affect the execution of cxref:
              LANG
                            Provide a default value for the internationalization variables that are unset or null.
11377
                            (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
11378
                            Internationalization Variables for the precedence of internationalization variables
11379
                            used to determine the values of locale categories.)
11380
              LC\_ALL
                            If set to a non-empty string value, override the values of all the other
11381
                            internationalization variables.
11382
              LC_COLLATE
11383
                            Determine the locale for the ordering of the output.
11384
              LC_CTYPE
                            Determine the locale for the interpretation of sequences of bytes of text data as
11385
11386
                            characters (for example, single-byte as opposed to multi-byte characters in
```

Utilities cxref

11387 arguments and input files). LC_MESSAGES 11388 11389 Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. 11390 11391 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 11392 ASYNCHRONOUS EVENTS 11393 Default. 11394 STDOUT 11395 The standard output shall be used for the cross-reference listing, unless the $-\mathbf{o}$ option is used to select a different output file. 11396 The format of standard output is unspecified, except that the following information shall be 11397 included: 11398 11399 • If the -c option is not specified, each portion of the listing shall start with the name of the input file on a separate line. 11400 • The name line shall be followed by a sorted list of symbols, each with its associated location 11401 pathname, the name of the function in which it appears (if it is not a function name itself), 11402 and line number references. 11403 • Each line number may be preceded by an asterisk ('*') flag, meaning that this is the 11404 declaring reference. Other single-character flags, with implementation-defined meanings, 11405 may be included. 11406 11407 STDERR The standard error shall be used only for diagnostic messages. 11408 11409 OUTPUT FILES The output file named by the $-\mathbf{o}$ option shall be used instead of standard output. 11410 11411 EXTENDED DESCRIPTION 11412 None. 11413 EXIT STATUS The following exit values shall be returned: 11414 Successful completion. 11415 >0 An error occurred. 11416 11417 CONSEQUENCES OF ERRORS 11418 Default. 11419 APPLICATION USAGE None. 11420 11421 EXAMPLES 11422 None. 11423 RATIONALE None. 11424 11425 FUTURE DIRECTIONS None. 11426

cxref Utilities

11427 **SEE ALSO** 11428 *c99*

11429 CHANGE HISTORY

First released in Issue 2.

11431 **Issue 5**

In the SYNOPSIS, [-U dir] is changed to [-U name].

11433 **Issue 6**

11434 The APPLICATION USAGE section is added.

Utilities date

11435 NAME			
11436	date — write	e the date	e and time
11437 SYNOP	SIS		
11438	date [-u]	[+form	nat]
11439 XSI	date [-u]	mmddhl	nmm[[cc]yy]
11440			
11441 DESCR			
11442 XSI			write the date and time to standard output or attempt to set the system date
11443 11444			, the current date and time shall be written. If an operand beginning with output format of <i>date</i> shall be controlled by the conversion specifications
11445	and other te		
11446 OPTIO	NS		
11447		lity shall	conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
11448	12.2, Utility	Syntax G	uidelines.
11449	The following	ng option	shall be supported:
11450	-u		n operations as if the TZ environment variable was set to the string "UTCO",
11451			equivalent historical value of "GMTO". Otherwise, date shall use the
11452 11453			ne indicated by the TZ environment variable or the system default if that is unset or null.
	NDC	variabi	is unset of fight.
11454 OPERA 11455		ng operar	nds shall be supported:
11456	+format	· .	the format is specified, each conversion specifier shall be replaced in the
11457	+101111at		d output by its corresponding value. All other characters shall be copied to
11458		the ou	tput without change. The output shall always be terminated with a
11459		<newli< td=""><td>ne>.</td></newli<>	ne>.
11460		Conver	sion Specifications
11461		%a	Locale's abbreviated weekday name.
11462		%A	Locale's full weekday name.
11463		%b	Locale's abbreviated month name.
11464		%B	Locale's full month name.
11465		%C	Locale's appropriate date and time representation.
11466		%C	Century (a year divided by 100 and truncated to an integer) as a decimal
11467			number [00,99].
11468		%d	Day of the month as a decimal number [01,31].
11469		%D	Date in the format $mm/dd/yy$.
11470 11471		%e	Day of the month as a decimal number [1,31] in a two-digit field with leading space character fill.
11472		%h	A synonym for %b.
11473		%H	Hour (24-hour clock) as a decimal number [00,23].
11474		%I	Hour (12-hour clock) as a decimal number [01,12].

date Utilities

11475	%j	Day of the year as a decimal number [001,366].
11476	%m	Month as a decimal number [01,12].
11477	%M	Minute as a decimal number [00,59].
11478	%n	A <newline>.</newline>
11479	%p	Locale's equivalent of either AM or PM.
11480 11481	%r	12-hour clock time [01,12] using the AM/PM notation; in the POSIX locale, this shall be equivalent to $\mbox{$\tt I:\M:\Sp.}$
11482	%S	Seconds as a decimal number [00,60].
11483	%t	A <tab>.</tab>
11484	%T	24-hour clock time [00,23] in the format HH:MM:SS.
11485	%u	Weekday as a decimal number [1,7] (1=Monday).
11486 11487 11488	%U	Week of the year (Sunday as the first day of the week) as a decimal number $[00,53]$. All days in a new year preceding the first Sunday shall be considered to be in week 0.
11489 11490 11491 11492	%V	Week of the year (Monday as the first day of the week) as a decimal number [01,53]. If the week containing January 1 has four or more days in the new year, then it shall be considered week 1; otherwise, it shall be the last week of the previous year, and the next week shall be week 1.
11493	%W	Weekday as a decimal number [0,6] (0=Sunday).
11494 11495 11496	%W	Week of the year (Monday as the first day of the week) as a decimal number [00,53]. All days in a new year preceding the first Monday shall be considered to be in week 0.
11497	%x	Locale's appropriate date representation.
11498	%X	Locale's appropriate time representation.
11499	%y	Year within century [00,99].
11500	%Y	Year with century as a decimal number.
11501	%Z	Timezone name, or no characters if no timezone is determinable.
11502	%%	A percent sign character.
11503 11504		Base Definitions volume of IEEE Std 1003.1-2001, Section 7.3.5, LC_TIME conversion specifier values in the POSIX locale.
11505	Modific	ed Conversion Specifications
11506 11507 11508 11509 11510 11511 11512	indicate descript LC_TIM alt_digi LC_TIM convers	onversion specifiers can be modified by the E and O modifier characters to a different format or specification as specified in the <i>LC_TIME</i> locale tion (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 7.3.5, ME). If the corresponding keyword (see era , era_year , era_d_fmt , and its in the Base Definitions volume of IEEE Std 1003.1-2001, Section 7.3.5, ME) is not specified or not supported for the current locale, the unmodified sion specifier value shall be used.
11513	%EC	Locale's alternative appropriate date and time representation.

Utilities date

11514 11515	%EC	The name of the base year (period) in the locale's alternative representation.
11516	%Ex	Locale's alternative date representation.
11517	%EX	Locale's alternative time representation.
11518	%Ey	Offset from %EC (year only) in the locale's alternative representation.
11519	%EY	Full alternative year representation.
11520	%Od	Day of month using the locale's alternative numeric symbols.
11521	%0e	Day of month using the locale's alternative numeric symbols.
11522	%OH	Hour (24-hour clock) using the locale's alternative numeric symbols.
11523	%OI	Hour (12-hour clock) using the locale's alternative numeric symbols.
11524	%Om	Month using the locale's alternative numeric symbols.
11525	%OM	Minutes using the locale's alternative numeric symbols.
11526	%OS	Seconds using the locale's alternative numeric symbols.
11527 11528	%Ou	Weekday as a number in the locale's alternative representation (Monday $= 1$).
11529 11530	%OU	Week number of the year (Sunday as the first day of the week) using the locale's alternative numeric symbols.
11531 11532	%OV	Week number of the year (Monday as the first day of the week, rules corresponding to %V), using the locale's alternative numeric symbols.
11533 11534	%Ow	Weekday as a number in the locale's alternative representation (Sunday = 0).
11535 11536	%OW	Week number of the year (Monday as the first day of the week) using the locale's alternative numeric symbols.
11537	%Oy	Year (offset from %C) in alternative representation.
11538 XSI	mmddhhmm[[cc]yy]	
11539	-	t to set the system date and time from the value given in the operand. This
11540		possible if the user has appropriate privileges and the system permits the
11541 11542		of the system date and time. The first <i>mm</i> is the month (number); <i>dd</i> is the umber); <i>hh</i> is the hour (number, 24-hour system); the second <i>mm</i> is the
11543		(number); cc is the century and is the first two digits of the year (this is
11544		l); yy is the last two digits of the year and is optional. If century is not
11545		d, then values in the range [69,99] shall refer to years 1969 to 1999 inclusive,
11546		lues in the range [00,68] shall refer to years 2000 to 2068 inclusive. The
11547		year is the default if <i>yy</i> is omitted.
11548	Note:	It is expected that in a future version of IEEE Std 1003.1-2001 the default
11549 11550		century inferred from a 2-digit year will change. (This would apply to all commands accepting a 2-digit year as input.)
11551 STDIN		
11551 STDIN 11552	Not used.	

date Utilities

11553 INPUT 1	FILES None.	
		ADIADI EC
11555 ENVIRO 11556		ng environment variables shall affect the execution of <i>date</i> :
11557 11558 11559 11560	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
11561 11562	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
11563 11564 11565	LC_CTYPE	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).
11566	LC_MESSA	GES
11567 11568		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
11569	LC_TIME	Determine the format and contents of date and time strings written by date.
11570 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .
11571 11572 11573	TZ	Determine the timezone in which the time and date are written, unless the $-\mathbf{u}$ option is specified. If the TZ variable is unset or null and $-\mathbf{u}$ is not specified, an unspecified system default timezone is used.
11574 ASYNC 11575	HRONOUS 1 Default.	EVENTS
11576 STDOU	T	
11577 11578	When no for specifying:	rmatting operand is specified, the output in the POSIX locale shall be equivalent to
11579	date "+%a	%b %e %H:%M:%S %Z %Y"
11580 STDER	R	
11581		d error shall be used only for diagnostic messages.
11582 OUTPU 11583	T FILES None.	
11584 EXTEN I	DED DESCR	IPTION
11585	None.	
11586 EXIT ST	TATUS	
11587	The following	ng exit values shall be returned:
11588	0 The date	e was written successfully.
11589	>0 An erro	r occurred.
11590 CONSE	QUENCES O	OF ERRORS
11501	Detaillf	

Default.

11591

Utilities date

11592 APPLICATION USAGE

11593 Conversion specifiers are of unspecified format when not in the POSIX locale. Some of them can 11594 contain <newline>s in some locales, so it may be difficult to use the format shown in standard 11595 output for parsing the output of *date* in those locales.

The range of values for \$S extends from 0 to 60 seconds to accommodate the occasional leap second.

Although certain of the conversion specifiers in the POSIX locale (such as the name of the month) are shown with initial capital letters, this need not be the case in other locales. Programs using these fields may need to adjust the capitalization if the output is going to be used at the beginning of a sentence.

The date string formatting capabilities are intended for use in Gregorian-style calendars, possibly with a different starting year (or years). The <code>%x</code> and <code>%c</code> conversion specifications, however, are intended for local representation; these may be based on a different, non-Gregorian calendar.

The C conversion specification was introduced to allow a fallback for the EC (alternative year format base year); it can be viewed as the base of the current subdivision in the Gregorian calendar. The century number is calculated as the year divided by 100 and truncated to an integer; it should not be confused with the use of ordinal numbers for centuries (for example, "twenty-first century".) Both the EY and Y can then be viewed as the offset from EC and Y, respectively.

The E and O modifiers modify the traditional conversion specifiers, so that they can always be used, even if the implementation (or the current locale) does not support the modifier.

The E modifier supports alternative date formats, such as the Japanese Emperor's Era, as long as these are based on the Gregorian calendar system. Extending the E modifiers to other date elements may provide an implementation-defined extension capable of supporting other calendar systems, especially in combination with the \circ modifier.

The O modifier supports time and date formats using the locale's alternative numerical symbols, such as Kanji or Hindi digits or ordinal number representation.

Non-European locales, whether they use Latin digits in computational items or not, often have local forms of the digits for use in date formats. This is not totally unknown even in Europe; a variant of dates uses Roman numerals for the months: the third day of September 1991 would be written as 3.IX.1991. In Japan, Kanji digits are regularly used for dates; in Arabic-speaking countries, Hindi digits are used. The %d, %e, %H, %I, %m, %S, %U, %w, %W, and %y conversion specifications always return the date and time field in Latin digits (that is, 0 to 9). The %O modifier was introduced to support the use for display purposes of non-Latin digits. In the *LC_TIME* category in *localedef*, the optional **alt_digits** keyword is intended for this purpose. As an example, assume the following (partial) *localedef* source:

```
11629 alt_digits "";"I";"II";"III";"IV";"V";"VI";"VII";"VIII" \
11630 "IX";"X";"XII"
11631 d_fmt "%e.%Om.%Y"
```

11632 With the above date, the command:

11633 date "+%x"

would yield 3.IX.1991. With the same **d_fmt**, but without the **alt_digits**, the command would yield 3.9.1991.

date **Utilities**

```
11636 EXAMPLES
              1. The following are input/output examples of date used at arbitrary times in the POSIX
11637
11638
                 locale:
11639
                 $ date
11640
                 Tue Jun 26 09:58:10 PDT 1990
                 $ date "+DATE: %m/%d/%y%nTIME: %H:%M:%S"
11641
11642
                 DATE: 11/02/91
                 TIME: 13:36:16
11643
11644
                 $ date "+TIME: %r"
                 TIME: 01:36:32 PM
11645
              2. Examples for Denmark, where the default date and time format is %a %d %b %Y %T %Z:
11646
                 $ LANG=da DK.iso 8859-1 date
11647
                 ons 02 okt 1991 15:03:32 CET
11648
11649
                 $ LANG=da DK.iso 8859-1 \
11650
                      date "+DATO: %A den %e. %B %Y%nKLOKKEN: %H:%M:%S"
11651
                 DATO: onsdag den 2. oktober 1991
                 KLOKKEN: 15:03:56
11652
11653
              3. Examples for Germany, where the default date and time format is %a %d.%h.%Y, %T %Z:
11654
                 $ LANG=De DE.88591 date
                 Mi 02.Okt.1991, 15:01:21 MEZ
11655
                 $ LANG=De_DE.88591 date "+DATUM: %A, %d. %B %Y%nZEIT: %H:%M:%S"
11656
                 DATUM: Mittwoch, 02. Oktober 1991
11657
                 ZEIT: 15:02:02
11658
              4. Examples for France, where the default date and time format is %a %d %h %Y %Z %T:
11659
11660
                 $ LANG=Fr FR.88591 date
                 Mer 02 oct 1991 MET 15:03:32
11661
                 $ LANG=Fr FR.88591 date "+JOUR: %A %d %B %Y%nHEURE: %H:%M:%S"
11662
                 JOUR: Mercredi 02 octobre 1991
11663
                 HEURE: 15:03:56
11664
11665 RATIONALE
11666
11667
11668
            implementation, where this option originated.
11669
11670
```

Some of the new options for formatting are from the ISO C standard. The -u option was introduced to allow portable access to Coordinated Universal Time (UTC). The string "GMTO" is allowed as an equivalent TZ value to be compatible with all of the systems using the BSD

The %e format conversion specification (adopted from System V) was added because the ISO C standard conversion specifications did not provide any way to produce the historical default date output during the first nine days of any month.

There are two varieties of day and week numbering supported (in addition to any others created with the locale-dependent %E and %O modifier characters):

 The historical variety in which Sunday is the first day of the week and the weekdays preceding the first Sunday of the year are considered week 0. These are represented by %w and &U. A variant of this is &W, using Monday as the first day of the week, but still referring to week 0. This view of the calendar was retained because so many historical applications depend on it and the ISO C standard strftime() function, on which many date

11671

11672 11673

11674

11675 11676

11677

11678 11679 **Utilities** date

11680 implementations are based, was defined in this way. • The international standard, based on the ISO 8601: 2000 standard where Monday is the first 11681 11682 weekday and the algorithm for the first week number is more complex: If the week (Monday to Sunday) containing January 1 has four or more days in the new year, then it is week 1; 11683 otherwise, it is week 53 of the previous year, and the next week is week 1. These are 11684 represented by the new conversion specifications %u and %V, added as a result of 11685 international comments. 11686 11687 FUTURE DIRECTIONS None. 11688 11689 SEE ALSO The System Interfaces volume of IEEE Std 1003.1-2001, printf(), strftime() 11690 11691 CHANGE HISTORY First released in Issue 2. 11692 11693 **Issue 5** Changes are made for Year 2000 alignment. 11694 11695 **Issue 6** The following new requirements on POSIX implementations derive from alignment with the 11696 Single UNIX Specification: 11697 • The %EX modified conversion specification is added. 11698 The Open Group Corrigendum U048/2 is applied, correcting the examples. 11699 The DESCRIPTION is updated to refer to conversion specifications, instead of field descriptors 11700 for consistency with the *LC_TIME* category. 11701 A clarification is made such that the current year is the default if the yy argument is omitted 11702 when setting the system date and time. 11703 IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/19 is applied, correcting the CHANGE 11704

11705

HISTORY section.

dd Utilities

11710 DESCRIPTION

 The *dd* utility shall copy the specified input file to the specified output file with possible conversions using specific input and output block sizes. It shall read the input one block at a time, using the specified input block size; it shall then process the block of data actually returned, which could be smaller than the requested block size. It shall apply any conversions that have been specified and write the resulting data to the output in blocks of the specified output block size. If the **bs**=*expr* operand is specified and no conversions other than **sync**, **noerror**, or **notrunc** are requested, the data returned from each input block shall be written as a separate output block; if the read returns less than a full block and the **sync** conversion is not specified, the resulting output block shall be the same size as the input block. If the **bs**=*expr* operand is not specified, or a conversion other than **sync**, **noerror**, or **notrunc** is requested, the input shall be processed and collected into full-sized output blocks until the end of the input is reached.

The processing order shall be as follows:

- 1. An input block is read.
- 2. If the input block is shorter than the specified input block size and the **sync** conversion is specified, null bytes shall be appended to the input data up to the specified size. (If either **block** or **unblock** is also specified, <space>s shall be appended instead of null bytes.) The remaining conversions and output shall include the pad characters as if they had been read from the input.
- 3. If the **bs**=*expr* operand is specified and no conversion other than **sync** or **noerror** is requested, the resulting data shall be written to the output as a single block, and the remaining steps are omitted.
- 4. If the **swab** conversion is specified, each pair of input data bytes shall be swapped. If there is an odd number of bytes in the input block, the last byte in the input record shall not be swapped.
- 5. Any remaining conversions (**block**, **unblock**, **lcase**, and **ucase**) shall be performed. These conversions shall operate on the input data independently of the input blocking; an input or output fixed-length record may span block boundaries.
- 6. The data resulting from input or conversion or both shall be aggregated into output blocks of the specified size. After the end of input is reached, any remaining output shall be written as a block without padding if **conv=sync** is not specified; thus, the final output block may be shorter than the output block size.

OPTIONS

11744 None.

11745 OPERANDS

All of the operands shall be processed before any input is read. The following operands shall be supported:

if=*file* Specify the input pathname; the default is standard input.

Specify the output pathname; the default is standard output. If the **seek**=*expr* conversion is not also specified, the output file shall be truncated before the copy begins if an explicit **of**=*file* operand is specified, unless **conv**=**notrunc** is specified.

Utilities dd

11752 11753 11754 11755 11756		preserve the the output f	is specified, but conv=notrunc is not, the effect of the copy shall be to blocks in the output file over which <i>dd</i> seeks, but no other portion of ile shall be preserved. (If the size of the seek plus the size of the input than the previous size of the output file, the output file shall be y the copy.)
11757	ibs=expr	Specify the i	nput block size, in bytes, by <i>expr</i> (default is 512).
11758	obs=expr	Specify the o	output block size, in bytes, by <i>expr</i> (default is 512).
11759 11760 11761	bs=expr	no conversion	ut and output block sizes to <i>expr</i> bytes, superseding ibs = and obs =. If on other than sync , noerror , and notrunc is specified, each input block ied to the output as a single block without aggregating short blocks.
11762 11763 11764	cbs=expr		conversion block size for block and unblock in bytes by <i>expr</i> (default is = is omitted or given a value of zero, using block or unblock produces results.
11765 XSI 11766 11767 11768 11769 11770		operand is swith an asci that characte conv= operathe block v	specified with a value of ascii , ebcdic , or ibm . For a conv = operand i value, the input is handled as described for the unblock value, except ers are converted to ASCII before any trailing <space>s are deleted. For ands with ebcdic or ibm values, the input is handled as described for value except that the characters are converted to EBCDIC or IBM spectively, after any trailing <space>s are added.</space></space>
11772 11773 11774	skip=n	On seekable	t blocks (using the specified input block size) before starting to copy. If files, the implementation shall read the blocks or seek past them; on the files, the blocks shall be read and the data shall be discarded.
11775 11776 11777 11778 11779	seek=n	output file by space from bytes; on se	ks (using the specified output block size) from the beginning of the before copying. On non-seekable files, existing blocks shall be read and the current end-of-file to the specified offset, if any, filled with null bekable files, the implementation shall seek to the specified offset or cks as described for non-seekable files.
11780	count=n	Copy only n	input blocks.
11781 11782	conv=value[-	s are comma-separated symbols from the following list:
11783 XSI		ascii	Convert EBCDIC to ASCII; see Table 4-6 (on page 305).
11784 XSI		ebcdic	Convert ASCII to EBCDIC; see Table 4-6 (on page 305).
11785 XSI		ibm	Convert ASCII to a different EBCDIC set; see Table 4-7 (on page 306).
11786		The ascii , eb	ocdic, and ibm values are mutually-exclusive.
11787 11788 11789 11790 11791 11792 11793 11794 11795 11796		block	Treat the input as a sequence of <newline>-terminated or end-of-file-terminated variable-length records independent of the input block boundaries. Each record shall be converted to a record with a fixed length specified by the conversion block size. Any <newline> shall be removed from the input line; <space>s shall be appended to lines that are shorter than their conversion block size to fill the block. Lines that are longer than the conversion block size shall be truncated to the largest number of characters that fit into that size; the number of truncated lines shall be reported (see the STDERR section).</space></newline></newline>

dd Utilities

11797		The block and unblock values are mutually-exclusive.
11798 11799 11800 11801	unblock	Convert fixed-length records to variable length. Read a number of bytes equal to the conversion block size (or the number of bytes remaining in the input, if less than the conversion block size), delete all trailing <space>s, and append a <newline>.</newline></space>
11802 11803 11804 11805	lcase	Map uppercase characters specified by the <i>LC_CTYPE</i> keyword tolower to the corresponding lowercase character. Characters for which no mapping is specified shall not be modified by this conversion.
11806		The lcase and ucase symbols are mutually-exclusive.
11807 11808 11809 11810	ucase	Map lowercase characters specified by the <i>LC_CTYPE</i> keyword toupper to the corresponding uppercase character. Characters for which no mapping is specified shall not be modified by this conversion.
11811	swab	Swap every pair of input bytes.
11812 11813 11814 11815 11816 11817 11818	noerror	Do not stop processing on an input error. When an input error occurs, a diagnostic message shall be written on standard error, followed by the current input and output block counts in the same format as used at completion (see the STDERR section). If the sync conversion is specified, the missing input shall be replaced with null bytes and processed normally; otherwise, the input block shall be omitted from the output.
11819 11820 11821	notrunc	Do not truncate the output file. Preserve blocks in the output file not explicitly written by this invocation of the dd utility. (See also the preceding of = $file$ operand.)
11822 11823 11824	sync	Pad every input block to the size of the ibs = buffer, appending null bytes. (If either block or unblock is also specified, append <space>s, rather than null bytes.)</space>
11825	The behavior is unspecifie	d if operands other than conv = are specified more than once.
11826 11827		, and obs = operands, the application shall supply an expression The expression, <i>expr</i> , can be:
11828	1. A positive decimal n	number
11829	2. A positive decimal n	number followed by k , specifying multiplication by 1 024
11830	3. A positive decimal n	number followed by b , specifying multiplication by 512
11831 11832	4. Two or more positiv the product of the in	e decimal numbers (with or without k or b) separated by x , specifying dicated values
11833	All of the operands are pro	ocessed before any input is read.
11834 XSI 11835 11836 11837 11838 11839 11840	conversions (first table) a values are the row and co For example, ASCII 0012 inverted tables (for EBC	display the octal number character values used for the ascii and ebcdic and for the ibm conversion (second table). In both tables, the ASCII lumn headers and the EBCDIC values are found at their intersections. (LF) is the second row, third column, yielding 0045 in EBCDIC. The DIC to ASCII conversion) are not shown, but are in one-to-one se tables. The differences between the two tables are highlighted by I five entries.

dd **Utilities**

11841

Table 4-6 ASCII to EBCDIC Conversion

11842

,	0		1	_	7		3		4	_	5		9			,
0000	0000	NUL	0001	SOH	0005	STX	0003	ETX	2900	EOT	0055	ENQ	9500	ACK	0057	BEL
010	0026	BS	9000	노	0045	5	0013		0014	Ή	0015	CR	0016	SO	0017	S
020	0020	DLE	0021	DC1	0022	DC2	0023	DC3	0074	DC4	0075	NAK	0062	SYN	0046	ETB
0030	0030	CAN	0031	E	2200	SUB	0047	ESC	0034	IFS	0035	IGS	9600	IRS	0037	IIB
040	0100	Sp	0132	_	0177		0173	#	0133	↔	0154	%	0120	∞	0175	_
020	0115	J	0135	<u> </u>	0134	*	0116	+	0153	•	0140		0113		0141	_
0900	0360	0	0361	_	0362	7	0363	3	0364	4	0365	2	9980	9	0367	7
020	0370	8	0371	6	0172		0136		0114	٧	0176	II	0156	٨	0157	5
100	0174	@	0301	4	0302	В	0303	ပ	0304	٥	0305	Е	9080	ш	0307	ŋ
110	0310	I	0311	_	0321	7	0322	~	0323	_	0324	Σ	0325	z	0326	0
120	0327	۵	0330	Ø	0331	<u>~</u>	0342	တ	0343	-	0344	⊃	0345	>	0346	≥
1130	0347	×	0320	>	0351	Z	0255	L	0340	_	0275		0232		0155	I
140	0171	,	0201	Ø	0202	q	0203	O	0204	Ф	0205	Φ	0206	Ţ	0207	g
150	0210	۲	0211		0221	_	0222	~	0223	_	0224	٤	0225	_	0226	0
091	0227	۵	0230	ъ	0231	_	0242	S	0243	t	0244	ם	0245	>	0246	>
170	0247	×	0250	у	0251	z	0300	}	0117		0320	}	0137	Г	0007	DEL
200	0040	DS	0041	sos	0042	FS	0043	MUS	0044	ВУР	0025	NL	9000	RNL	0027	POC
210	0020	SA	0051	SFE	0052	SM	0053	CSP	0054	MFA	0011	SPS	0012	RPT	0033	CQ1
220	0900		0061		0032	NBS	0063	<u>~</u>	0064	ЬР	900	TRN	9900	NBS	0010	ЭE
230	0020	SBS	0071	⊨	0072	RFF	0073	CU3	0004	SEL	0024	RES	9200		0341	
240	0101		0102		0103		0104		0105		0106		0107		0110	
1250	0111		0121		0122		0123		0124		0125		0126		0127	
097	0130		0131		0142		0143		0144		0145		0146		0147	
0220	0150		0151		0160		0161		0162		0163		0164		0165	
300	0166		0167		0170		0200		0212		0213		0214		0215	
310	0216		0217		0220		0152		0233		0234		0235		0236	
320	0237		0240		0252		0253		0254		0112	e	0256		0257	
330	0260		0261		0262		0263		0264		0265		0266		0267	
340	0270		0271		0272		0273		0274		0241		0276		0277	
350	0312		0313		0314	5	0315		0316	>-	0317		0332		0333	
360	0334		0335		0336		0337		0352		0353		0354	ℸ	0355	
370	0356		0357		0372	_	0373		0374		0375		0376		0377	Ю

11843 11844

Table 4-7 ASCII to IBM EBCDIC Conversion

	0		7		7		ဗ	_	4		2		9		'	
0000	0000	NOL	0001	SOH	0005	STX	0003	ETX	2900	EOT	0055	ENQ	9500	ACK	0057	BEL
0100	0026	BS	0002	노	0045	느	0013	L	0014	止	0015	S	0016	SO	0017	S
0070	0020	DLE	0021	DC1	0022	DC2	0023	DC3	0074	DC4	0075	NAK	0062	SYN	0046	ETB
0030	0030	CAN	0031	M	0077	SUB	0047	ESC	0034	IFS	0035	IGS	0036	RS	0037	ITB
0040	0100	Sp	0132		0177	=	0173	#	0133	↔	0154	%	0120	৵	0175	_
0020	0115		0135	<u> </u>	0134	*	0116	+	0153	•	0140	ı	0113	•	0141	
0900	0360	0	0361	_	0362	7	0363	ю	0364	4	0365	2	0366	9	0367	7
0020	0370	8	0371	6	0172		0136		0114	٧	0176	=	0156	^	0157	5
0100	0174	@	0301	٧	0302	В	0303	C	0304	D	0305	Ш	9080	ь	0307	Ð
0110	0310	I	0311	_	0321	7	0322	~	0323	_	0324	Σ	0325	z	0326	0
0120	0327	۵	0330	Ø	0331	<u>~</u>	0342	S	0343	-	0344	⊃	0345	>	0346	>
0130	0347	×	0320	>	0351	Z	0255		0340	_	0275	_	0137	Г	0155	ı
0140	0171	,	0201	Ø	0202	q	0203	O	0204	р	0205	Φ	0206	-	0207	D
0150	0210	۲	0211		0221	<u>.</u>	0222	*	0223		0224	٤	0225	د	0226	0
0160	0227	۵	0230	ь	0231	_	0242	S	0243	+	0244	ם	0245	>	0246	>
0110	0247	×	0250	у	0251	z	0300	{	0117		0320	}	0241		0007	DEL
0200	0040	DS	0041	sos	0042	FS	0043	MUS	0044	ВУР	0025	N	9000	RNL	0027	POC
0210	0020	SA	0051	SFE	0052	SM	0053	CSP	0054	MFA	0011	SPS	0012	RPT	0033	CU1
0220	0900		0061		0032	UBS	0063	<u>∝</u>	0064	ЬР	0065	TRN	9900	NBS	0010	GE
0230	000	SBS	0071	⊨	0072	RFF	0073	CU3	0004	SEL	0024	RES	9200		0341	
0240	0101		0102		0103		0104		0105		0106		0107		0110	
0220	0111		0121		0122		0123		0124		0125		0126		0127	
0260	0130		0131		0142		0143		0144		0145		0146		0147	
0220	0150		0151		0160		0161		0162		0163		0164		0165	
0300	0166		0167		0170		0200		0212		0213		0214		0215	
0310	0216		0217		0220		0232		0233		0234		0235		0236	
0320	0237		0240		0252		0253		0254		0255		0256		0257	
0330	0260		0261		0262		0263		0264		0265		0266		0267	
0340	0270		0271		0272		0273		0274		0275		0276		0277	
0320	0312		0313		0314	-	0315		0316	>-	0317		0332		0333	
0360	0334		0335		0336		0337		0352		0353		0354	ェ	0355	
0370	0356		0357		0372		0373		0374		0375		0376		0377	Ш

Utilities dd

11845 **STDIN**

If no **if**= operand is specified, the standard input shall be used. See the INPUT FILES section.

11847 INPUT FILES

The input file can be any file type.

11849 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *dd*:

Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)

11855 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

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), the classification of characters as uppercase or lowercase, and the mapping of characters from one case to the other.

11861 LC_MESSAGES

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

11865 XSI NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

11866 ASYNCHRONOUS EVENTS

For SIGINT, the *dd* utility shall interrupt its current processing, write status information to standard error, and exit as though terminated by SIGINT. It shall take the standard action for all other signals; see the ASYNCHRONOUS EVENTS section in Section 1.11 (on page 20).

11870 STDOUT

If no **of**= operand is specified, the standard output shall be used. The nature of the output depends on the operands selected.

11873 STDERR

On completion, *dd* shall write the number of input and output blocks to standard error. In the POSIX locale the following formats shall be used:

A partial input block is one for which *read*() returned less than the input block size. A partial output block is one that was written with fewer bytes than specified by the output block size.

In addition, when there is at least one truncated block, the number of truncated blocks shall be written to standard error. In the POSIX locale, the format shall be:

```
"%u truncated %s\n", <number of truncated blocks>, "record" (if <number of truncated blocks> is one) "records" (otherwise)
```

Diagnostic messages may also be written to standard error.

dd Utilities

11887 OUTPUT FILES

11888 If the **of**= operand is used, the output shall be the same as described in the STDOUT section.

11889 EXTENDED DESCRIPTION

None. 11890

11891 EXIT STATUS

The following exit values shall be returned: 11892

11893 The input file was copied successfully.

>0 An error occurred. 11894

11895 CONSEQUENCES OF ERRORS

If an input error is detected and the noerror conversion has not been specified, any partial 11896 output block shall be written to the output file, a diagnostic message shall be written, and the 11897 copy operation shall be discontinued. If some other error is detected, a diagnostic message shall 11898 be written and the copy operation shall be discontinued. 11899

11900 APPLICATION USAGE

The input and output block size can be specified to take advantage of raw physical I/O. 11901

There are many different versions of the EBCDIC codesets. The ASCII and EBCDIC conversions 11902 11903 specified for the *dd* utility perform conversions for the version specified by the tables.

11904 EXAMPLES

The following command: 11905

dd if=/dev/rmt0h of=/dev/rmt1h 11906

11907 copies from tape drive 0 to tape drive 1, using a common historical device naming convention.

The following command: 11908

dd ibs=10 skip=1 11909

strips the first 10 bytes from standard input. 11910

This example reads an EBCDIC tape blocked ten 80-byte EBCDIC card images per block into the 11911

ASCII file x: 11912

11913 dd if=/dev/tape of=x ibs=800 cbs=80 conv=ascii,lcase

11914 RATIONALE

11915 11916

11917

11918

11919

11920

11921

11922

11923

11924

11925

11926

The OPTIONS section is listed as "None" because there are no options recognized by historical dd utilities. Certainly, many of the operands could have been designed to use the Utility Syntax Guidelines, which would have resulted in the classic hyphenated option letters. In this version of this volume of IEEE Std 1003.1-2001, dd retains its curious JCL-like syntax due to the large number of applications that depend on the historical implementation.

A suggested implementation technique for conv=noerror, sync is to zero (or <space>-fill, if blocking or unblocking) the input buffer before each read and to write the contents of the input buffer to the output even after an error. In this manner, any data transferred to the input buffer before the error was detected is preserved. Another point is that a failed read on a regular file or a disk generally does not increment the file offset, and dd must then seek past the block on which the error occurred; otherwise, the input error occurs repetitively. When the input is a magnetic tape, however, the tape normally has passed the block containing the error when the error is reported, and thus no seek is necessary.

11927

11928 The default **ibs**= and **obs**= sizes are specified as 512 bytes because there are historical (largely portable) scripts that assume these values. If they were left unspecified, unusual results could 11929

Utilities dd

occur if an implementation chose an odd block size.

Historical implementations of *dd* used *creat()* when processing **of**=*file*. This makes the **seek**= operand unusable except on special files. The **conv=notrunc** feature was added because more recent BSD-based implementations use *open()* (without O_TRUNC) instead of *creat()*, but they fail to delete output file contents after the data copied.

The *w* multiplier (historically meaning *word*), is used in System V to mean 2 and in 4.2 BSD to mean 4. Since *word* is inherently non-portable, its use is not supported by this volume of IEEE Std 1003.1-2001.

Standard EBCDIC does not have the characters '[' and ']'. The values used in the table are taken from a common print train that does contain them. Other than those characters, the print train values are not filled in, but appear to provide some of the motivation for the historical choice of translations reflected here.

The Standard EBCDIC table provides a 1:1 translation for all 256 bytes.

The IBM EBCDIC table does not provide such a translation. The marked cells in the tables differ in such a way that:

- 2. EBCDIC 0137 (' \neg ') translates to/from ASCII 0236 (' $^{\circ}$ '). In the standard table, EBCDIC 0232 (no graphic) is used.
- 3. EBCDIC 0241 ($' \sim '$) translates to/from ASCII 0176 ($' \sim '$). In the standard table, EBCDIC 0137 ($' \sim '$) is used.
- 4. 0255 ('[') and 0275 (']') appear twice, once in the same place as for the standard table and once in place of 0112 (' \diamondsuit ') and 0241 (' \sim ').

In net result:

EBCDIC 0275 (']') displaced EBCDIC 0241 ('~') in cell 0345.

That displaced EBCDIC 0137 ($'\neg'$) in cell 0176.

That displaced EBCDIC 0232 (no graphic) in cell 0136.

That replaced EBCDIC 0152 (broken pipe) in cell 0313.

EBCDIC 0255 (' [') replaced EBCDIC 0112 ('¢').

This translation, however, reflects historical practice that (ASCII) $' \sim '$ and $' \neg '$ were often mapped to each other, as were ' [' and $' \Leftrightarrow '$; and '] ' and (EBCDIC) $' \sim '$.

The **cbs** operand is required if any of the **ascii**, **ebcdic**, or **ibm** operands are specified. For the **ascii** operand, the input is handled as described for the **unblock** operand except that characters are converted to ASCII before the trailing <space>s are deleted. For the **ebcdic** and **ibm** operands, the input is handled as described for the **block** operand except that the characters are converted to EBCDIC or IBM EBCDIC after the trailing <space>s are added.

The **block** and **unblock** keywords are from historical BSD practice.

The consistent use of the word **record** in standard error messages matches most historical practice. An earlier version of System V used **block**, but this has been updated in more recent releases.

Early proposals only allowed two numbers separated by **x** to be used in a product when specifying **bs**=, **cbs**=, **ibs**=, and **obs**= sizes. This was changed to reflect the historical practice of allowing multiple numbers in the product as provided by Version 7 and all releases of System V

dd Utilities

11972 and BSD. A change to the swab conversion is required to match historical practice and is the result of IEEE 11973 11974 PASC Interpretations 1003.2 #03 and #04, submitted for the ISO POSIX-2: 1993 standard. A change to the handling of SIGINT is required to match historical practice and is the result of 11975 11976 IEEE PASC Interpretation 1003.2 #06 submitted for the ISO POSIX-2: 1993 standard. 11977 FUTURE DIRECTIONS 11978 None. 11979 SEE ALSO 11980 Section 1.11 (on page 20), sed, tr 11981 CHANGE HISTORY First released in Issue 2. 11982 11983 Issue 5 11984 The second paragraph of the **cbs**= description is reworded and marked EX. The FUTURE DIRECTIONS section is added. 11985 11986 Issue 6 Changes are made to swab conversion and SIGINT handling to align with the IEEE P1003.2b 11987 11988 draft standard. 11989 The normative text is reworded to avoid use of the term "must" for application requirements. IEEE PASC Interpretation 1003.2 #209 is applied, clarifying the interaction between dd of=file and 11990 conv=notrunc. 11991

delta **Utilities**

11992 **NAME**

11993 delta — make a delta (change) to an SCCS file (**DEVELOPMENT**)

11994 SYNOPSIS

delta [-nps] [-g list] [-m mrlist] [-r SID] [-y[comment]] file... 11995 XSI 11996

11997 **DESCRIPTION**

The delta utility shall be used to permanently introduce into the named SCCS files changes that 11998 were made to the files retrieved by *get* (called the *g-files*, or generated files). 11999

12000 OPTIONS

12004

12015

12016

12017

12018

12019

12020

12021 12022

12023 12024

12025

12026

12027

12028

12029

12030

12031

12032

12033

12034

12035

The delta utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12001 12.2, Utility Syntax Guidelines, except that the -y option has an optional option-argument. This 12002 optional option-argument shall not be presented as a separate argument. 12003

The following options shall be supported:

12005	–r SID	Uniquely identify which delta is to be made to the SCCS file. The use of this option
12006		shall be necessary only if two or more outstanding get commands for editing (get
12007		−e) on the same SCCS file were done by the same person (login name). The SID
12008		value specified with the -r option can be either the SID specified on the get
12009		command line or the SID to be made as reported by the <i>get</i> utility; see <i>get</i> (on page
12010		476).

Suppress the report to standard output of the activity associated with each file. 12011 -sSee the STDOUT section. 12012

Specify retention of the edited g-file (normally removed at completion of delta 12013 -n processing). 12014

> −g list Specify a list (see get for the definition of list) of deltas that shall be ignored when the file is accessed at the change level (SID) created by this delta.

> -m mrlist Specify a modification request (MR) number that the application shall supply as the reason for creating the new delta. This shall be used if the SCCS file has the v flag set; see admin.

> > If -m is not used and '-' is not specified as a file argument, and the standard input is a terminal, the prompt described in the STDOUT section shall be written to standard output before the standard input is read; if the standard input is not a terminal, no prompt shall be issued.

> > MRs in a list shall be separated by

> >

> > blank>s or escaped <newline>s. An unescaped <newline> shall terminate the MR list. The escape character is <backslash>.

> > If the v flag has a value, it shall be taken to be the name of a program which validates the correctness of the MR numbers. If a non-zero exit status is returned from the MR number validation program, the delta utility shall terminate. (It is

assumed that the MR numbers were not all valid.)

-y[comment] Describe the reason for making the delta. The comment shall be an arbitrary group of lines that would meet the definition of a text file. Implementations shall support comments from zero to 512 bytes and may support longer values. A null string (specified as either -y, -y", or in response to a prompt for a comment) shall be considered a valid comment.

delta **Utilities**

12036 12037 12038 12039 12040		If $-y$ is not specified and '-' is not specified as a file argument, and the standard input is a terminal, the prompt described in the STDOUT section shall be written to standard output before the standard input is read; if the standard input is not a terminal, no prompt shall be issued. An unescaped <newline> shall terminate the comment text. The escape character is backslash>.</newline>
12041		The $-y$ option shall be required if the <i>file</i> operand is specified as $'-'$.
12042 12043	- p	Write (to standard output) the SCCS file differences before and after the delta is applied in <i>diff</i> format; see <i>diff</i> .
12044 OPER		1 1 101
12045		ng operand shall be supported:
12046 12047 12048 12049	file	A pathname of an existing SCCS file or a directory. If <i>file</i> is a directory, the <i>delta</i> utility shall behave as though each file in the directory were specified as a named file, except that non-SCCS files (last component of the pathname does not begin with s.) and unreadable files shall be silently ignored.
12050 12051 12052		If exactly one <i>file</i> operand appears, and it is $'-'$, the standard input shall be read; each line of the standard input shall be taken to be the name of an SCCS file to be processed. Non-SCCS files and unreadable files shall be silently ignored.
12053 STDIN		
12054		d input shall be a text file used only in the following cases:
12055		an <i>mrlist</i> or a <i>comment</i> (see the $-\mathbf{m}$ and $-\mathbf{y}$ options).
12056 12057 12058	the com	erand shall be specified as $'-'$. In this case, the $-\mathbf{y}$ option must be used to specify ment, and if the SCCS file has the \mathbf{v} flag set, the $-\mathbf{m}$ option must also be used to he MR list.
12059 INPUT	FILES	
12060 12061 12062 12063	any line of	hall be text files whose data is to be included in the SCCS files. If the first character of an input file is <soh> in the POSIX locale, the results are unspecified. If this file re than 99 999 lines, the number of lines recorded in the header for this file shall be is delta.</soh>
12064 ENVIR	ONMENT VA	
12065	The following	ng environment variables shall affect the execution of <i>delta</i> :
12066 12067 12068 12069	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
12070 12071	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
12072 12073 12074	LC_CTYPE	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).
12075 12076 12077 12078	LC_MESSA	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error, and informative messages written to standard output.

NLSPATH

12079

Determine the location of message catalogs for the processing of *LC_MESSAGES*.

Utilities delta

Determine the timezone in which the time and date are written in the SCCS file. If the TZ variable is unset or NULL, an unspecified system default timezone is used.

12082 ASYNCHRONOUS EVENTS

If SIGINT is caught, temporary files shall be cleaned up and *delta* shall exit with a non-zero exit code. The standard action shall be taken for all other signals; see Section 1.11 (on page 20).

12085 STDOUT

12086 The standard output shall be used only for the following messages in the POSIX locale:

• Prompts (see the -**m** and -**y** options) in the following formats:

```
12088 "MRs? "
12089 "comments? "
```

12090 The MR prompt, if written, shall always precede the comments prompt.

• A report of each file's activities (unless the –s option is specified) in the following format:

```
"%s\n%d inserted\n%d deleted\n%d unchanged\n", <New SID>, <number of lines inserted>, <number of lines deleted>, <number of lines unchanged>
```

12095 STDERR

12091

12092

12093

12094

The standard error shall be used only for diagnostic messages.

12097 OUTPUT FILES

12098 Any SCCS files updated shall be files of an unspecified format.

12099 EXTENDED DESCRIPTION

12100 System Date and Time

When a delta is added to an SCCS file, the system date and time shall be recorded for the new delta. If a *get* is performed using an SCCS file with a date recorded apparently in the future, the behavior is unspecified.

12104 EXIT STATUS

12105 The following exit values shall be returned:

12106 0 Successful completion.

12107 >0 An error occurred.

12108 CONSEQUENCES OF ERRORS

12109 Default.

12110 APPLICATION USAGE

Problems can arise if the system date and time have been modified (for example, put forward and then back again, or unsynchronized clocks across a network) and can also arise when different values of the *TZ* environment variable are used.

Problems of a similar nature can also arise for the operation of the *get* utility, which records the date and time in the file body.

12116 EXAMPLES

12117 None.

delta Utilities

12118 RATIO	NALE				
12119	None.				
12120 FUTURE DIRECTIONS					
12121	None.				
12122 SEE ALSO					
12123	Section 1.11 (on page 20), admin, diff, get, prs, rmdel				
12124 CHANGE HISTORY					
12125	First released in Issue 2.				
12126 Issue 5					
12127	The output format description in the STDOUT section is corrected.				
12128 Issue 6					
12129	The APPLICATION USAGE section is added.				
12130	The normative text is reworded to avoid use of the term "must" for application requirements.				
12131	The Open Group Base Resolution bwg2001-007 is applied as follows:				
12132	• The use of $'-'$ as a file argument is clarified.				
12133	The use of STDIN is added.				
12134	• The ASYNCHRONOUS EVENTS section is updated to remove the implicit requirement that				
12135	implementations re-signal themselves when catching a normally fatal signal.				
12136	• New text is added to the INPUT FILES section warning that the maximum lines recorded in				
12137	the file is 99 999.				
12138	New text is added to the EXTENDED DESCRIPTION and APPLICATION USAGE sections				
12139	regarding how the system date and time may be taken into account, and the TZ environment				
12140 12141	variable is added to the ENVIRONMENT VARIABLES section as per The Open Group Base Resolution bwg2001-007.				
	200000000000000000000000000000000000000				

Utilities df

12142 **NAME** 12143 df — report free disk space 12144 SYNOPSIS 12145 UP XSI df [-k][-P-t][file...]12146 12147 DESCRIPTION The df utility shall write the amount of available space and file slots for file systems on which the 12148 XSI invoking user has appropriate read access. File systems shall be specified by the *file* operands; 12149 when none are specified, information shall be written for all file systems. The format of the 12150 12151 default output from df is unspecified, but all space figures are reported in 512-byte units, unless the -k option is specified. This output shall contain at least the file system names, amount of 12152 available space on each of these file systems, and the number of free file slots, or inodes, 12153 XSI 12154 available; when -t is specified, the output shall contain the total allocated space as well. 12155 OPTIONS The df utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, 12156 Utility Syntax Guidelines. 12157 The following options shall be supported: 12158 $-\mathbf{k}$ Use 1024-byte units, instead of the default 512-byte units, when writing space 12159 figures. 12160 $-\mathbf{P}$ Produce output in the format described in the STDOUT section. 12161 -t Include total allocated-space figures in the output. 12162 XSI 12163 OPERANDS The following operand shall be supported: 12164 12165 file A pathname of a file within the hierarchy of the desired file system. If a file other than a FIFO, a regular file, a directory, or a special file representing the device 12166 XSI containing the file system (for example, /dev/dsk/0s1) is specified, the results are 12167 unspecified. Otherwise, df shall write the amount of free space in the file system 12168 12169 containing the specified *file* operand. 12170 STDIN 12171 Not used. 12172 INPUT FILES 12173 None. 12174 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of df: 12175 LANG Provide a default value for the internationalization variables that are unset or null. 12176 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 12177 12178 Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 12179 12180 LC ALL If set to a non-empty string value, override the values of all the other

internationalization variables.

arguments).

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

12181

12182

12183

12184

 LC_CTYPE

df Utilities

12185 12186 12187 12188	LC_MESSAC	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.			
12189 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.			
12190 ASYNC 12191	HRONOUS I Default.	EVENTS			
12192 STDOU 12193 12194	When both the $-\mathbf{k}$ and $-\mathbf{P}$ options are specified, the following header line shall be written (in the POSIX locale):				
12195	"Filesyste	em 1024-blocks Used Available Capacity Mounted on\n"			
12196 12197	When the $-P$ option is specified without the $-k$ option, the following header line shall be written (in the POSIX locale):				
12198	"Filesyste	em 512-blocks Used Available Capacity Mounted on\n"			
12199 12200	The implementation may adjust the spacing of the header line and the individual data lines so that the information is presented in orderly columns.				
12201 12202	The remaining output with $-P$ shall consist of one line of information for each specified file system. These lines shall be formatted as follows:				
12203 12204 12205	<space< td=""><td>%d %d%% %s\n", <file name="" system="">, <total space="">, e used>, <space free="">, <percentage used="">, system root></percentage></space></total></file></td></space<>	%d %d%% %s\n", <file name="" system="">, <total space="">, e used>, <space free="">, <percentage used="">, system root></percentage></space></total></file>			
12206 12207	In the following list, all quantities expressed in 512-byte units (1 024-byte when $-\mathbf{k}$ is specified) shall be rounded up to the next higher unit. The fields are:				
12208 12209	<file n<="" system="" td=""><td>name> The name of the file system, in an implementation-defined format.</td></file>	name> The name of the file system, in an implementation-defined format.			
12210 12211 12212	<total space=""></total>	The total size of the file system in 512-byte units. The exact meaning of this figure is implementation-defined, but should include <i><space used=""></space></i> , <i><space free=""></space></i> , plus any space reserved by the system not normally available to a user.			
12213 12214	<space used=""></space>	The total amount of space allocated to existing files in the file system, in 512-byte units.			
12215 12216 12217 12218 12219	<space free=""></space>	The total amount of space available within the file system for the creation of new files by unprivileged users, in 512-byte units. When this figure is less than or equal to zero, it shall not be possible to create any new files on the file system without first deleting others, unless the process has appropriate privileges. The figure written may be less than zero.			
12220 12221 12222	<pre><percentage <="" pre="" u=""></percentage></pre>	used> The percentage of the normally available space that is currently allocated to all files on the file system. This shall be calculated using the fraction:			
12223		<pre><space used="">/(<space used="">+ <space free="">)</space></space></space></pre>			
12224 12225 12226		expressed as a percentage. This percentage may be greater than 100 if <i><space free=""></space></i> is less than zero. The percentage value shall be expressed as a positive integer, with any fractional result causing it to be rounded to the next highest integer.			

Utilities df

12227 <file system root> 12228 The directory below which the file system hierarchy appears. 12229 XSI The output format is unspecified when **–t** is used. 12230 STDERR 12231 The standard error shall be used only for diagnostic messages. 12232 OUTPUT FILES 12233 None. 12234 EXTENDED DESCRIPTION 12235 None. 12236 EXIT STATUS 12237 The following exit values shall be returned: Successful completion. 12238 12239 >0 An error occurred. 12240 CONSEQUENCES OF ERRORS 12241 Default. 12242 APPLICATION USAGE 12243 On most systems, the "name of the file system, in an implementation-defined format" is the 12244 special file on which the file system is mounted. 12245 On large file systems, the calculation specified for percentage used can create huge rounding 12246 errors. 12247 EXAMPLES 1. The following example writes portable information about the /usr file system: 19948 12249 df -P /usr 12250 2. Assuming that /usr/src is part of the /usr file system, the following produces the same output as the previous example: 12251 12252 df -P /usr/src 12253 RATIONALE 12254 The behavior of df with the **-P** option is the default action of the 4.2 BSD df utility. The uppercase **−P** was selected to avoid collision with a known industry extension using **−p**. 12255 12256 Historical df implementations vary considerably in their default output. It was therefore 12257 necessary to describe the default output in a loose manner to accommodate all known historical implementations and to add a portable option (-P) to provide information in a portable format. 12258 The use of 512-byte units is historical practice and maintains compatibility with *ls* and other 12259 utilities in this volume of IEEE Std 1003.1-2001. This does not mandate that the file system itself 12260 12261 be based on 512-byte blocks. The -k option was added as a compromise measure. It was agreed by the standard developers that 512 bytes was the best default unit because of its complete 12262 12263 historical consistency on System V (versus the mixed 512/1024-byte usage on BSD systems), and that a -k option to switch to 1024-byte units was a good compromise. Users who prefer the 12264 more logical 1024-byte quantity can easily alias df to $df - \mathbf{k}$ without breaking many historical 12265 scripts relying on the 512-byte units. 12266

It was suggested that *df* and the various related utilities be modified to access a *BLOCKSIZE* environment variable to achieve consistency and user acceptance. Since this is not historical practice on any system, it is left as a possible area for system extensions and will be re-evaluated

12267 12268

12269

df Utilities

in a future version if it is widely implemented.

12271 **FUTURE DIRECTIONS**

12272 None.

12273 SEE ALSO

12274 *find*

12275 CHANGE HISTORY

First released in Issue 2.

12277 **Issue 6**

12278 This utility is marked as part of the User Portability Utilities option.

Utilities diff

12279 NAME							
12280	diff — comp	diff — compare two files					
	12281 SYNOPSIS						
12282	diff $[-c \mid -e \mid -f \mid -C n]$ [-br] file1 file2						
12283 DESCR 12284	B DESCRIPTION The difficultity shall compare the contents of file1 and file2 and swrite to standard output a list of						
12285	The <i>diff</i> utility shall compare the contents of <i>file1</i> and <i>file2</i> and write to standard output a list of changes necessary to convert <i>file1</i> into <i>file2</i> . This list should be minimal. No output shall be						
12286		produced if the files are identical.					
12287 OPTIONS							
12288 12289		The <i>diff</i> utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.					
12290	The following	The following options shall be supported:					
12291	- b	Cause any amount of white space at the end of a line to be treated as a single					
12292		<newline> (that is, the white-space characters preceding the <newline> are ignored) and other strings of white space characters pret including convolines state.</newline></newline>					
12293 12294		ignored) and other strings of white-space characters, not including <newline>s, to compare equal.</newline>					
12295	-c	Produce output in a form that provides three lines of context.					
12296 12297	–C n	Produce output in a form that provides n lines of context (where n shall be interpreted as a positive decimal integer).					
12298 12299	−e	Produce output in a form suitable as input for the <i>ed</i> utility, which can then be used to convert <i>file1</i> into <i>file2</i> .					
12300 12301	−f	Produce output in an alternative form, similar in format to $-\mathbf{e}$, but not intended to be suitable as input for the ed utility, and in the opposite order.					
12302 12303	-r	Apply diff recursively to files and directories of the same name when file1 and file2 are both directories.					
12304 OPER							
12305	The following	ng operands shall be supported:					
12306 12307	file1, file2	A pathname of a file to be compared. If either the <i>file1</i> or <i>file2</i> operand is $'-'$, the standard input shall be used in its place.					
12308		If both file1 and file2 are directories, diff shall not compare block special files, character special					
12309	files, or FIFO special files to any files and shall not compare regular files to directories. Further						
12310 12311	details are as specified in Diff Directory Comparison Format (on page 320). The behavior of <i>diff</i> on other file types is implementation-defined when found in directories.						
12312		of <i>file1</i> and <i>file2</i> is a directory, <i>diff</i> shall be applied to the non-directory file and the file					
12313		contained in the directory file with a filename that is the same as the last component of the non-					
12314	directory file.						
12315 STDIN							
12316 12317		rd input shall be used only if one of the <i>file1</i> or <i>file2</i> operands references standard he INPUT FILES section.					
16011	ութաւ թեժ և	mpac dec die 1 11 O 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1					

12318 INPUT FILES

12319

The input files may be of any type.

diff Utilities

12320 ENVIRONMENT VARIABLES					
12321	The following environment variables shall affect the execution of <i>diff</i> :				
12322 12323 12324 12325	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)			
12326 12327	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.			
12328 12329 12330	LC_CTYPE	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).			
12331	LC_MESSAGES				
12332 12333 12334		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.			
12335 12336	LC_TIME	Determine the locale for affecting the format of file timestamps written with the $-C$ and $-c$ options.			
12337 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.			
12338 12339 12340	TZ	Determine the timezone used for calculating file timestamps written with the $-C$ and $-c$ options. If TZ is unset or null, an unspecified default timezone shall be used.			

12341 ASYNCHRONOUS EVENTS

40000 ENIX/IDONIN/ENIT VADIADI EC

12342 Default.

12343 STDOUT

Diff Directory Comparison Format 12344 If both *file1* and *file2* are directories, the following output formats shall be used. 12345 12346 In the POSIX locale, each file that is present in only one directory shall be reported using the 12347 following format: 12348 "Only in %s: %s\n", <directory pathname>, <filename> In the POSIX locale, subdirectories that are common to the two directories may be reported with 12349 12350 the following format: 12351 "Common subdirectories: %s and %s\n", <directory1 pathname>, 12352 <directory2 pathname> For each file common to the two directories if the two files are not to be compared, the following 12353 12354 format shall be used in the POSIX locale: "File %s is a %s while file %s is a %s\n", <directory1 pathname>, 12355 12356 <file type of directory1 pathname>, <directory2 pathname>, <file type of directory2 pathname> 12357 For each file common to the two directories, if the files are compared and are identical, no output 12358 shall be written. If the two files differ, the following format is written: 12359 12360 "diff %s %s %s\n", <diff options>, <filename1>, <filename2>

Utilities diff

12361 where *<diff_options>* are the options as specified on the command line. All directory pathnames listed in this section shall be relative to the original command line 12362 12363 arguments. All other names of files listed in this section shall be filenames (pathname components). 12364 **Diff Binary Output Format** 12365 In the POSIX locale, if one or both of the files being compared are not text files, an unspecified 12366 format shall be used that contains the pathnames of two files being compared and the string 12367 "differ". 12368 If both files being compared are text files, depending on the options specified, one of the 12369 following formats shall be used to write the differences. 12370 **Diff Default Output Format** 12371 12372 The default (without $-\mathbf{e}$, $-\mathbf{f}$, $-\mathbf{c}$, or $-\mathbf{C}$ options) diff utility output shall contain lines of these forms: 12373 "%da%d\n", <num1>, <num2> 12374 "%da%d,%d\n", <num1>, <num2>, <num3> 12375 "%dd%d\n", <num1>, <num2> 12376 "%d,%dd%d\n", <num1>, <num2>, <num3> 12377 "%dc%d\n", <num1>, <num2> 12378 12379 "%d,%dc%d\n", <num1>, <num2>, <num3> "%dc%d,%d\n", <num1>, <num2>, <num3> 12380 "%d,%dc%d,%d\n", <num1>, <num2>, <num3>, <num4> 12381 These lines resemble ed subcommands to convert file1 into file2. The line numbers before the 12382 action letters shall pertain to *file1*; those after shall pertain to *file2*. Thus, by exchanging a for d 12383 12384 and reading the line in reverse order, one can also determine how to convert *file2* into *file1*. As in *ed*, identical pairs (where *num1*= *num2*) are abbreviated as a single number. 12385 Following each of these lines, diff shall write to standard output all lines affected in the first file 12386 12387 using the format: 12388 $"<\Delta$ %s", <line> and all lines affected in the second file using the format: 12389 "> Δ %s", <line> 12390

If there are lines affected in both file1 and file2 (as with the c subcommand), the changes are

separated with a line consisting of three hyphens:

12391

12392

12393

"---\n"

diff Utilities

Diff –e Output Format

12394

12395 12396

12397 12398

12399

12400

12401

12402

12403

12404

12405

12406

12407 12408

12409

12410

12411 12412

12413 12414

12415

12418 12419

12420

12421

12426

With the –e option, a script shall be produced that shall, when provided as input to *ed*, along with an appended w (write) command, convert *file1* into *file2*. Only the a (append), c (change), d (delete), i (insert), and s (substitute) commands of *ed* shall be used in this script. Text lines, except those consisting of the single character period (' . '), shall be output as they appear in the file.

Diff –f Output Format

With the $-\mathbf{f}$ option, an alternative format of script shall be produced. It is similar to that produced by $-\mathbf{e}$, with the following differences:

- 1. It is expressed in reverse sequence; the output of **–e** orders changes from the end of the file to the beginning; the **–f** from beginning to end.
- 2. The command form < lines > < command-letter > used by -e is reversed. For example, 10c with -e would be c10 with -f.
- 3. The form used for ranges of line numbers is <space>-separated, rather than commaseparated.

Diff –c or –C Output Format

With the -c or -C option, the output format shall consist of affected lines along with surrounding lines of context. The affected lines shall show which ones need to be deleted or changed in *file1*, and those added from *file2*. With the -c option, three lines of context, if available, shall be written before and after the affected lines. With the -C option, the user can specify how many lines of context are written. The exact format follows.

The name and last modification time of each file shall be output in the following format:

```
12416 "*** %s %s\n", file1, <file1 timestamp>
12417 "--- %s %s\n", file2, <file2 timestamp>
```

Each < file > field shall be the pathname of the corresponding file being compared. The pathname written for standard input is unspecified.

In the POSIX locale, each *<timestamp>* field shall be equivalent to the output from the following command:

```
12422 date "+%a %b %e %T %Y"
```

without the trailing <newline>, executed at the time of last modification of the corresponding file (or the current time, if the file is standard input).

12425 Then, the following output formats shall be applied for every set of changes.

First, a line shall be written in the following format:

```
12427 "**********\n"
```

Next, the range of lines in *file1* shall be written in the following format if the range contains two or more lines:

```
12430 "*** %d,%d ****\n", <beginning line number>, <ending line number>
```

and the following format otherwise:

```
12432 "*** %d ****\n", <ending line number>
```

Utilities diff

12433 The ending line number of an empty range shall be the number of the preceding line, or 0 if the range is at the start of the file. 12434 12435 Next, the affected lines along with lines of context (unaffected lines) shall be written. Unaffected lines shall be written in the following format: 12436 12437 " $\Delta\Delta$ %s", <unaffected line> Deleted lines shall be written as: 12438 "- Δ %s", <deleted line> 12439 Changed lines shall be written as: 12440 "! Δ %s", < changed_line> 12441 Next, the range of lines in *file2* shall be written in the following format if the range contains two 12442 or more lines: 12443 12444 "--- %d,%d ----\n", <beginning line number>, <ending line number> and the following format otherwise: 12445 "--- %d ----\n", <ending line number> 12446 Then, lines of context and changed lines shall be written as described in the previous formats. 12447 Lines added from *file2* shall be written in the following format: 12448 "+ Δ %s", <added_line> 12449 12450 STDERR The standard error shall be used only for diagnostic messages. 12451 12452 OUTPUT FILES 12453 None 12454 EXTENDED DESCRIPTION 12455 None. 12456 EXIT STATUS 12457 The following exit values shall be returned: No differences were found. 12458 12459 Differences were found. 12460 >1 An error occurred. 12461 CONSEQUENCES OF ERRORS Default. 12462 12463 APPLICATION USAGE If lines at the end of a file are changed and other lines are added, diff output may show this as a 12464 12465 delete and add, as a change, or as a change and add; diff is not expected to know which happened and users should not care about the difference in output as long as it clearly shows the 12466 differences between the files. 12467 12468 EXAMPLES If dir1 is a directory containing a directory named x, dir2 is a directory containing a directory 12469 12470 named x, dir1/x and dir2/x both contain files named date.out, and dir2/x contains a file named y, the command: 12471

diff -r dir1 dir2

12472

diff Utilities

```
12473 could produce output similar to:

12474 Common subdirectories: dir1/x and dir2/x
12475 Only in dir2/x: y
12476 diff -r dir1/x/date.out dir2/x/date.out
12477 1c1
12478 < Mon Jul 2 13:12:16 PDT 1990
12479 ---
12480 > Tue Jun 19 21:41:39 PDT 1990
```

12481 RATIONALE

 The **-h** option was omitted because it was insufficiently specified and does not add to applications portability.

Historical implementations employ algorithms that do not always produce a minimum list of differences; the current language about making every effort is the best this volume of IEEE Std 1003.1-2001 can do, as there is no metric that could be employed to judge the quality of implementations against any and all file contents. The statement "This list should be minimal" clearly implies that implementations are not expected to provide the following output when comparing two 100-line files that differ in only one character on a single line:

```
1,100c1,100
all 100 lines from file1 preceded with "< "
---
all 100 lines from file2 preceded with "> "
```

The "Only in" messages required when the —r option is specified are not used by most historical implementations if the —e option is also specified. It is required here because it provides useful information that must be provided to update a target directory hierarchy to match a source hierarchy. The "Common subdirectories" messages are written by System V and 4.3 BSD when the —r option is specified. They are allowed here but are not required because they are reporting on something that is the same, not reporting a difference, and are not needed to update a target hierarchy.

The -c option, which writes output in a format using lines of context, has been included. The format is useful for a variety of reasons, among them being much improved readability and the ability to understand difference changes when the target file has line numbers that differ from another similar, but slightly different, copy. The patch utility is most valuable when working with difference listings using the context format. The BSD version of -c takes an optional argument specifying the amount of context. Rather than overloading -c and breaking the Utility Syntax Guidelines for diff, the standard developers decided to add a separate option for specifying a context diff with a specified amount of context (-C). Also, the format for context diffs was extended slightly in 4.3 BSD to allow multiple changes that are within context lines from each other to be merged together. The output format contains an additional four asterisks after the range of affected lines in the first filename. This was to provide a flag for old programs (like old versions of *patch*) that only understand the old context format. The version of context described here does not require that multiple changes within context lines be merged, but it does not prohibit it either. The extension is upwards-compatible, so any vendors that wish to retain the old version of diff can do so by adding the extra four asterisks (that is, utilities that currently use diff and understand the new merged format will also understand the old unmerged format, but not vice versa).

The substitute command was added as an additional format for the **–e** option. This was added to provide implementations with a way to fix the classic "dot alone on a line" bug present in many versions of *diff.* Since many implementations have fixed this bug, the standard developers decided not to standardize broken behavior, but rather to provide the necessary tool for fixing

Utilities diff

12522 the bug. One way to fix this bug is to output two periods whenever a lone period is needed, then 12523 terminate the append command with a period, and then use the substitute command to convert 12524 the two periods into one period. The BSD-derived –r option was added to provide a mechanism for using diff to compare two file 12525 12526 system trees. This behavior is useful, is standard practice on all BSD-derived systems, and is not easily reproducible with the *find* utility. 12527 The requirement that diff not compare files in some circumstances, even though they have the 12528 same name, is based on the actual output of historical implementations. The message specified 12529 12530 here is already in use when a directory is being compared to a non-directory. It is extended here 12531 to preclude the problems arising from running into FIFOs and other files that would cause diff to hang waiting for input with no indication to the user that diff was hung. In most common usage, 12532 diff -r should indicate differences in the file hierarchies, not the difference of contents of devices 12533 12534 pointed to by the hierarchies. Many early implementations of *diff* require seekable files. Since the System Interfaces volume of 12535 12536 IEEE Std 1003.1-2001 supports named pipes, the standard developers decided that such a 12537 restriction was unreasonable. Note also that the allowed filename – almost always refers to a 12538 pipe. No directory search order is specified for diff. The historical ordering is, in fact, not optimal, in 12539 that it prints out all of the differences at the current level, including the statements about all 12540 12541 common subdirectories before recursing into those subdirectories. 12542 The message: 12543 "diff %s %s %sn", <diff options>, <filename1>, <filename2>does not vary by locale because it is the representation of a command, not an English sentence. 12544 12545 FUTURE DIRECTIONS None. 12546 12547 SEE ALSO 12548 cmp, comm, ed, find 12549 CHANGE HISTORY First released in Issue 2. 12550 12551 Issue 5 The FUTURE DIRECTIONS section is added. 12552 12553 **Issue 6** The following new requirements on POSIX implementations derive from alignment with the 12554 Single UNIX Specification: 12555 12556 The –f option is added. The output format for -c or -c format is changed to align with changes to the IEEE P1003.2b 12557 12558 draft standard resulting from IEEE PASC Interpretation 1003.2 #71. The normative text is reworded to avoid use of the term "must" for application requirements. 12559 IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/20 is applied, changing the STDOUT 12560

section. This changes the specification of diff -c so that it agrees with existing practice when

contexts contain zero lines or one line.

12561

12562

dirname Utilities

12563 **NAME** dirname — return the directory portion of a pathname 12564 12565 SYNOPSIS 12566 dirname string 12567 DESCRIPTION The string operand shall be treated as a pathname, as defined in the Base Definitions volume of 12568 IEEE Std 1003.1-2001, Section 3.266, Pathname. The string string shall be converted to the name 12569 of the directory containing the filename corresponding to the last pathname component in 12570 12571 *string*, performing actions equivalent to the following steps in order: 1. If *string* is //, skip steps 2 to 5. 12572 12573 2. If string consists entirely of slash characters, string shall be set to a single slash character. In 12574 this case, skip steps 3 to 8. If there are any trailing slash characters in *string*, they shall be removed. 12575 If there are no slash characters remaining in string, string shall be set to a single period 12576 character. In this case, skip steps 5 to 8. 12577 12578 If there are any trailing non-slash characters in *string*, they shall be removed. If the remaining string is //, it is implementation-defined whether steps 7 and 8 are skipped 12579 12580 or processed. 12581 If there are any trailing slash characters in *string*, they shall be removed. 12582 If the remaining *string* is empty, *string* shall be set to a single slash character. The resulting string shall be written to standard output. 12583 12584 OPTIONS None. 12585 12586 **OPERANDS** The following operand shall be supported: 12587 string 12588 A string. 12589 **STDIN** Not used. 12590 12591 INPUT FILES None. 12592 12593 ENVIRONMENT VARIABLES 12594 The following environment variables shall affect the execution of *dirname*: LANG Provide a default value for the internationalization variables that are unset or null. 12595 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 12596 Internationalization Variables for the precedence of internationalization variables 12597 12598 used to determine the values of locale categories.)

12599

12600

12601

12602 12603 LC_ALL

LC_CTYPE

internationalization variables.

arguments).

If set to a non-empty string value, override the values of all the other

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

Utilities dirname

12604 LC_MESSAGES

Determine the locale that should be used to affect the format and contents of

12606 diagnostic messages written to standard error.

12607 XSI NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

12608 ASYNCHRONOUS EVENTS

12609 Default.

12610 STDOUT

12611 The *dirname* utility shall write a line to the standard output in the following format:

12612 "%s\n", <resulting string>

12613 STDERR

12614 The standard error shall be used only for diagnostic messages.

12615 OUTPUT FILES

12616 None.

12617 EXTENDED DESCRIPTION

12618 None.

12619 EXIT STATUS

12620 The following exit values shall be returned:

12621 0 Successful completion.

12622 >0 An error occurred.

12623 CONSEQUENCES OF ERRORS

12624 Default.

12625 APPLICATION USAGE

The definition of *pathname* specifies implementation-defined behavior for pathnames starting with two slash characters. Therefore, applications shall not arbitrarily add slashes to the beginning of a pathname unless they can ensure that there are more or less than two or are prepared to deal with the implementation-defined consequences.

12630 EXAMPLES

12626

12627 12628

12629

12631	Command	Results
12632	dirname /	/
12633	dirname //	/ or //
12634	dirname /a/b/	/a
12635	dirname //a//b//	//a
12636	dirname	Unspecified
12637	dirname a	. (\$? = 0)
12638	dirname ""	. (\$? = 0)
12639	dirname /a	/
12640	dirname /a/b	/a
12641	dirname a/b	a

12642 RATIONALE

The *dirname* utility originated in System III. It has evolved through the System V releases to a version that matches the requirements specified in this description in System V Release 3. 4.3 BSD and earlier versions did not include *dirname*.

The behaviors of *basename* and *dirname* in this volume of IEEE Std 1003.1-2001 have been coordinated so that when *string* is a valid pathname:

dirname Utilities

12648	<pre>\$(basename "string")</pre>
12649	would be a valid filename for the file in the directory:
12650	<pre>\$(dirname "string")</pre>
12651 12652 12653 12654 12655	This would not work for the versions of these utilities in early proposals due to the way processing of trailing slashes was specified. Consideration was given to leaving processing unspecified if there were trailing slashes, but this cannot be done; the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.266, Pathname allows trailing slashes. The <i>basename</i> and <i>dirname</i> utilities have to specify consistent handling for all valid pathnames.
12656 FUTUF 12657	RE DIRECTIONS None.
12658 SEE AI 12659	
12660 CHAN 12661	GE HISTORY First released in Issue 2.

12662 **NAME**

12663 du — estimate file space usage

12664 SYNOPSIS

12668

12669

12670 12671

12672

12673

12674

12675

12676

12677 12678

12665 UP du [-a | -s] [-kx] [-H | -L] [file ...]
12666

12667 DESCRIPTION

By default, the *du* utility shall write to standard output the size of the file space allocated to, and the size of the file space allocated to each subdirectory of, the file hierarchy rooted in each of the specified files. By default, when a symbolic link is encountered on the command line or in the file hierarchy, *du* shall count the size of the symbolic link (rather than the file referenced by the link), and shall not follow the link to another portion of the file hierarchy. The size of the file space allocated to a file of type directory shall be defined as the sum total of space allocated to all files in the file hierarchy rooted in the directory plus the space allocated to the directory itself.

When *du* cannot *stat*() files or *stat*() or read directories, it shall report an error condition and the final exit status is affected. Files with multiple links shall be counted and written for only one entry. The directory entry that is selected in the report is unspecified. By default, file sizes shall be written in 512-byte units, rounded up to the next 512-byte unit.

12679 OPTIONS

The *du* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

12682 The following options shall be supported:

In addition to the default output, report the size of each file not of type directory in the file hierarchy rooted in the specified file. Regardless of the presence of the **–a** option, non-directories given as *file* operands shall always be listed.

-H If a symbolic link is specified on the command line, du shall count the size of the file or file hierarchy referenced by the link.

12688 —**k** Write the files sizes in units of 1 024 bytes, rather than the default 512-byte units.

If a symbolic link is specified on the command line or encountered during the traversal of a file hierarchy, du shall count the size of the file or file hierarchy referenced by the link.

12692 —s Instead of the default output, report only the total sum for each of the specified files.

12694 -**x** When evaluating file sizes, evaluate only those files that have the same device as the file specified by the *file* operand.

Specifying more than one of the mutually-exclusive options –**H** and –**L** shall not be considered an error. The last option specified shall determine the behavior of the utility.

12698 OPERANDS

12699 The following operand shall be supported:

12700 *file* The pathname of a file whose size is to be written. If no *file* is specified, the current directory shall be used.

12702 **STDIN**

12696

12697

12703 Not used.

du Utilities

12704 INPUT FILES 12705 None. 12706 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *du*: 12707 LANG 12708 Provide a default value for the internationalization variables that are unset or null. 12709 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables 12710 used to determine the values of locale categories.) 12711 LC ALL If set to a non-empty string value, override the values of all the other 12712 12713 internationalization variables. Determine the locale for the interpretation of sequences of bytes of text data as 12714 LC_CTYPE characters (for example, single-byte as opposed to multi-byte characters in 12715 12716 arguments). LC MESSAGES 12717 12718 Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. 12719 12720 XSI **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 12721 ASYNCHRONOUS EVENTS 12722 Default. 12723 STDOUT The output from du shall consist of the amount of space allocated to a file and the name of the 12724 file, in the following format: 12725 "%d %s\n", <size>, <pathname> 12726 12727 STDERR The standard error shall be used only for diagnostic messages. 12728 12729 OUTPUT FILES None. 12730 12731 EXTENDED DESCRIPTION 12732 None. 12733 EXIT STATUS The following exit values shall be returned: 12734 12735 Successful completion. >0 An error occurred. 12736 12737 CONSEQUENCES OF ERRORS

12738

Default.

12739 APPLICATION USAGE

2740 None.

12741 EXAMPLES

12742 None.

12743 RATIONALE

The use of 512-byte units is historical practice and maintains compatibility with ls and other utilities in this volume of IEEE Std 1003.1-2001. This does not mandate that the file system itself be based on 512-byte blocks. The $-\mathbf{k}$ option was added as a compromise measure. It was agreed by the standard developers that 512 bytes was the best default unit because of its complete historical consistency on System V (*versus* the mixed 512/1024-byte usage on BSD systems), and that a $-\mathbf{k}$ option to switch to 1024-byte units was a good compromise. Users who prefer the 1024-byte quantity can easily alias du to du $-\mathbf{k}$ without breaking the many historical scripts relying on the 512-byte units.

The $-\mathbf{b}$ option was added to an early proposal to provide a resolution to the situation where System V and BSD systems give figures for file sizes in *blocks*, which is an implementation-defined concept. (In common usage, the block size is 512 bytes for System V and 1024 bytes for BSD systems.) However, $-\mathbf{b}$ was later deleted, since the default was eventually decided as 512-byte units.

Historical file systems provided no way to obtain exact figures for the space allocation given to files. There are two known areas of inaccuracies in historical file systems: cases of *indirect blocks* being used by the file system or *sparse* files yielding incorrectly high values. An indirect block is space used by the file system in the storage of the file, but that need not be counted in the space allocated to the file. A *sparse* file is one in which an *lseek()* call has been made to a position beyond the end of the file and data has subsequently been written at that point. A file system need not allocate all the intervening zero-filled blocks to such a file. It is up to the implementation to define exactly how accurate its methods are.

The **-a** and **-s** options were mutually-exclusive in the original version of *du*. The POSIX Shell and Utilities description is implied by the language in the SVID where **-s** is described as causing "only the grand total" to be reported. Some systems may produce output for **-sa**, but a Strictly Conforming POSIX Shell and Utilities Application cannot use that combination.

The $-\mathbf{a}$ and $-\mathbf{s}$ options were adopted from the SVID except that the System V behavior of not listing non-directories explicitly given as operands, unless the $-\mathbf{a}$ option is specified, was considered a bug; the BSD-based behavior (report for all operands) is mandated. The default behavior of du in the SVID with regard to reporting the failure to read files (it produces no messages) was considered counter-intuitive, and thus it was specified that the POSIX Shell and Utilities default behavior shall be to produce such messages. These messages can be turned off with shell redirection to achieve the System V behavior.

The $-\mathbf{x}$ option is historical practice on recent BSD systems. It has been adopted by this volume of IEEE Std 1003.1-2001 because there was no other historical method of limiting the du search to a single file hierarchy. This limitation of the search is necessary to make it possible to obtain file space usage information about a file system on which other file systems are mounted, without having to resort to a lengthy *find* and *awk* script.

12781 FUTURE DIRECTIONS

12782 None.

du Utilities

SO
ls, the System Interfaces volume of IEEE Std 1003.1-2001, stat()
GE HISTORY
First released in Issue 2.
This utility is marked as part of the User Portability Utilities option.
The APPLICATION USAGE section is added.
The obsolescent – r option has been removed.
The Open Group Corrigendum $U025/3$ is applied. The du utility is reinstated, as it had
incorrectly been marked LEGACY in Issue 5.
The -H and -L options for symbolic links are added as described in the IEEE P1003.2b draft
standard.

Utilities echo

12795 **NAME** 12796 echo — write arguments to standard output 12797 SYNOPSIS 12798 echo [string ...] 12799 **DESCRIPTION** The echo utility writes its arguments to standard output, followed by a <newline>. If there are 12800 no arguments, only the <newline> is written. 12801 12802 OPTIONS The *echo* utility shall not recognize the "--" argument in the manner specified by Guideline 10 12803 of the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines; 12804 "--" shall be recognized as a string operand. 12805 Implementations shall not support any options. 12806 12807 OPERANDS The following operands shall be supported: 12808 string A string to be written to standard output. If the first operand is $-\mathbf{n}$, or if any of the 12809 operands contain a backslash ('\') character, the results are implementation-12810 defined. 12811 12812 XSI On XSI-conformant systems, if the first operand is $-\mathbf{n}$, it shall be treated as a string, 12813 not an option. The following character sequences shall be recognized on XSIconformant systems within any of the arguments: 12814 Write an <alert>. 12815 ∖a \b Write a <backspace>. 12816 \c Suppress the <newline> that otherwise follows the final argument in the 12817 output. All characters following the '\c' in the arguments shall be 12818 ignored. 12819 \f Write a <form-feed>. 12820 Write a <newline>. 12821 \n \r Write a <carriage-return>. 12822 12823 \t Write a <tab>. Write a <vertical-tab>. 12824 \v // Write a backslash character. 12825 $\backslash 0num$ Write an 8-bit value that is the zero, one, two, or three-digit octal number 12826 12827 12828 12829 **STDIN** Not used. 12830 12831 INPUT FILES None. 12832 12833 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *echo*: 12834 LANG 12835 Provide a default value for the internationalization variables that are unset or null.

12836

the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,

echo Utilities

12837 12838		Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)				
12839 12840	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.				
12841 XSI 12842 12843	LC_CTYPE	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).				
12844 12845 12846	LC_MESSA	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.				
12847 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.				
12848 ASYN (12849	CHRONOUS Default.	EVENTS				
	12850 STDOUT					
12851 12852 XSI 12853	The <i>echo</i> utility arguments shall be separated by single <space>s and a <newline> shall follow the last argument. Output transformations shall occur based on the escape sequences in the input. See the OPERANDS section.</newline></space>					
12854 STDER						
12855		d error shall be used only for diagnostic messages.				
12856 OUTPU 12857	U T FILES None.					
12858 EXTENDED DESCRIPTION 12859 None.						
12860 EXIT S 12861		ng exit values shall be returned:				
12862	0 Success	ful completion.				
12863	>0 An erro	or occurred.				

12864 CONSEQUENCES OF ERRORS

12865 Default.

12869

12870

12871

12872

12873

12874

12866 APPLICATION USAGE

It is not possible to use *echo* portably across all POSIX systems unless both $-\mathbf{n}$ (as the first argument) and escape sequences are omitted.

The *printf* utility can be used portably to emulate any of the traditional behaviors of the *echo* utility as follows (assuming that *IFS* has its standard value or is unset):

 \bullet The historic System V *echo* and the requirements on XSI implementations in this volume of IEEE Std 1003.1-2001 are equivalent to:

```
printf "%b\n" "$*"
```

• The BSD *echo* is equivalent to:

```
12875 if [ "X$1" = "X-n" ]

12876 then

12877 shift

12878 printf "%s" "$*"

12879 else
```

Utilities echo

12880

12913

12914

printf "%s\n" "\$*" 12881 fi 12882 New applications are encouraged to use *printf* instead of *echo*. 12883 EXAMPLES 12884 None 12885 RATIONALE The echo utility has not been made obsolescent because of its extremely widespread use in 12886 historical applications. Conforming applications that wish to do prompting without <newline>s 12887 or that could possibly be expecting to echo a -n, should use the *printf* utility derived from the 12888 Ninth Edition system. 12889 As specified, echo writes its arguments in the simplest of ways. The two different historical 12890 12891 versions of *echo* vary in fatally incompatible ways. The BSD echo checks the first argument for the string -n which causes it to suppress the 12892 12893 <newline> that would otherwise follow the final argument in the output. The System V echo does not support any options, but allows escape sequences within its 12894 operands, as described for XSI implementations in the OPERANDS section. 12895 The echo utility does not support Utility Syntax Guideline 10 because historical applications 12896 depend on *echo* to echo *all* of its arguments, except for the **-n** option in the BSD version. 12897 12898 FUTURE DIRECTIONS 12899 None. 12900 SEE ALSO 12901 printf 12902 CHANGE HISTORY First released in Issue 2. 12903 12904 **Issue 5** In the OPTIONS section, the last sentence is changed to indicate that implementations "do not" 12905 12906 support any options; in the previous issue this said "need not". 12907 Issue 6 12908 The following new requirements on POSIX implementations derive from alignment with the 12909 Single UNIX Specification: 12910 A set of character sequences is defined as string operands. 12911 LC_CTYPE is added to the list of environment variables affecting echo. 12912 • In the OPTIONS section, implementations shall not support any options.

IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/21 is applied, so that the echo utility can

accommodate historical BSD behavior.

 \mathbf{ed} Utilities

12916	12915 NAME				
DESCRIPTION The of utility is a line-oriented text editor that uses two modes: command mode and input mode. In command mode the input characters shall be interpreted as commands, and in input mode they shall be interpreted as text. See the EXTENDED DESCRIPTION section.		ed — edit tex	xt		
The defutility is a line-oriented text editor that uses two modes: command mode and input mode they shall be interpreted as text. See the EXTENDED DESCRIPTION section.	12917 SYNOP	SIS			
The dt utility is a line-oriented text editor that uses two modes: command mode and input mode the command mode the input characters shall be interpreted as commands, and in input mode they shall be interpreted as text. See the EXTENDED DESCRIPTION section. 12923 OPTIONS 12924 The dutility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines. 12926 The following options shall be supported: 12927 — p string Use string as the prompt string when in command mode. By default, there shall be no prompt string. 12928 — s Suppress the writing of byte counts by e, E, r, and w commands and of the '!' prompt after a !command. 12930 The following operand shall be supported: 12931 If the file argument is given, ed shall simulate an e command on the file named by the pathname, file, before accepting commands from the standard input. If the file operand is '-', the results are unspecified. 12936 STDIN 12937 The standard input shall be a text file consisting of commands, as described in the EXTENDED DESCRIPTION section. 12939 INPUT FILES 12940 The input files shall be text files. 12941 ENVIRONMENT VARIABLES 12942 The following environment variables shall affect the execution of ed: 12943 LANG Provide a default value for the internationalization variables that are unset or null. See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables of the precedence of internationalization variables used to determine the values of locale categories.) 12944 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 12956 LC_COLLATE 12957 Determine the locale for the behavior of ranges, equivalence classes, and multi-character collating elements within regular expressions. 12958 LC_MESSACES Determine the locale for the interpretation of sequences of bytes of text data as characters (for example, single-byte as opposed to multi-byte char	12918	ed [-p st	ring] [-s] [file]		
In command mode the input characters shall be interpreted as commands, and in input mode they shall be interpreted as text. See the EXTENDED DESCRIPTION section. In the ed utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines. In the following options shall be supported: In posting options shall be supported: In prompt string. In prompt string of byte counts by e, E, r, and w commands and of the '!' prompt after a lcommand. In prompt after a lcomman					
12922 The year with the standard input shall be a text file consisting of commands, as described in the EXTENDED DESCRIPTION section. 12.2, 12925 The following options shall be supported: 12926					
The following options shall be supported: 12926					
The ed utility Syntax Guidelines. The following options shall be supported: -p string Use string as the prompt string when in command mode. By default, there shall be no prompt string. -s Suppress the writing of byte counts by e, E, r, and w commands and of the '1' prompt after a lcommand. 12929 -s Suppress the writing of byte counts by e, E, r, and w commands and of the '1' prompt after a lcommand. 12931 OPERANDS 12932		-	r		
The following options shall be supported: - p string Use string as the prompt string when in command mode. By default, there shall be no prompt string. - s Suppress the writing of byte counts by e, E, r, and w commands and of the '!' prompt after a !command. 2931 OPERANDS The following operand shall be supported: 12932 The following operand shall be supported: 12933 file If the file argument is given, ed shall simulate an e command on the file named by the pathname, file, before accepting commands from the standard input. If the file operand is '-', the results are unspecified. 12936 STDIN 12936 STDIN 12937 The standard input shall be a text file consisting of commands, as described in the EXTENDED DESCRIPTION section. 12939 INPUT FILES 12940 The input files shall be text files. 12941 ENVIRONMENT VARIABLES 12942 The following environment variables shall affect the execution of ed: 12944 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 12946 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables internationalization variables 12947 Determine the locale for the behavior of ranges, equivalence classes, and multi-character collating elements within regular expressions. 12953 LC_CTYPE 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) and the behavior of character classes within regular expressions.			y shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2,		
12927	12925	Utility Synta	x Guidelines.		
12928	12926	The followin	g options shall be supported:		
Suppress the writing of byte counts by e, E, r, and w commands and of the '!' prompt after a !command. 12931 OPERANDS 12932 The following operand shall be supported: 12933 file	12927	– p string	Use <i>string</i> as the prompt string when in command mode. By default, there shall be		
12931 OPERANDS 12932 The following operand shall be supported: 12933 file If the file argument is given, ed shall simulate an e command on the file named by the pathname, file, before accepting commands from the standard input. If the file 12934 the pathname, file, before accepting commands from the standard input. If the file 12935 operand is '-', the results are unspecified. 12936 STDIN 12937 The standard input shall be a text file consisting of commands, as described in the EXTENDED 12938 DESCRIPTION section. 12939 INPUT FILES 12940 The input files shall be text files. 12941 ENVIRONMENT VARIABLES 12942 The following environment variables shall affect the execution of ed: 12943 HOME Determine the pathname of the user's home directory. 12944 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 11 Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 12948 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 12949 LC_COLLATE 12950 LC_COLLATE 12951 Determine the locale for the behavior of ranges, equivalence classes, and multi-character collating elements within regular expressions. 12953 LC_CTYPE 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) and the behavior of character classes within regular expressions.	12928		no prompt string.		
The following operand shall be supported: 12933 file	12929	-s			
The following operand shall be supported: 12933	12930		prompt after a !command.		
file If the file argument is given, ed shall simulate an e command on the file named by the pathname, file, before accepting commands from the standard input. If the file operand is '-', the results are unspecified. 12936 STDIN 12937 The standard input shall be a text file consisting of commands, as described in the EXTENDED DESCRIPTION section. 12939 INPUT FILES 12940 The input files shall be text files. 12941 ENVIRONMENT VARIABLES 12942 The following environment variables shall affect the execution of ed: 12943 HOME Determine the pathname of the user's home directory. 12944 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 12948 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 12950 LC_COLLATE 12951 Determine the locale for the behavior of ranges, equivalence classes, and multi-character collating elements within regular expressions. 12953 LC_CTYPE 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) and the behavior of character classes within regular expressions. 12957 LC_MESSAGES			or an arrand also II has arrand and also		
the pathname, file, before accepting commands from the standard input. If the file operand is '-', the results are unspecified. 12936 STDIN 12937 The standard input shall be a text file consisting of commands, as described in the EXTENDED DESCRIPTION section. 12938 INPUT FILES 12940 The input files shall be text files. 12941 ENVIRONMENT VARIABLES 12942 The following environment variables shall affect the execution of ed: 12943 HOME Determine the pathname of the user's home directory. 12944 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 12948 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 12950 LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multicharacter collating elements within regular expressions. 12953 LC_CTYPE 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) and the behavior of character classes within regular expressions.			· · · · · · · · · · · · · · · · · · ·		
12936 STDIN 12937 The standard input shall be a text file consisting of commands, as described in the EXTENDED DESCRIPTION section. 12938 INPUT FILES 12940 The input files shall be text files. 12941 ENVIRONMENT VARIABLES 12942 The following environment variables shall affect the execution of ed: 12943 HOME Determine the pathname of the user's home directory. 12944 LANG Provide a default value for the internationalization variables that are unset or null. 12945 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 12948 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 12950 LC_COLLATE 12951 Determine the locale for the behavior of ranges, equivalence classes, and multicharacter collating elements within regular expressions. 12953 LC_CTYPE 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) and the behavior of character classes within regular expressions.		file			
The standard input shall be a text file consisting of commands, as described in the EXTENDED DESCRIPTION section. 12939 INPUT FILES 12940 The input files shall be text files. 12941 ENVIRONMENT VARIABLES 12942 The following environment variables shall affect the execution of ed: 12943 HOME Determine the pathname of the user's home directory. 12944 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 12948 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 12950 LC_COLLATE 12951 Determine the locale for the behavior of ranges, equivalence classes, and multi-character collating elements within regular expressions. 12953 LC_CTYPE 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) and the behavior of character classes within regular expressions. 12957 LC_MESSAGES					
The standard input shall be a text file consisting of commands, as described in the EXTENDED DESCRIPTION section. 12939 INPUT FILES 12940 The input files shall be text files. 12941 ENVIRONMENT VARIABLES 12942 The following environment variables shall affect the execution of ed: 12943 HOME Determine the pathname of the user's home directory. 12944 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 12948 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 12950 LC_COLLATE 12951 Determine the locale for the behavior of ranges, equivalence classes, and multicharacter collating elements within regular expressions. 12953 LC_CTYPE 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) and the behavior of character classes within regular expressions. 12957 LC_MESSAGES	12936 STDIN				
The input files shall be text files. 12941 ENVIRONMENT VARIABLES 12942 The following environment variables shall affect the execution of ed: 12943 HOME Determine the pathname of the user's home directory. 12944 LANG Provide a default value for the internationalization variables that are unset or null. 12945 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 12946 Internationalization Variables for the precedence of internationalization variables 12947 used to determine the values of locale categories.) 12948 LC_ALL If set to a non-empty string value, override the values of all the other 12949 internationalization variables. 12950 LC_COLLATE 12951 Determine the locale for the behavior of ranges, equivalence classes, and multi- 12952 character collating elements within regular expressions. 12953 LC_CTYPE 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) and the behavior of character classes within regular expressions. 12957 LC_MESSAGES		The standard	d input shall be a text file consisting of commands, as described in the EXTENDED		
The input files shall be text files. 12941 ENVIRONMENT VARIABLES 12942 The following environment variables shall affect the execution of ed: 12943 HOME Determine the pathname of the user's home directory. 12944 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 12948 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 12950 LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multi-character collating elements within regular expressions. 12953 LC_CTYPE 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) and the behavior of character classes within regular expressions. 12957 LC_MESSAGES LC_M	12938	DESCRIPTION	ON section.		
The following environment variables shall affect the execution of ed: 12943 HOME Determine the pathname of the user's home directory. 12944 LANG Provide a default value for the internationalization variables that are unset or null. 12945 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 12946 Internationalization Variables for the precedence of internationalization variables 12947 used to determine the values of locale categories.) 12948 LC_ALL If set to a non-empty string value, override the values of all the other 12949 internationalization variables. 12950 LC_COLLATE 12951 Determine the locale for the behavior of ranges, equivalence classes, and multi- 12952 character collating elements within regular expressions. 12953 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 12954 characters (for example, single-byte as opposed to multi-byte characters in 12955 arguments and input files) and the behavior of character classes within regular 12957 LC_MESSAGES	12939 INPUT				
The following environment variables shall affect the execution of ed: HOME Determine the pathname of the user's home directory. LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multi-character collating elements within regular expressions. LC_CTYPE 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) and the behavior of character classes within regular expressions. LC_MESSAGES	12940	The input fil	es shall be text files.		
HOME Determine the pathname of the user's home directory. LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multi-character collating elements within regular expressions. LC_CTYPE 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) and the behavior of character classes within regular expressions. LC_MESSAGES					
LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multi- character collating elements within regular expressions. LC_CTYPE 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) and the behavior of character classes within regular expressions. LC_MESSAGES	12942				
(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 12948	12943		•		
Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multi-character collating elements within regular expressions. LC_CTYPE 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) and the behavior of character classes within regular expressions. LC_MESSAGES		LANG			
used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multi-character collating elements within regular expressions. LC_CTYPE 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) and the behavior of character classes within regular expressions. LC_MESSAGES					
internationalization variables. LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multi- character collating elements within regular expressions. LC_CTYPE 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) and the behavior of character classes within regular expressions. LC_MESSAGES					
internationalization variables. LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multi- character collating elements within regular expressions. LC_CTYPE 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) and the behavior of character classes within regular expressions. LC_MESSAGES	12948	LC ALL	If set to a non-empty string value, override the values of all the other		
Determine the locale for the behavior of ranges, equivalence classes, and multi- character collating elements within regular expressions. LC_CTYPE 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) and the behavior of character classes within regular expressions. LC_MESSAGES			* 0		
character collating elements within regular expressions. LC_CTYPE 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) and the behavior of character classes within regular expressions. LC_MESSAGES	12950	LC_COLLAT	E		
12953 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 12954 characters (for example, single-byte as opposed to multi-byte characters in 12955 arguments and input files) and the behavior of character classes within regular 12956 expressions. 12957 LC_MESSAGES			ÿ .		
characters (for example, single-byte as opposed to multi-byte characters in arguments and input files) and the behavior of character classes within regular expressions. LC_MESSAGES	12952				
arguments and input files) and the behavior of character classes within regular expressions. LC_MESSAGES		LC_CTYPE			
12956 expressions. 12957 <i>LC_MESSAGES</i>					
	12957 LC MESSAGES				
	12958				

12959 12960		diagnostic messages written to standard error and informative messages written to standard output.
12961 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
12962 ASYNC 12963 12964	The <i>ed</i> utility	EVENTS y shall take the standard action for all signals (see the ASYNCHRONOUS EVENTS ction 1.11 (on page 20)) with the following exceptions:
12965 12966 12967	SIGINT	The <i>ed</i> utility shall interrupt its current activity, write the string "?\n" to standard output, and return to command mode (see the EXTENDED DESCRIPTION section).
12968 12969 12970 12971 12972	SIGHUP	If the buffer is not empty and has changed since the last write, the <i>ed</i> utility shall attempt to write a copy of the buffer in a file. First, the file named ed.hup in the current directory shall be used; if that fails, the file named ed.hup in the directory named by the <i>HOME</i> environment variable shall be used. In any case, the <i>ed</i> utility shall exit without returning to command mode.
12973	SIGQUIT	The ed utility shall ignore this event.

12974 STDOUT

Various editing commands and the prompting feature (see **-p**) write to standard output, as described in the EXTENDED DESCRIPTION section.

12977 STDERR

12984

12985

12986

12987

12988

12989 12990

12991

12992

12993

12994

12978 The standard error shall be used only for diagnostic messages.

12979 OUTPUT FILES

12980 The output files shall be text files whose formats are dependent on the editing commands given.

12981 EXTENDED DESCRIPTION

The *ed* utility shall operate on a copy of the file it is editing; changes made to the copy shall have no effect on the file until a **w** (write) command is given. The copy of the text is called the *buffer*.

Commands to *ed* have a simple and regular structure: zero, one, or two *addresses* followed by a single-character *command*, possibly followed by parameters to that command. These addresses specify one or more lines in the buffer. Every command that requires addresses has default addresses, so that the addresses very often can be omitted. If the **–p** option is specified, the prompt string shall be written to standard output before each command is read.

In general, only one command can appear on a line. Certain commands allow text to be input. This text is placed in the appropriate place in the buffer. While *ed* is accepting text, it is said to be in *input mode*. In this mode, no commands shall be recognized; all input is merely collected. Input mode is terminated by entering a line consisting of two characters: a period ('.') followed by a <newline>. This line is not considered part of the input text.

Regular Expressions in ed

The *ed* utility shall support basic regular expressions, as described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.3, Basic Regular Expressions. Since regular expressions in *ed* are always matched against single lines (excluding the terminating <newline>s), never against any larger section of text, there is no way for a regular expression to match a <newline>.

13000 A null RE shall be equivalent to the last RE encountered.

13001 Regular expressions are used in addresses to specify lines, and in some commands (for example, 13002 the **s** substitute command) to specify portions of a line to be substituted.

ed Utilities

Addresses in ed

Addressing in *ed* relates to the current line. Generally, the current line is the last line affected by a command. The current line number is the address of the current line. If the edit buffer is not empty, the initial value for the current line shall be the last line in the edit buffer; otherwise, zero.

Addresses shall be constructed as follows:

- 1. The period character (' . ') shall address the current line.
- 2. The dollar sign character $(' \ \sharp')$ shall address the last line of the edit buffer.
- 3. The positive decimal number *n* shall address the *n*th line of the edit buffer.
- 4. The apostrophe-x character pair ("'x") shall address the line marked with the mark name character x, which shall be a lowercase letter from the portable character set. It shall be an error if the character has not been set to mark a line or if the line that was marked is not currently present in the edit buffer.
- 5. A BRE enclosed by slash characters ('/') shall address the first line found by searching forwards from the line following the current line toward the end of the edit buffer and stopping at the first line for which the line excluding the terminating <newline> matches the BRE. The BRE consisting of a null BRE delimited by a pair of slash characters shall address the next line for which the line excluding the terminating <newline> matches the last BRE encountered. In addition, the second slash can be omitted at the end of a command line. Within the BRE, a backslash-slash pair ("\/") shall represent a literal slash instead of the BRE delimiter. If necessary, the search shall wrap around to the beginning of the buffer and continue up to and including the current line, so that the entire buffer is searched.
- 6. A BRE enclosed by question-mark characters ('?') shall address the first line found by searching backwards from the line preceding the current line toward the beginning of the edit buffer and stopping at the first line for which the line excluding the terminating <newline> matches the BRE. The BRE consisting of a null BRE delimited by a pair of question-mark characters ("??") shall address the previous line for which the line excluding the terminating <newline> matches the last BRE encountered. In addition, the second question-mark can be omitted at the end of a command line. Within the BRE, a backslash-question-mark pair ("\?") shall represent a literal question mark instead of the BRE delimiter. If necessary, the search shall wrap around to the end of the buffer and continue up to and including the current line, so that the entire buffer is searched.
- 7. A plus-sign ('+') or hyphen character ('-') followed by a decimal number shall address the current line plus or minus the number. A plus-sign or hyphen character not followed by a decimal number shall address the current line plus or minus 1.

Addresses can be followed by zero or more address offsets, optionally

 dlank>-separated. Address offsets are constructed as follows:

- A plus-sign or hyphen character followed by a decimal number shall add or subtract, respectively, the indicated number of lines to or from the address. A plus-sign or hyphen character not followed by a decimal number shall add or subtract 1 to or from the address.
- A decimal number shall add the indicated number of lines to the address.

It shall not be an error for an intermediate address value to be less than zero or greater than the last line in the edit buffer. It shall be an error for the final address value to be less than zero or greater than the last line in the edit buffer. It shall be an error if a search for a BRE fails to find a matching line.

Commands accept zero, one, or two addresses. If more than the required number of addresses are provided to a command that requires zero addresses, it shall be an error. Otherwise, if more than the required number of addresses are provided to a command, the addresses specified first shall be evaluated and then discarded until the maximum number of valid addresses remain, for the specified command.

Addresses shall be separated from each other by a comma (', ') or semicolon character (', '). In the case of a semicolon separator, the current line (', ') shall be set to the first address, and only then will the second address be calculated. This feature can be used to determine the starting line for forwards and backwards searches; see rules 5. and 6.

Addresses can be omitted on either side of the comma or semicolon separator, in which case the resulting address pairs shall be as follows:

Specified	Resulting
,	1 , \$
, addr	1 , addr
addr ,	addr , addr
;	. ; \$
; addr	. ; addr
addr ;	addr ; addr

Any

 sincluded between addresses, address separators, or address offsets shall be
 ignored.

Commands in ed

In the following list of *ed* commands, the default addresses are shown in parentheses. The number of addresses shown in the default shall be the number expected by the command. The parentheses are not part of the address; they show that the given addresses are the default.

Each address component can be preceded by zero or more <blank>s. The command letter can be preceded by zero or more <blank>s. If a suffix letter (l, n, or p) is given, the application shall ensure that it immediately follows the command.

The **e**, **E**, **f**, **r**, and **w** commands shall take an optional *file* parameter, separated from the command letter by one or more

blank>s.

If changes have been made in the buffer since the last \mathbf{w} command that wrote the entire buffer, ed shall warn the user if an attempt is made to destroy the editor buffer via the \mathbf{e} or \mathbf{q} commands. The ed utility shall write the string:

13089 "?\n"

(followed by an explanatory message if *help mode* has been enabled via the **H** command) to standard output and shall continue in command mode with the current line number unchanged.

If the **e** or **q** command is repeated with no intervening command, it shall take effect.

ed Utilities

If a terminal disconnect is detected:

• If the buffer is not empty and has changed since the last write, the *ed* utility shall attempt to write a copy of the buffer to a file named **ed.hup** in the current directory. If this write fails, *ed* shall attempt to write a copy of the buffer to a filename **ed.hup** in the directory named by the *HOME* environment variable. If both these attempts fail, *ed* shall exit without saving the buffer.

• The *ed* utility shall not write the file to the currently remembered pathname or return to command mode, and shall terminate with a non-zero exit status.

If an end-of-file is detected on standard input:

- If the *ed* utility is in input mode, *ed* shall terminate input mode and return to command mode. It is unspecified if any partially entered lines (that is, input text without a terminating <newline>) are discarded from the input text.
- If the *ed* utility is in command mode, it shall act as if a **q** command had been entered.

If the closing delimiter of an RE or of a replacement string (for example, '/') in a g, G, s, v, or V command would be the last character before a <newline>, that delimiter can be omitted, in which case the addressed line shall be written. For example, the following pairs of commands are equivalent:

```
s/s1/s2 s/s1/s2/p
g/s1 g/s1/p
?s1 ?s1?
```

If an invalid command is entered, *ed* shall write the string:

13114 "?\n"

 (followed by an explanatory message if *help mode* has been enabled via the **H** command) to standard output and shall continue in command mode with the current line number unchanged.

Append Command

```
Synopsis: (.)a
<text>
```

The a command shall read the given text and append it after the addressed line; the current line number shall become the address of the last inserted line or, if there were none, the addressed line. Address 0 shall be valid for this command; it shall cause the appended text to be placed at the beginning of the buffer.

Change Command

```
Synopsis: (.,.)c <text>
```

The c command shall delete the addressed lines, then accept input text that replaces these lines; the current line shall be set to the address of the last line input; or, if there were none, at the line after the last line deleted; if the lines deleted were originally at the end of the buffer, the current line number shall be set to the address of the new last line; if no lines remain in the buffer, the current line number shall be set to zero. Address 0 shall be valid for this command; it shall be interpreted as if address 1 were specified.

ne after the nally at the line; if no named by line of the e used (see less the -s
nally at the line; if no named by line of the used (see
line of the used (see
line of the used (see
line of the used (see
nent e, E, r, be a shell membered command. er shall be
pt that the the last w
e; whether membered
luding the ginning of line, with mmand list

13174

13175

value assigned by the last command in the command list. If there were no matching lines, the

current line number shall not be changed. A single command or the first of a list of commands

ed Utilities

shall appear on the same line as the global command. All lines of a multi-line list except the last line shall be ended with a backslash preceding the terminating <newline>; the **a**, **i**, and **c** commands and associated input are permitted. The '.' terminating input mode can be omitted if it would be the last line of the *command list*. An empty *command list* shall be equivalent to the **p** command. The use of the **g**, **G**, **v**, **V**, and ! commands in the *command list* produces undefined results. Any character other than <space> or <newline> can be used instead of a slash to delimit the RE. Within the RE, the RE delimiter itself can be used as a literal character if it is preceded by a backslash.

Interactive Global Command

Synopsis: (1,\$)G/RE/

In the G command, the first step shall be to mark every line for which the line excluding the terminating <newline> matches the given RE. Then, for every such line, that line shall be written, the current line number shall be set to the address of that line, and any one command (other than one of the a, c, i, g, G, v, and V commands) shall be read and executed. A <newline> shall act as a null command (causing no action to be taken on the current line); an '&' shall cause the re-execution of the most recent non-null command executed within the current invocation of G. Note that the commands input as part of the execution of the G command can address and affect any lines in the buffer. Any line modified by the command shall be unmarked. The final value of the current line number shall be the value set by the last command successfully executed. (Note that the last command successfully executed shall be the G command itself if a command fails or the null command is specified.) If there were no matching lines, the current line number shall not be changed. The G command can be terminated by a SIGINT signal. Any character other than <space> or <newline> can be used instead of a slash to delimit the RE and the replacement. Within the RE, the RE delimiter itself can be used as a literal character if it is preceded by a backslash.

Help Command

Synopsis: h

The **h** command shall write a short message to standard output that explains the reason for the most recent '?' notification. The current line number shall be unchanged.

Help-Mode Command

Synopsis: H

The **H** command shall cause *ed* to enter a mode in which help messages (see the **h** command) shall be written to standard output for all subsequent '?' notifications. The **H** command alternately shall turn this mode on and off; it is initially off. If the help-mode is being turned on, the **H** command also explains the previous '?' notification, if there was one. The current line number shall be unchanged.

Insert Command

Synopsis: (.)i <text>

The **i** command shall insert the given text before the addressed line; the current line is set to the last inserted line or, if there was none, to the addressed line. This command differs from the **a** command only in the placement of the input text. Address 0 shall be valid for this command; it shall be interpreted as if address 1 were specified.

13220 Join Command (.,.+1)j13221 Synopsis: The **j** command shall join contiguous lines by removing the appropriate <newline>s. If exactly 13222 13223 one address is given, this command shall do nothing. If lines are joined, the current line number 13224 shall be set to the address of the joined line; otherwise, the current line number shall be unchanged. 13225 **Mark Command** 13226 Synopsis: 13227 The k command shall mark the addressed line with name x, which the application shall ensure is 13228 a lowercase letter from the portable character set. The address "'x" shall then refer to this line; 13229 the current line number shall be unchanged. 13230 List Command 13231 13232 Synopsis: (.,.)1The I command shall write to standard output the addressed lines in a visually unambiguous 13233 form. The characters listed in the Base Definitions volume of IEEE Std 1003.1-2001, Table 5-1, 13234 Escape Sequences and Associated Actions ('\\', '\a', '\b', '\f', '\r', '\t', '\v') shall 13235 be written as the corresponding escape sequence; the '\n' in that table is not applicable. Non-13236 printable characters not in the table shall be written as one three-digit octal number (with a 13237 preceding backslash character) for each byte in the character (most significant byte first). If the 13238 size of a byte on the system is greater than nine bits, the format used for non-printable characters 13239 13240 is implementation-defined. Long lines shall be folded, with the point of folding indicated by <newline> preceded by a 13241 13242 backslash; the length at which folding occurs is unspecified, but should be appropriate for the 13243 output device. The end of each line shall be marked with a '\$', and '\$' characters within the 13244 text shall be written with a preceding backslash. An I command can be appended to any other command other than e, E, f, q, Q, r, w, or !. The current line number shall be set to the address of 13245 13246 the last line written. **Move Command** 13247 13248 Synopsis: (.,.) maddress 13249 The **m** command shall reposition the addressed lines after the line addressed by address. 13250 Address 0 shall be valid for address and cause the addressed lines to be moved to the beginning of the buffer. It shall be an error if address address falls within the range of moved lines. The 13251 current line number shall be set to the address of the last line moved. 13252 **Number Command** 13253 13254 Synopsis:

The **n** command shall write to standard output the addressed lines, preceding each line by its

line number and a <tab>; the current line number shall be set to the address of the last line

written. The n command can be appended to any command other than e, E, f, q, Q, r, w, or !.

13255

13256

13257

ed Utilities

13258 **Print Command** (.,.)p 13259 Synopsis: The **p** command shall write to standard output the addressed lines; the current line number shall 13260 13261 be set to the address of the last line written. The \mathbf{p} command can be appended to any command 13262 other than e, E, f, q, Q, r, w, or !. **Prompt Command** 13263 Synopsis: 13264 13265 The **P** command shall cause ed to prompt with an asterisk ('*') (or string, if $-\mathbf{p}$ is specified) for all subsequent commands. The P command alternatively shall turn this mode on and off; it shall 13266 be initially on if the $-\mathbf{p}$ option is specified; otherwise, off. The current line number shall be 13267 unchanged. 13268 **Quit Command** 13269 13270 Synopsis: The **q** command shall cause *ed* to exit. If the buffer has changed since the last time the entire 13271 buffer was written, the user shall be warned, as described previously. 13272 **Quit Without Checking Command** 13273 13274 Synopsis: The \mathbf{Q} command shall cause *ed* to exit without checking whether changes have been made in the 13275 13276 buffer since the last w command. **Read Command** 13277 Synopsis: (\$)r [file] 13278 13279 The r command shall read in the file named by the pathname file and append it after the addressed line. If no file argument is given, the currently remembered pathname, if any, shall be 13280 used (see the e and f commands). The currently remembered pathname shall not be changed 13281 unless there is no remembered pathname. Address 0 shall be valid for r and shall cause the file to 13282 be read at the beginning of the buffer. If the read is successful, and -s was not specified, the 13283 number of bytes read shall be written to standard output in the following format: 13284 13285 "%d\n", <number of bytes read> The current line number shall be set to the address of the last line read in. If *file* is replaced by 13286 '!', the rest of the line shall be taken to be a shell command line whose output is to be read. 13287 Such a shell command line shall not be remembered as the current pathname. 13288 **Substitute Command** 13289 Synopsis: (.,.)s/RE/replacement/flags 13290 13291 The s command shall search each addressed line for an occurrence of the specified RE and replace either the first or all (non-overlapped) matched strings with the replacement; see the 13292 following description of the g suffix. It is an error if the substitution fails on every addressed 13293 line. Any character other than <space> or <newline> can be used instead of a slash to delimit the 13294 RE and the replacement. Within the RE, the RE delimiter itself can be used as a literal character 13295

which a substitution occurred.

13296

13297

if it is preceded by a backslash. The current line shall be set to the address of the last line on

An ampersand ('&') appearing in the *replacement* shall be replaced by the string matching the RE on the current line. The special meaning of '&' in this context can be suppressed by preceding it by backslash. As a more general feature, the characters '\n', where n is a digit, shall be replaced by the text matched by the corresponding back-reference expression. When the character '%' is the only character in the *replacement*, the *replacement* used in the most recent substitute command shall be used as the *replacement* in the current substitute command; if there was no previous substitute command, the use of '%' in this manner shall be an error. The '%' shall lose its special meaning when it is in a replacement string of more than one character or is preceded by a backslash. For each backslash ('\') encountered in scanning *replacement* from beginning to end, the following character shall lose its special meaning (if any). It is unspecified what special meaning is given to any character other than '&', '\', '\', '\', or digits.

A line can be split by substituting a <newline> into it. The application shall ensure it escapes the <newline> in the *replacement* by preceding it by backslash. Such substitution cannot be done as part of a g or v *command list*. The current line number shall be set to the address of the last line on which a substitution is performed. If no substitution is performed, the current line number shall be unchanged. If a line is split, a substitution shall be considered to have been performed on each of the new lines for the purpose of determining the new current line number. A substitution shall be considered to have been performed even if the replacement string is identical to the string that it replaces.

The application shall ensure that the value of *flags* is zero or more of:

count Substitute for the *count*th occurrence only of the RE found on each addressed line.

- g Globally substitute for all non-overlapping instances of the RE rather than just the first one. If both g and *count* are specified, the results are unspecified.
- Write to standard output the final line in which a substitution was made. The line shall be written in the format specified for the I command.
- **n** Write to standard output the final line in which a substitution was made. The line shall be written in the format specified for the **n** command.
- **p** Write to standard output the final line in which a substitution was made. The line shall be written in the format specified for the **p** command.

Copy Command

13328 Synopsis: (.,.)taddress

The **t** command shall be equivalent to the **m** command, except that a copy of the addressed lines shall be placed after address (which can be 0); the current line number shall be set to the address of the last line added.

Undo Command

Synopsis: u

 The **u** command shall nullify the effect of the most recent command that modified anything in the buffer, namely the most recent **a**, **c**, **d**, **g**, **i**, **j**, **m**, **r**, **s**, **t**, **u**, **v**, **G**, or **V** command. All changes made to the buffer by a **g**, **G**, **v**, or **V** global command shall be undone as a single change; if no changes were made by the global command (such as with g/RE/p), the **u** command shall have no effect. The current line number shall be set to the value it had immediately before the command being undone started.

ed Utilities

Global Non-Matched Command

13341 Synopsis: (1,\$) v/RE/command list

This command shall be equivalent to the global command **g** except that the lines that are marked during the first step shall be those for which the line excluding the terminating <newline> does not match the RE.

Interactive Global Not-Matched Command

13346 Synopsis: (1, \$) V/RE/

13340

13345

13347

13348

13349

13350

13352

13353

13354

13355

13356

13357

13358

13360

13361

13362

13363 13364

13365

13366

13372

13374

13375

13376

13377

13378 13379

13380

This command shall be equivalent to the interactive global command **G** except that the lines that are marked during the first step shall be those for which the line excluding the terminating <newline> does not match the RE.

Write Command

13351 Synopsis: (1,\$) w [file]

The w command shall write the addressed lines into the file named by the pathname *file*. The command shall create the file, if it does not exist, or shall replace the contents of the existing file. The currently remembered pathname shall not be changed unless there is no remembered pathname. If no pathname is given, the currently remembered pathname, if any, shall be used (see the e and f commands); the current line number shall be unchanged. If the command is successful, the number of bytes written shall be written to standard output, unless the -s option was specified, in the following format:

13359 "%d\n", <number of bytes written>

If *file* begins with '!', the rest of the line shall be taken to be a shell command line whose standard input shall be the addressed lines. Such a shell command line shall not be remembered as the current pathname. This usage of the write command with '!' shall not be considered as a "last w command that wrote the entire buffer", as described previously; thus, this alone shall not prevent the warning to the user if an attempt is made to destroy the editor buffer via the e or q commands.

Line Number Command

13367 *Synopsis*: (\$) =

The line number of the addressed line shall be written to standard output in the following format:

13370 "%d\n", <line number>

13371 The current line number shall be unchanged by this command.

Shell Escape Command

13373 Synopsis: !command

The remainder of the line after the '!' shall be sent to the command interpreter to be interpreted as a shell command line. Within the text of that shell command line, the unescaped character '%' shall be replaced with the remembered pathname; if a '!' appears as the first character of the command, it shall be replaced with the text of the previous shell command executed via '!'. Thus, "!!" shall repeat the previous !command. If any replacements of '%' or '!' are performed, the modified line shall be written to the standard output before command is executed. The ! command shall write:

13381 "!\n" to standard output upon completion, unless the -s option is specified. The current line number 13382 13383 shall be unchanged. **Null Command** 13384 Synopsis: 13385 13386 An address alone on a line shall cause the addressed line to be written. A <newline> alone shall be equivalent to "+1p". The current line number shall be set to the address of the written line. 13387 13388 EXIT STATUS The following exit values shall be returned: 13389 Successful completion without any file or command errors. 13390 >0 An error occurred. 13391 13392 CONSEQUENCES OF ERRORS When an error in the input script is encountered, or when an error is detected that is a 13393 13394 consequence of the data (not) present in the file or due to an external condition such as a read or 13395 write error: • If the standard input is a terminal device file, all input shall be flushed, and a new command 13396 13397 13398 • If the standard input is a regular file, *ed* shall terminate with a non-zero exit status. 13399 APPLICATION USAGE 13400 Because of the extremely terse nature of the default error messages, the prudent script writer 13401 begins the ed input commands with an H command, so that if any errors do occur at least some 13402 clue as to the cause is made available. 13403 In previous versions, an obsolescent - option was described. This is no longer specified. Applications should use the -s option. Using - as a file operand now produces unspecified 13404 results. This allows implementations to continue to support the former required behavior. 13405 13406 EXAMPLES 13407 None. 13408 RATIONALE The initial description of this utility was adapted from the SVID. It contains some features not 13409 found in Version 7 or BSD-derived systems. Some of the differences between the POSIX and 13410 13411 BSD *ed* utilities include, but need not be limited to: • The BSD – option does not suppress the '!' prompt after a! command. 13412 13413

- BSD does not support the special meanings of the '%' and '!' characters within a ! command.
 - BSD does not support the addresses ';' and ','.

13414

13415

- BSD allows the command/suffix pairs **pp**, **ll**, and so on, which are unspecified in this volume of IEEE Std 1003.1-2001.
- BSD does not support the '!' character part of the e, r, or w commands.
- \bullet A failed **g** command in BSD sets the line number to the last line searched if there are no matches.

ed Utilities

- BSD does not default the *command list* to the \mathbf{p} command.
 - BSD does not support the **G**, **h**, **H**, **n**, or **V** commands.
 - On BSD, if there is no inserted text, the insert command changes the current line to the referenced line –1; that is, the line before the specified line.
 - On BSD, the join command with only a single address changes the current line to that address.
 - ullet BSD does not support the P command; moreover, in BSD it is synonymous with the p command.
 - BSD does not support the *undo* of the commands **j**, **m**, **r**, **s**, or **t**.
 - The Version 7 *ed* command **W**, and the BSD *ed* commands **W**, **wq**, and **z** are not present in this volume of IEEE Std 1003.1-2001.

The -s option was added to allow the functionality of the now withdrawn – option in a manner compatible with the Utility Syntax Guidelines.

In early proposals there was a limit, {ED_FILE_MAX}, that described the historical limitations of some *ed* utilities in their handling of large files; some of these have had problems with files larger than 100 000 bytes. It was this limitation that prompted much of the desire to include a *split* command in this volume of IEEE Std 1003.1-2001. Since this limit was removed, this volume of IEEE Std 1003.1-2001 requires that implementations document the file size limits imposed by *ed* in the conformance document. The limit {ED_LINE_MAX} was also removed; therefore, the global limit {LINE_MAX} is used for input and output lines.

The manner in which the I command writes non-printable characters was changed to avoid the historical backspace-overstrike method. On video display terminals, the overstrike is ambiguous because most terminals simply replace overstruck characters, making the I format not useful for its intended purpose of unambiguously understanding the content of the line. The historical backslash escapes were also ambiguous. (The string "a\0011" could represent a line containing those six characters or a line containing the three characters 'a', a byte with a binary value of 1, and a 1.) In the format required here, a backslash appearing in the line is written as "\\" so that the output is truly unambiguous. The method of marking the ends of lines was adopted from the ex editor and is required for any line ending in <space>s; the '\$' is placed on all lines so that a real '\$' at the end of a line cannot be misinterpreted.

Systems with bytes too large to fit into three octal digits must devise other means of displaying non-printable characters. Consideration was given to requiring that the number of octal digits be large enough to hold a byte, but this seemed to be too confusing for applications on the vast majority of systems where three digits are adequate. It would be theoretically possible for the application to use the *getconf* utility to find out the CHAR_BIT value and deal with such an algorithm; however, there is really no portable way that an application can use the octal values of the bytes across various coded character sets, so the additional specification was not worthwhile.

The description of how a NUL is written was removed. The NUL character cannot be in text files, and this volume of IEEE Std 1003.1-2001 should not dictate behavior in the case of undefined, erroneous input.

Unlike some of the other editing utilities, the filenames accepted by the E, e, R, and r commands are not patterns.

Early proposals stated that the $-\mathbf{p}$ option worked only when standard input was associated with a terminal device. This has been changed to conform to historical implementations, thereby allowing applications to interpose themselves between a user and the ed utility.

The form of the substitute command that uses the $\bf n$ suffix was limited in some historical documentation (where this was described incorrectly as "backreferencing"). This limit has been omitted because there is no reason why an editor processing lines of {LINE_MAX} length should have this restriction. The command $\bf s/x/X/2047$ should be able to substitute the 2 047th occurrence of 'x' on a line.

The use of printing commands with printing suffixes (such as **pn**, **lp**, and so on) was made unspecified because BSD-based systems allow this, whereas System V does not.

Some BSD-based systems exit immediately upon receipt of end-of-file if all of the lines in the file have been deleted. Since this volume of IEEE Std 1003.1-2001 refers to the $\bf q$ command in this instance, such behavior is not allowed.

Some historical implementations returned exit status zero even if command errors had occurred; this is not allowed by this volume of IEEE Std 1003.1-2001.

Some historical implementations contained a bug that allowed a single period to be entered in input mode as
backslash> period> <newline>. This is not allowed by ed because there is no description of escaping any of the characters in input mode; backslashes are entered into the buffer exactly as typed. The typical method of entering a single period has been to precede it with another character and then use the substitute command to delete that character.

It is difficult under some modes of some versions of historical operating system terminal drivers to distinguish between an end-of-file condition and terminal disconnect. IEEE Std 1003.1-2001 does not require implementations to distinguish between the two situations, which permits historical implementations of the *ed* utility on historical platforms to conform. Implementations are encouraged to distinguish between the two, if possible, and take appropriate action on terminal disconnect.

Historically, ed accepted a zero address for the a and r commands in order to insert text at the start of the edit buffer. When the buffer was empty the command .= returned zero. IEEE Std 1003.1-2001 requires conformance to historical practice.

For consistency with the \mathbf{a} and \mathbf{r} commands and better user functionality, the \mathbf{i} and \mathbf{c} commands must also accept an address of 0, in which case 0i is treated as 1i and likewise for the \mathbf{c} command.

All of the following are valid addresses:

13497 +++ Three lines after the current line.

13498 /pattern/- One line before the next occurrence of pattern.

13499 −2 Two lines before the current line.

3 ---- 2 Line one (note the intermediate negative address).

13501 1 2 3 Line six.

Any number of addresses can be provided to commands taking addresses; for example, "1,2,3,4,5p" prints lines 4 and 5, because two is the greatest valid number of addresses accepted by the **print** command. This, in combination with the semicolon delimiter, permits users to create commands based on ordered patterns in the file. For example, the command "3;/foo/;+2p" will display the first line after line 3 that contains the pattern *foo*, plus the next two lines. Note that the address "3;" must still be evaluated before being discarded, because the search origin for the "/foo/" command depends on this.

Historically, *ed* disallowed address chains, as discussed above, consisting solely of comma or semicolon separators; for example, ", , , " or ";;" were considered an error. For consistency of address specification, this restriction is removed. The following table lists some of the address

ed Utilities

13512	forms now	possible:
13512	forms now	/ possib

13513	Address	Addr1	Addr2	Status	Comment
13514	7,	7	7	Historical	
13515	7,5,	5	5	Historical	
13516	7,5,9	5	9	Historical	
13517	7,9	7	9	Historical	
13518	7,+	7	8	Historical	
13519	,	1	\$	Historical	
13520	, 7	1	7	Extension	
13521	, ,	\$	\$	Extension	
13522	, ;	\$	\$	Extension	
13523	7;	7	7	Historical	
13524	7;5;	5	5	Historical	
13525	7;5;9	5	9	Historical	
13526	7;5,9	5	9	Historical	
13527	7;\$;4	\$	4	Historical	Valid, but erroneous.
13528	7;9	7	9	Historical	
13529	7;+	7	8	Historical	
13530	;	•	\$	Historical	
13531	;7	•	7	Extension	
13532	;;	\$	\$	Extension	
13533	; ,	\$	\$	Extension	

Historically, values could be added to addresses by including them after one or more

silonk>s; for example, "3 – 5p" wrote the seventh line of the file, and "/foo/ 5" was the same as

"5 /foo/". However, only absolute values could be added; for example, "5 /foo/" was an error. IEEE Std 1003.1-2001 requires conformance to historical practice.

Historically, *ed* accepted the '^' character as an address, in which case it was identical to the hyphen character. IEEE Std 1003.1-2001 does not require or prohibit this behavior.

13540 FUTURE DIRECTIONS

13541 None.

13542 **SEE ALSO**

13534

13535

1353613537

13543 Section 1.11 (on page 20), ex, sed, sh, vi

13544 CHANGE HISTORY

First released in Issue 2.

13546 **Issue 5**

13547 In the OPTIONS section, the meaning of -s and - is clarified.

13548 A second FUTURE DIRECTION is added.

13549 **Issue 6**

13550 The obsolescent single-minus form is removed.

13551 A second APPLICATION USAGE note is added.

The Open Group Corrigendum U025/2 is applied, correcting the description of the Edit section.

The *ed* utility is updated to align with the IEEE P1003.2b draft standard. This includes addition of the treatment of the SIGQUIT signal, changes to *ed* addressing, and changes to processing when end-of-file is detected and when terminal disconnect is detected.

13556	The normative text is reworded to avoid use of the term—must—for application requirements.
13557	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/22 is applied, adding the text: "Any line
13558	modified by the <i>command list</i> shall be unmarked." to the G command. This change corresponds
13559	to a similar change made to the ${f g}$ command in the first version of IEEE Std 1003.1-2001.

env Utilities

13560 NAME		
13561	env — set th	e environment for command invocation
13562 SYNOP		
13563	env [-1][name=value] [utility [argument]]
13564 DESCR		
13565 13566		ity shall obtain the current environment, modify it according to its arguments, then itility named by the <i>utility</i> operand with the modified environment.
13567	Optional arg	guments shall be passed to <i>utility</i> .
13568 13569	v	operand is specified, the resulting environment shall be written to the standard one <i>name=value</i> pair per line.
13570 OPTIO	NS	
13571 13572		ity shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section Syntax Guidelines.
13573	The followin	ng options shall be supported:
13574 13575	- i	Invoke <i>utility</i> with exactly the environment specified by the arguments; the inherited environment shall be ignored completely.
13576 OPERA	NDS	
13577	The following	ng operands shall be supported:
13578 13579	name=value	Arguments of the form <i>name=value</i> shall modify the execution environment, and shall be placed into the inherited environment before the <i>utility</i> is invoked.
13580 13581	utility	The name of the utility to be invoked. If the <i>utility</i> operand names any of the special built-in utilities in Section 2.14 (on page 64), the results are undefined.
13582	argument	A string to pass as an argument for the invoked utility.
13583 STDIN		
13584	Not used.	
13585 INPUT 13586	FILES None.	
13587 ENVIR	ONMENT VA	ARIABLES
13588	The following	ng environment variables shall affect the execution of env:
13589	LANG	Provide a default value for the internationalization variables that are unset or null.
13590		(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
13591 13592		Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
	LC_ALL	<u> </u>
13593 13594	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
13595 13596	LC_CTYPE	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
13597		arguments).
13598	LC_MESSA	GES
13599		Determine the locale that should be used to affect the format and contents of
13600		diagnostic messages written to standard error.
13601 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .

Utilities **env**

Determine the location of the *utility*, as described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables. If *PATH* is specified as a name=value operand to *env*, the value given shall be used in the search for *utility*.

13605 ASYNCHRONOUS EVENTS

13606 Default.

13607 STDOUT

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

13610 "%s=%s\n", <name>, <value>

13611 If the *utility* operand is specified, the *env* utility shall not write to standard output.

13612 STDERR

13613 The standard error shall be used only for diagnostic messages.

13614 OUTPUT FILES

13615 None.

13616 EXTENDED DESCRIPTION

13617 None.

13618 EXIT STATUS

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

13621 0 The *env* utility completed successfully.

13622 1–125 An error occurred in the *env* utility.

13623 126 The utility specified by *utility* was found but could not be invoked.

13624 127 The utility specified by *utility* could not be found.

13625 CONSEQUENCES OF ERRORS

13626 Default.

13628

13629

13630

13631

13632

13633

13634

13635

13636

13637

13638

13639

13640

13641

13642

13627 APPLICATION USAGE

The *command, env, nice, nohup, time,* 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.

Historical implementations of the *env* utility use the *execvp()* or *execlp()* functions defined in the System Interfaces volume of IEEE Std 1003.1-2001 to invoke the specified utility; this provides better performance and keeps users from having to escape characters with special meaning to the shell. Therefore, shell functions, special built-ins, and built-ins that are only provided by the shell are not found.

env Utilities

13643 EXAMPLES 13644 The following command: 13645 env -i PATH=/mybin mygrep xyz myfile invokes the command mygrep with a new PATH value as the only entry in its environment. In 13646 13647 this case, *PATH* is used to locate *mygrep*, which then must reside in /**mybin**. 13648 RATIONALE 13649 As with all other utilities that invoke other utilities, this volume of IEEE Std 1003.1-2001 only 13650 specifies what env does with standard input, standard output, standard error, input files, and output files. If a utility is executed, it is not constrained by the specification of input and output 13651 by env. 13652 The -i option was added to allow the functionality of the withdrawn - option in a manner 13653 compatible with the Utility Syntax Guidelines. 13654 Some have suggested that *env* is redundant since the same effect is achieved by: 13655 name=value ... utility [argument ...] 13656 The example is equivalent to env when an environment variable is being added to the 13657 environment of the command, but not when the environment is being set to the given value. 13658 13659 The env utility also writes out the current environment if invoked without arguments. There is sufficient functionality beyond what the example provides to justify inclusion of env. 13660 13661 FUTURE DIRECTIONS None. 13662 13663 SEE ALSO Section 2.5 (on page 33), Section 2.14 (on page 64) 13664 13665 CHANGE HISTORY First released in Issue 2. 13666

13667 NAME 13668 ex — text editor 13669 SYNOPSIS ex [-rR] $[-s \mid -v]$ [-c command] [-t tagstring] [-w size] [file ...]13670 UP 13671 13672 **DESCRIPTION** The ex utility is a line-oriented text editor. There are two other modes of the editor—open and 13673 visual—in which screen-oriented editing is available. This is described more fully by the *ex* open 13674 and **visual** commands and in vi. 13675 This section uses the term edit buffer to describe the current working text. No specific 13676 implementation is implied by this term. All editing changes are performed on the edit buffer, 13677 13678 and no changes to it shall affect any file until an editor command writes the file. Certain terminals do not have all the capabilities necessary to support the complete *ex* definition, 13679 such as the full-screen editing commands (visual mode or open mode). When these commands 13680 cannot be supported on such terminals, this condition shall not produce an error message such 13681 as "not an editor command" or report a syntax error. The implementation may either accept the 13682 commands and produce results on the screen that are the result of an unsuccessful attempt to 13683 meet the requirements of this volume of IEEE Std 1003.1-2001 or report an error describing the 13684 terminal-related deficiency. 13685 13686 OPTIONS 13687 The ex utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines. 13688 13689 The following options shall be supported: -c command Specify an initial command to be executed in the first edit buffer loaded from an 13690 13691 existing file (see the EXTENDED DESCRIPTION section). Implementations may support more than a single -c option. In such implementations, the specified 13692 commands shall be executed in the order specified on the command line. 13693 Recover the named files (see the EXTENDED DESCRIPTION section). Recovery 13694 $-\mathbf{r}$ information for a file shall be saved during an editor or system crash (for example, 13695 when the editor is terminated by a signal which the editor can catch), or after the 13696 use of an ex preserve command. 13697 A *crash* in this context is an unexpected failure of the system or utility that requires 13698 13699

A *crash* in this context is an unexpected failure of the system or utility that requires restarting the failed system or utility. A system crash implies that any utilities running at the time also crash. In the case of an editor or system crash, the number of changes to the edit buffer (since the most recent **preserve** command) that will be recovered is unspecified.

If no *file* operands are given and the –t option is not specified, all other options, the *EXINIT* variable, and any .exrc files shall be ignored; a list of all recoverable files available to the invoking user shall be written, and the editor shall exit normally without further action.

13707 — R Set **readonly** edit option.

13700

13701

13702

13703

13704

13705

13706

13708

13709 13710

13711

-s Prepare *ex* for batch use by taking the following actions:

- Suppress writing prompts and informational (but not diagnostic) messages.
- Ignore the value of *TERM* and any implementation default terminal type and assume the terminal is a type incapable of supporting open or visual modes;

ex Utilities

13712		see the visual command and the description of <i>vi</i> .		
13713 13714		 Suppress the use of the EXINIT environment variable and the reading of any .exrc file; see the EXTENDED DESCRIPTION section. 		
13715		• Suppress autoindentation, ignoring the value of the autoindent edit option.		
13716 13717 13718 13719 13720	–t tagstring	Edit the file containing the specified <i>tagstring</i> ; see <i>ctags</i> . The tags feature represented by –t <i>tagstring</i> and the tag command is optional. It shall be provided on any system that also provides a conforming implementation of <i>ctags</i> ; otherwise, the use of –t produces undefined results. On any system, it shall be an error to specify more than a single –t option.		
13721	- v	Begin in visual mode (see vi).		
13722	−w size	Set the value of the window editor option to size.		
13723 OPERANDS 13724 The following operand shall be supported:				
13725	file	A pathname of a file to be edited.		
13725 13726 STDIN	me	A patiname of a file to be cured.		
13727 13728 13729	The standard input consists of a series of commands and input text, as described in the EXTENDED DESCRIPTION section. The implementation may limit each line of standard input to a length of {LINE_MAX}.			
13730	If the standard input is not a terminal device, it shall be as if the $-\mathbf{s}$ option had been specified.			
13731 13732	If a read from the standard input returns an error, or if the editor detects an end-of-file condition from the standard input, it shall be equivalent to a SIGHUP asynchronous event.			
13733 INPUT 13734 13735 13736 13737	Input files shall be text files or files that would be text files except for an incomplete last line that is not longer than {LINE_MAX}–1 bytes in length and contains no NUL characters. By default, any incomplete last line shall be treated as if it had a trailing <newline>. The editing of other forms of files may optionally be allowed by <i>ex</i> implementations.</newline>			
13738 13739	The .exrc files and source files shall be text files consisting of <i>ex</i> commands; see the EXTENDED DESCRIPTION section.			
13740 13741	By default, the editor shall read lines from the files to be edited without interpreting any of those lines as any form of editor command.			
13742 ENVIRONMENT VARIABLES 13743 The following environment variables shall affect the execution of <i>ex</i> :				
13744 13745 13746	COLUMNS	Override the system-selected horizontal screen size. See the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables for valid values and results when it is unset or null.		
13747 13748	EXINIT	Determine a list of <i>ex</i> commands that are executed on editor start-up. See the EXTENDED DESCRIPTION section for more details of the initialization phase.		
13749 13750	HOME	Determine a pathname of a directory that shall be searched for an editor start-up file named .exrc; see the EXTENDED DESCRIPTION section.		
13751 13752 13753 13754	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)		

13755 13756	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.		
13757	LC_COLLATE			
13758 13759		Determine the locale for the behavior of ranges, equivalence classes, and multi- character collating elements within regular expressions.		
13760 13761 13762 13763 13764	LC_CTYPE	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), the behavior of character classes within regular expressions, the classification of characters as uppercase or lowercase letters, the case conversion of letters, and the detection of word boundaries.		
13765 13766 13767	LC_MESSAGES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.			
13768 13769 13770 13771	LINES	Override the system-selected vertical screen size, used as the number of lines in a screenful and the vertical screen size in visual mode. See the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables for valid values and results when it is unset or null.		
13772 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .		
13773 13774 13775	PATH	Determine the search path for the shell command specified in the <i>ex</i> editor commands !, shell , read , and write , and the open and visual mode command !; see the description of command search and execution in Section 2.9.1.1 (on page 48).		
13776 13777	SHELL	Determine the preferred command line interpreter for use as the default value of the shell edit option.		
13778 13779	TERM	Determine the name of the terminal type. If this variable is unset or null, an unspecified default terminal type shall be used.		
13780 ASYNCHRONOUS EVENTS				
13781 13782	The following term is used in this and following sections to specify command and asynchronous event actions:			
13783 13784 13785 13786 13787 13788	complete writ	A complete write is a write of the entire contents of the edit buffer to a file of a type other than a terminal device, or the saving of the edit buffer caused by the user executing the <i>ex</i> preserve command. Writing the contents of the edit buffer to a temporary file that will be removed when the editor exits shall not be considered a complete write.		
13789	The following actions shall be taken upon receipt of signals:			
13790 13791	SIGINT	If the standard input is not a terminal device, <i>ex</i> shall not write the file or return to command or text input mode, and shall exit with a non-zero exit status.		
13792 13793		Otherwise, if executing an open or visual text input mode command, <i>ex</i> in receipt of SIGINT shall behave identically to its receipt of the <esc> character.</esc>		
13794		Otherwise:		
13795 13796 13797		1. If executing an <i>ex</i> text input mode command, all input lines that have been completely entered shall be resolved into the edit buffer, and any partially entered line shall be discarded.		

ex Utilities

13798 2. If there is a currently executing command, it shall be aborted and a message displayed. Unless otherwise specified by the ex or vi command descriptions, 13799 it is unspecified whether any lines modified by the executing command 13800 appear modified, or as they were before being modified by the executing 13801 command, in the buffer. 13802 If the currently executing command was a motion command, its associated 13803 command shall be discarded. 13804 If in open or visual command mode, the terminal shall be alerted. 13805 The editor shall then return to command mode. 13806 **SIGCONT** The screen shall be refreshed if in open or visual mode. 13807 **SIGHUP** If the edit buffer has been modified since the last complete write, ex shall attempt 13808 to save the edit buffer so that it can be recovered later using the -r option or the ex 13809 recover command. The editor shall not write the file or return to command or text 13810 input mode, and shall terminate with a non-zero exit status. 13811 **SIGTERM** Refer to SIGHUP. 13812 The action taken for all other signals is unspecified. 13813

13814 STDOUT

The standard output shall be used only for writing prompts to the user, for informational messages, and for writing lines from the file.

13817 STDERR

13824

13825

13826

13827

13828 13829

13830

13831

13832

13833

13834

13835

13836

13837

13838 13839

13818 The standard error shall be used only for diagnostic messages.

13819 OUTPUT FILES

13820 The output from *ex* shall be text files.

13821 EXTENDED DESCRIPTION

Only the *ex* mode of the editor is described in this section. See *vi* for additional editing capabilities available in *ex*.

When an error occurs, *ex* shall write a message. If the terminal supports a standout mode (such as inverse video), the message shall be written in standout mode. If the terminal does not support a standout mode, and the edit option **errorbells** is set, an alert action shall precede the error message.

By default, *ex* shall start in command mode, which shall be indicated by a: prompt; see the **prompt** command. Text input mode can be entered by the **append**, **insert**, or **change** commands; it can be exited (and command mode re-entered) by typing a period (' . ') alone at the beginning of a line.

Initialization in ex and vi

The following symbols are used in this and following sections to specify locations in the edit buffer:

alternate and current pathnames

Two pathnames, named *current* and *alternate*, are maintained by the editor. Any *ex* commands that take filenames as arguments shall set them as follows:

1. If a *file* argument is specified to the *ex* **edit**, **ex**, or **recover** commands, or if an *ex* **tag** command replaces the contents of the edit buffer.

 a. If the command replaces the contents of the edit buffer, the current pathname shall be set to the *file* argument or the file indicated by the tag, and the alternate pathname shall be set to the previous value of the current pathname.

b. Otherwise, the alternate pathname shall be set to the *file* argument.

2. If a *file* argument is specified to the *ex* **next** command:

 a. If the command replaces the contents of the edit buffer, the current pathname shall be set to the first *file* argument, and the alternate pathname shall be set to the previous value of the current pathname.

 3. If a *file* argument is specified to the *ex* **file** command, the current pathname shall be set to the *file* argument, and the alternate pathname shall be set to the previous value of the current pathname.

 4. If a *file* argument is specified to the *ex* **read** and **write** commands (that is, when reading or writing a file, and not to the program named by the **shell** edit option), or a *file* argument is specified to the *ex* **xit** command:

file argument is specified to the ex xit command:a. If the current pathname has no value, the current pathname shall be set to the file

b. Otherwise, the alternate pathname shall be set to the *file* argument.

 If the alternate pathname is set to the previous value of the current pathname when the current pathname had no previous value, then the alternate pathname shall have no value as a result.

current line

argument.

 The line of the edit buffer referenced by the cursor. Each command description specifies the current line after the command has been executed, as the *current line value*. When the edit buffer contains no lines, the current line shall be zero; see **Addressing in ex** (on page 361).

current column

The current display line column occupied by the cursor. (The columns shall be numbered beginning at 1.) Each command description specifies the current column after the command has been executed, as the *current column* value. This column is an *ideal* column that is remembered over the lifetime of the editor. The actual display line column upon which the cursor rests may be different from the current column; see the cursor positioning discussion in **Command Descriptions in vi** (on page 988).

set to non-<blank>

A description for a current column value, meaning that the current column shall be set to the last display line column on which is displayed any part of the first non-
blank> of the line. If the line has no non-
blank> non-<newline>s, the current column shall be set to the last display line column on which is displayed any part of the last non-<newline> in the line. If the line is empty, the current column shall be set to column position 1.

The length of lines in the edit buffer may be limited to {LINE_MAX} bytes. In open and visual mode, the length of lines in the edit buffer may be limited to the number of characters that will fit in the display. If either limit is exceeded during editing, an error message shall be written. If either limit is exceeded by a line read in from a file, an error message shall be written and the edit session may be terminated.

 If the editor stops running due to any reason other than a user command, and the edit buffer has been modified since the last complete write, it shall be equivalent to a SIGHUP asynchronous event. If the system crashes, it shall be equivalent to a SIGHUP asynchronous event.

During initialization (before the first file is copied into the edit buffer or any user commands from the terminal are processed) the following shall occur:

- 1. If the environment variable *EXINIT* is set, the editor shall execute the *ex* commands contained in that variable.
- 2. If the *EXINIT* variable is not set, and all of the following are true:
 - a. The *HOME* environment variable is not null and not empty.
 - b. The file **.exrc** in the directory referred to by the *HOME* environment variable:
 - 1. Exists
 - 2. Is owned by the same user ID as the real user ID of the process or the process has appropriate privileges
 - 3. Is not writable by anyone other than the owner

the editor shall execute the ex commands contained in that file.

- 3. If and only if all of the following are true:
 - a. The current directory is not referred to by the *HOME* environment variable.
 - b. A command in the *EXINIT* environment variable or a command in the .exrc file in the directory referred to by the *HOME* environment variable sets the editor option exrc.
 - c. The **.exrc** file in the current directory:
 - 1. Exists
 - 2. Is owned by the same user ID as the real user ID of the process, or by one of a set of implementation-defined user IDs
 - 3. Is not writable by anyone other than the owner

the editor shall attempt to execute the ex commands contained in that file.

Lines in any .exrc file that are blank lines shall be ignored. If any .exrc file exists, but is not read for ownership or permission reasons, it shall be an error.

After the *EXINIT* variable and any **.exrc** files are processed, the first file specified by the user shall be edited, as follows:

- 1. If the user specified the —t option, the effect shall be as if the ex tag command was entered with the specified argument, with the exception that if tag processing does not result in a file to edit, the effect shall be as described in step 3. below.
- 2. Otherwise, if the user specified any command line *file* arguments, the effect shall be as if the *ex* **edit** command was entered with the first of those arguments as its *file* argument.
- 3. Otherwise, the effect shall be as if the *ex* **edit** command was entered with a nonexistent filename as its *file* argument. It is unspecified whether this action shall set the current pathname. In an implementation where this action does not set the current pathname, any editor command using the current pathname shall fail until an editor command sets the current pathname.

If the —r option was specified, the first time a file in the initial argument list or a file specified by the —t option is edited, if recovery information has previously been saved about it, that information shall be recovered and the editor shall behave as if the contents of the edit buffer have already been modified. If there are multiple instances of the file to be recovered, the one most recently saved shall be recovered, and an informational message that there are previous

versions of the file that can be recovered shall be written. If no recovery information about a file is available, an informational message to this effect shall be written, and the edit shall proceed as usual.

If the -c option was specified, the first time a file that already exists (including a file that might not exist but for which recovery information is available, when the -r option is specified) replaces or initializes the contents of the edit buffer, the current line shall be set to the last line of the edit buffer, the current column shall be set to non-<blank>, and the ex commands specified with the -c option shall be executed. In this case, the current line and current column shall not be set as described for the command associated with the replacement or initialization of the edit buffer contents. However, if the -t option or a tag command is associated with this action, the -c option commands shall be executed and then the movement to the tag shall be performed.

The current argument list shall initially be set to the filenames specified by the user on the command line. If no filenames are specified by the user, the current argument list shall be empty. If the <code>-t</code> option was specified, it is unspecified whether any filename resulting from tag processing shall be prepended to the current argument list. In the case where the filename is added as a prefix to the current argument list, the current argument list reference shall be set to that filename. In the case where the filename is not added as a prefix to the current argument list, the current argument list reference shall logically be located before the first of the filenames specified on the command line (for example, a subsequent <code>ex next</code> command shall edit the first filename from the command line). If the <code>-t</code> option was not specified, the current argument list reference shall be to the first of the filenames on the command line.

Addressing in ex

 Addressing in *ex* relates to the current line and the current column; the address of a line is its 1-based line number, the address of a column is its 1-based count from the beginning of the line. Generally, the current line is the last line affected by a command. The current line number is the address of the current line. In each command description, the effect of the command on the current line number and the current column is described.

Addresses are constructed as follows:

- 1. The character '.' (period) shall address the current line.
- 2. The character '\$' shall address the last line of the edit buffer.
- 3. The positive decimal number *n* shall address the *n*th line of the edit buffer.
- 4. The address "'x" refers to the line marked with the mark name character 'x', which shall be a lowercase letter from the portable character set or one of the characters ''' or '''. It shall be an error if the line that was marked is not currently present in the edit buffer or the mark has not been set. Lines can be marked with the *ex* mark or **k** commands, or the *vi* m command.
- 5. A regular expression enclosed by slashes ('/') shall address the first line found by searching forwards from the line following the current line toward the end of the edit buffer and stopping at the first line for which the line excluding the terminating <newline> matches the regular expression. As stated in **Regular Expressions in ex** (on page 391), an address consisting of a null regular expression delimited by slashes "//" shall address the next line for which the line excluding the terminating <newline> matches the last regular expression encountered. In addition, the second slash can be omitted at the end of a command line. If the **wrapscan** edit option is set, the search shall wrap around to the beginning of the edit buffer and continue up to and including the current line, so that the entire edit buffer is searched. Within the regular expression, the sequence "\/" shall represent a literal slash instead of the regular expression delimiter.

6. A regular expression enclosed in question marks ('?') shall address the first line found by searching backwards from the line preceding the current line toward the beginning of the edit buffer and stopping at the first line for which the line excluding the terminating <newline> matches the regular expression. An address consisting of a null regular expression delimited by question marks "??" shall address the previous line for which the line excluding the terminating <newline> matches the last regular expression encountered. In addition, the second question mark can be omitted at the end of a command line. If the wrapscan edit option is set, the search shall wrap around from the beginning of the edit buffer to the end of the edit buffer and continue up to and including the current line, so that the entire edit buffer is searched. Within the regular expression, the sequence "\?" shall represent a literal question mark instead of the RE delimiter.

7. A plus sign ('+') or a minus sign ('-') followed by a decimal number shall address the current line plus or minus the number. A '+' or '-' not followed by a decimal number shall address the current line plus or minus 1.

Addresses can be followed by zero or more address offsets, optionally <blank>-separated. Address offsets are constructed as follows:

- 1. A '+' or '-' immediately followed by a decimal number shall add (subtract) the indicated number of lines to (from) the address. A '+' or '-' not followed by a decimal number shall add (subtract) 1 to (from) the address.
- 2. A decimal number shall add the indicated number of lines to the address.

It shall not be an error for an intermediate address value to be less than zero or greater than the last line in the edit buffer. It shall be an error for the final address value to be less than zero or greater than the last line in the edit buffer.

Commands take zero, one, or two addresses; see the descriptions of *1addr* and *2addr* in **Command Descriptions in ex** (on page 368). If more than the required number of addresses are provided to a command that requires zero addresses, it shall be an error. Otherwise, if more than the required number of addresses are provided to a command, the addresses specified first shall be evaluated and then discarded until the maximum number of valid addresses remain.

Addresses shall be separated from each other by a comma (',') or a semicolon (';'). If no address is specified before or after a comma or semicolon separator, it shall be as if the address of the current line was specified before or after the separator. In the case of a semicolon separator, the current line (',') shall be set to the first address, and only then will the next address be calculated. This feature can be used to determine the starting line for forwards and backwards searches (see rules 5. and 6.).

A percent sign (' %') shall be equivalent to entering the two addresses "1, \$".

Any delimiting

blank>s between addresses, address separators, or address offsets shall be discarded.

Command Line Parsing in ex

The following symbol is used in this and following sections to describe parsing behavior:

If a character is referred to as "backslash-escaped" or "<control>-V-escaped," it shall mean that the character acquired or lost a special meaning by virtue of being preceded, respectively, by a backslash or <control>-V character. Unless otherwise specified, the escaping character shall be discarded at that time and shall not be further considered for any purpose.

14017 Command-line parsing shall be done in the following steps. For each step, characters already evaluated shall be ignored; that is, the phrase "leading character" refers to the next character that has not yet been evaluated.

- 1. Leading colon characters shall be skipped.
- 2. Leading <blank>s shall be skipped.
- 3. If the leading character is a double-quote character, the characters up to and including the next non-backslash-escaped <newline> shall be discarded, and any subsequent characters shall be parsed as a separate command.
- 4. Leading characters that can be interpreted as addresses shall be evaluated; see **Addressing** in ex (on page 361).
- 5. Leading <blank>s shall be skipped.
- 6. If the next character is a vertical-line character or a <newline>:
 - a. If the next character is a <newline>:
 - 1. If *ex* is in open or visual mode, the current line shall be set to the last address specified, if any.
 - 2. Otherwise, if the last command was terminated by a vertical-line character, no action shall be taken; for example, the command "||<newline>" shall execute two implied commands, not three.
 - 3. Otherwise, step 6.b. shall apply.
 - b. Otherwise, the implied command shall be the **print** command. The last #, **p**, and **l** flags specified to any *ex* command shall be remembered and shall apply to this implied command. Executing the *ex* **number**, **print**, or **list** command shall set the remembered flags to #, nothing, and **l**, respectively, plus any other flags specified for that execution of the **number**, **print**, or **list** command.
 - If *ex* is not currently performing a **global** or **v** command, and no address or count is specified, the current line shall be incremented by 1 before the command is executed. If incrementing the current line would result in an address past the last line in the edit buffer, the command shall fail, and the increment shall not happen.
 - c. The <newline> or vertical-line character shall be discarded and any subsequent characters shall be parsed as a separate command.
- 7. The command name shall be comprised of the next character (if the character is not alphabetic), or the next character and any subsequent alphabetic characters (if the character is alphabetic), with the following exceptions:
 - a. Commands that consist of any prefix of the characters in the command name **delete**, followed immediately by any of the characters 'l', 'p', '+', '-', or '#' shall be interpreted as a **delete** command, followed by a <blank>, followed by the characters that were not part of the prefix of the **delete** command. The maximum number of characters shall be matched to the command name **delete**; for example, "del" shall not be treated as "de" followed by the flag **l**.
 - b. Commands that consist of the character 'k', followed by a character that can be used as the name of a mark, shall be equivalent to the mark command followed by a
blank>, followed by the character that followed the 'k'.
 - c. Commands that consist of the character 's', followed by characters that could be interpreted as valid options to the s command, shall be the equivalent of the s

command, without any pattern or replacement values, followed by a <blank>, followed by the characters after the 's'.

8. The command name shall be matched against the possible command names, and a command name that contains a prefix matching the characters specified by the user shall be the executed command. In the case of commands where the characters specified by the user could be ambiguous, the executed command shall be as follows:

a	append	n	next	t	t
c	change	р	print	u	undo
ch	change	pr	print	un	undo
e	edit	r	read	v	v
m	move	re	read	w	write
ma	mark	s	s		

Implementation extensions with names causing similar ambiguities shall not be checked for a match until all possible matches for commands specified by IEEE Std 1003.1-2001 have been checked.

- 10. Otherwise, if the command is an **edit**, **ex**, or **next** command, or a **visual** command while in open or visual mode, the next part of the command shall be parsed as follows:
 - a. Any '!' character immediately following the command shall be skipped and be part of the command.
 - b. Any leading

 shall be skipped and be part of the command.
 - c. If the next character is a '+', characters up to the first non-backslash-escaped <newline> or non-backslash-escaped <blank> shall be skipped and be part of the command.
 - d. The rest of the command shall be determined by the steps specified in paragraph 12.
- 11. Otherwise, if the command is a **global**, **open**, **s**, or **v** command, the next part of the command shall be parsed as follows:
 - a. Any leading

 blank>s shall be skipped and be part of the command.
 - b. If the next character is not an alphanumeric, double-quote, <newline>, backslash, or vertical-line character:
 - 1. The next character shall be used as a command delimiter.
 - 2. If the command is a **global**, **open**, or **v** command, characters up to the first non-backslash-escaped <newline>, or first non-backslash-escaped delimiter character, shall be skipped and be part of the command.
 - 3. If the command is an **s** command, characters up to the first non-backslash-escaped <newline>, or second non-backslash-escaped delimiter character, shall be skipped and be part of the command.
 - c. If the command is a **global** or **v** command, characters up to the first non-backslash-escaped <newline> shall be skipped and be part of the command.

d. Otherwise, the rest of the command shall be determined by the steps specified in paragraph 12.

12. Otherwise:

- a. If the command was a **map**, **unmap**, **abbreviate**, or **unabbreviate** command, characters up to the first non-<control>-V-escaped <newline>, vertical-line, or double-quote character shall be skipped and be part of the command.
- b. Otherwise, characters up to the first non-backslash-escaped <newline>, vertical-line, or double-quote character shall be skipped and be part of the command.
- c. If the command was an **append**, **change**, or **insert** command, and the step 12.b. ended at a vertical-line character, any subsequent characters, up to the next non-backslash-escaped <newline> shall be used as input text to the command.
- d. If the command was ended by a double-quote character, all subsequent characters, up to the next non-backslash-escaped <newline>, shall be discarded.
- e. The terminating <newline> or vertical-line character shall be discarded and any subsequent characters shall be parsed as a separate *ex* command.

Command arguments shall be parsed as described by the Synopsis and Description of each individual *ex* command. This parsing shall not be <blank>-sensitive, except for the ! argument, which must follow the command name without intervening <blank>s, and where it would otherwise be ambiguous. For example, *count* and *flag* arguments need not be <blank>-separated because "d22p" is not ambiguous, but *file* arguments to the *ex* **next** command must be separated by one or more
blank>s. Any <blank> in command arguments for the **abbreviate**, **unabbreviate**, **map**, and **unmap** commands can be <control>-V-escaped, in which case the
blank> shall not be used as an argument delimiter. Any
blank> in the command argument for any other command can be backslash-escaped, in which case that
blank> shall not be used as an argument delimiter.

Within command arguments for the **abbreviate**, **unabbreviate**, **map**, and **unmap** commands, any character can be <control>-V-escaped. All such escaped characters shall be treated literally and shall have no special meaning. Within command arguments for all other *ex* commands that are not regular expressions or replacement strings, any character that would otherwise have a special meaning can be backslash-escaped. Escaped characters shall be treated literally, without special meaning as shell expansion characters or '!', '%', and '#' expansion characters. See **Regular Expressions in ex** (on page 391) and **Replacement Strings in ex** (on page 391) for descriptions of command arguments that are regular expressions or replacement strings.

Non-backslash-escaped '%' characters appearing in *file* arguments to any *ex* command shall be replaced by the current pathname; unescaped '#' characters shall be replaced by the alternate pathname. It shall be an error if '%' or '#' characters appear unescaped in an argument and their corresponding values are not set.

If an error occurs during the parsing or execution of an *ex* command:

- An informational message to this effect shall be written. Execution of the *ex* command shall stop, and the cursor (for example, the current line and column) shall not be further modified.
 - If the *ex* command resulted from a map expansion, all characters from that map expansion shall be discarded, except as otherwise specified by the **map** command.
 - Otherwise, if the *ex* command resulted from the processing of an *EXINIT* environment variable, a .exrc file, a :source command, a -c option, or a +command specified to an *ex* edit, ex, next, or visual command, no further commands from the source of the commands shall be executed.
 - Otherwise, if the ex command resulted from the execution of a buffer or a global or v command, no further commands caused by the execution of the buffer or the global or v command shall be executed.
 - Otherwise, if the *ex* command was not terminated by a <newline>, all characters up to and including the next non-backslash-escaped <newline> shall be discarded.

Input Editing in ex

The following symbol is used in this and the following sections to specify command actions:

word In the POSIX locale, a word consists of a maximal sequence of letters, digits, and underscores, delimited at both ends by characters other than letters, digits, or underscores, or by the beginning or end of a line or the edit buffer.

When accepting input characters from the user, in either *ex* command mode or *ex* text input mode, *ex* shall enable canonical mode input processing, as defined in the System Interfaces volume of IEEE Std 1003.1-2001.

If in *ex* text input mode:

- 1. If the **number** edit option is set, *ex* shall prompt for input using the line number that would be assigned to the line if it is entered, in the format specified for the *ex* **number** command.
- 2. If the **autoindent** edit option is set, *ex* shall prompt for input using **autoindent** characters, as described by the **autoindent** edit option. **autoindent** characters shall follow the line number, if any.

If in *ex* command mode:

- 1. If the **prompt** edit option is set, input shall be prompted for using a single ':' character; otherwise, there shall be no prompt.
- The input characters in the following sections shall have the following effects on the input line.

14181 Scroll

14152

14153

14155

14156

14157

14158

14159

14160

14161

14162

14163 14164

14165

14166 14167

14168

14169

14170

14171 14172

14173

14174

14175 14176

14177 14178

- 14182 *Synopsis*: eof
- See the description of the *stty eof* character in *stty*.
- 14184 If in *ex* command mode:
- If the *eof* character is the first character entered on the line, the line shall be evaluated as if it contained two characters: a <control>-D and a <newline>.
- 14187 Otherwise, the *eof* character shall have no special meaning.

14188 If in *ex* text input mode: If the cursor follows an autoindent character, the autoindent characters in the line shall be 14189 14190 modified so that a part of the next text input character will be displayed on the first column in the line after the previous **shiftwidth** edit option column boundary, and the user shall be 14191 prompted again for input for the same line. 14192 Otherwise, if the cursor follows a '0', which follows an autoindent character, and the '0' 14193 was the previous text input character, the '0' and all autoindent characters in the line shall 14194 be discarded, and the user shall be prompted again for input for the same line. 14195 Otherwise, if the cursor follows a '^', which follows an **autoindent** character, and the '^' 14196 was the previous text input character, the ' ^ ' and all **autoindent** characters in the line shall 14197 be discarded, and the user shall be prompted again for input for the same line. In addition, 14198 the autoindent level for the next input line shall be derived from the same line from which 14199 the **autoindent** level for the current input line was derived. 14200 Otherwise, if there are no **autoindent** or text input characters in the line, the *eof* character 14201 14202 shall be discarded. 14203 Otherwise, the *eof* character shall have no special meaning. 14204 <newline> Synopsis: <newline> 14205 14206 <control>-J 14207 If in *ex* command mode: Cause the command line to be parsed; <control>-J shall be mapped to the <newline> for this 14208 14209 purpose. If in *ex* text input mode: 14210 Terminate the current line. If there are no characters other than autoindent characters on the 14211 14212 line, all characters on the line shall be discarded. Prompt for text input on a new line after the current line. If the **autoindent** edit option is set, 14213 an appropriate number of autoindent characters shall be added as a prefix to the line as 14214 14215 described by the ex autoindent edit option. <base>backslash> 14216 Synopsis: 14217 <backslash> Allow the entry of a subsequent <newline> or <control>-J as a literal character, removing any 14218 special meaning that it may have to the editor during text input mode. The backslash character 14219 14220 shall be retained and evaluated when the command line is parsed, or retained and included

when the input text becomes part of the edit buffer.

14222	<control>-V</control>	I
14223	Synopsis:	<pre><control>-V</control></pre>
14224	-	ntry of any subsequent character as a literal character, removing any special meaning
14225		have to the editor during text input mode. The <control>-V character shall be</control>
14226		efore the command line is parsed or the input text becomes part of the edit buffer.
14227	If the "liter:	al next'' functionality is performed by the underlying system, it is implementation-
14228		ether a character other than <control>-V performs this function.</control>
		•
14229	<control>-V</control>	V
14230	Synopsis:	<control>-W</control>
14231		<control>-W, and the word previous to it in the input line, including any <blank>s</blank></control>
14232		he word and preceding the <control>-W. If the "word erase" functionality is</control>
14233		by the underlying system, it is implementation-defined whether a character other ol>-W performs this function.
14234	tilali <toliti< td=""><td>bi>-w performs this function.</td></toliti<>	bi>-w performs this function.
14235	Command 1	Descriptions in ex
14236	The following	ng symbols are used in this section to represent command modifiers. Some of these
14237	modifiers ca	an be omitted, in which case the specified defaults shall be used.
14238	1addr	A single line address, given in any of the forms described in Addressing in ex (on
14239		page 361); the default shall be the current line ($^{\prime}$. $^{\prime}$), unless otherwise specified.
14240		If the line address is zero, it shall be an error, unless otherwise specified in the
14241		following command descriptions.
14242		If the edit buffer is empty, and the address is specified with a command other than
14243		=, append, insert, open, put, read, or visual, or the address is not zero, it shall be
14244		an error.
14245	2addr	Two addresses specifying an inclusive range of lines. If no addresses are specified,
14246		the default for <i>2addr</i> shall be the current line only (".,."), unless otherwise
14247		specified in the following command descriptions. If one address is specified, <i>2addr</i>
14248 14249		shall specify that line only, unless otherwise specified in the following command descriptions.
14243		It shall be an error if the first address is greater than the second address.
		<u> </u>
14251		If the edit buffer is empty, and the two addresses are specified with a command
14252 14253		other than the !, write, wq, or xit commands, or either address is not zero, it shall be an error.
14254 14255	count	A positive decimal number. If <i>count</i> is specified, it shall be equivalent to specifying an additional address to the command, unless otherwise specified by the following
14256		command descriptions. The additional address shall be equal to the last address
14257		specified to the command (either explicitly or by default) plus <i>count</i> -1.
14258		If this would result in an address greater than the last line of the edit buffer, it shall
14259		be corrected to equal the last line of the edit buffer.
14260	flags	One or more of the characters $'+'$, $'-'$, $'\#'$, $'p'$, or $'l'$ (ell). The flag characters
14261		can be <blank>-separated, and in any order or combination. The characters '#',</blank>
14262		'p', and 'l' shall cause lines to be written in the format specified by the print
14263		command with the specified <i>flags</i> .

14264

The lines to be written are as follows:

14265 14266 1. All edit buffer lines written during the execution of the ex &, ~, list, number, open, print, s, visual, and z commands shall be written as specified by flags.

14267 14268 14269 After the completion of an ex command with a flag as an argument, the current line shall be written as specified by flags, unless the current line was the last line written by the command.

14270 14271 14272 The characters '+' and '-' cause the value of the current line after the execution of the ex command to be adjusted by the offset address as described in **Addressing** in ex (on page 361). This adjustment shall occur before the current line is written as described in 2. above.

14273 14274

14276

14277

14278

14279

The default for *flags* shall be none.

14275

buffer

One of a number of named areas for holding text. The named buffers are specified by the alphanumeric characters of the POSIX locale. There shall also be one "unnamed" buffer. When no buffer is specified for editor commands that use a buffer, the unnamed buffer shall be used. Commands that store text into buffers shall store the text as it was before the command took effect, and shall store text occurring earlier in the file before text occurring later in the file, regardless of how the text region was specified. Commands that store text into buffers shall store the text into the unnamed buffer as well as any specified buffer.

14280 14281 14282

> In ex commands, buffer names are specified as the name by itself. In open or visual mode commands the name is preceded by a double quote (' "') character.

14283 14284

> If the specified buffer name is an uppercase character, and the buffer contents are to be modified, the buffer shall be appended to rather than being overwritten. If the buffer is not being modified, specifying the buffer name in lowercase and

14285 14286 14287

uppercase shall have identical results.

14288 14289 14290

14291

14292

There shall also be buffers named by the numbers 1 through 9. In open and visual mode, if a region of text including characters from more than a single line is being modified by the vi c or d commands, the motion character associated with the c or d commands specifies that the buffer text shall be in line mode, or the commands %, ', /, ?, (,), N, n, $\{$, or $\}$ are used to define a region of text for the c or d commands, the contents of buffers 1 through 8 shall be moved into the buffer named by the next numerically greater value, the contents of buffer 9 shall be discarded, and the region of text shall be copied into buffer 1. This shall be in addition to copying the text into a user-specified buffer or unnamed buffer, or both. Numeric buffers can be specified as a source buffer for open and visual mode commands; however, specifying a numeric buffer as the write target of an open or visual mode

14297

14298

command shall have unspecified results.

14299 14300 14301

The text of each buffer shall have the characteristic of being in either line or character mode. Appending text to a non-empty buffer shall set the mode to match the characteristic of the text being appended. Appending text to a buffer shall cause the creation of at least one additional line in the buffer. All text stored into buffers by ex commands shall be in line mode. The ex commands that use buffers as the source of text specify individually how buffers of different modes are handled. Each open or visual mode command that uses buffers for any purpose specifies individually the mode of the text stored into the buffer and how buffers

14306 14307 14308

14309

of different modes are handled.

14310	file	Command text used to derive a pathname. The default shall be the current
14311		pathname, as defined previously, in which case, if no current pathname has yet
14312		been established it shall be an error, except where specifically noted in the
14313		individual command descriptions that follow. If the command text contains any of
14314		the characters $\prime \sim \prime$, $\prime \{\prime, \prime [\prime, \prime *\prime, \prime ?\prime, \prime *\prime, \prime \prime, \prime \prime, \prime \prime, \prime \prime,$
14315		subjected to the process of "shell expansions", as described below; if more than a
14316		single pathname results and the command expects only one, it shall be an error.
14317		The process of shell expansions in the editor shall be done as follows. The ex utility
14318		shall pass two arguments to the program named by the shell edit option; the first
14319		shall be -c, and the second shall be the string "echo" and the command text as a
14320		single argument. The standard output and standard error of that command shall
14321		replace the command text.
14021		replace the communicatext.
14322	!	A character that can be appended to the command name to modify its operation,
14323		as detailed in the individual command descriptions. With the exception of the ex
14324		read , write, and! commands, the '!' character shall only act as a modifier if there
14325		are no <blank>s between it and the command name.</blank>
14326	remember	red search direction

The vi commands N and n begin searching in a forwards or backwards direction in the edit buffer based on a remembered search direction, which is initially unset, and is set by the ex **global**, v, s, and tag commands, and the vi and tag commands.

Abbreviate

Synopsis: ab[breviate][lhs rhs]

If *lhs* and *rhs* are not specified, write the current list of abbreviations and do nothing more.

Implementations may restrict the set of characters accepted in *lhs* or *rh*, except that printable characters and <blank>s shall not be restricted. Additional restrictions shall be implementation-defined.

In both *lhs* and *rhs*, any character may be escaped with a <control>-V, in which case the character shall not be used to delimit *lhs* from *rhs*, and the escaping <control>-V shall be discarded.

In open and visual text input mode, if a non-word or <ESC> character that is not escaped by a <control>-V character is entered after a word character, a check shall be made for a set of characters matching *lhs*, in the text input entered during this command. If it is found, the effect shall be as if *rhs* was entered instead of *lhs*.

The set of characters that are checked is defined as follows:

- 1. If there are no characters inserted before the word and non-word or <ESC> characters that triggered the check, the set of characters shall consist of the word character.
- 2. If the character inserted before the word and non-word or <ESC> characters that triggered the check is a word character, the set of characters shall consist of the characters inserted immediately before the triggering characters that are word characters, plus the triggering word character.
- 3. If the character inserted before the word and non-word or <ESC> characters that triggered the check is not a word character, the set of characters shall consist of the characters that were inserted before the triggering characters that are neither <blank>s nor word characters, plus the triggering word character.

Utilities $\mathbf{e}\mathbf{x}$

14354 14355 14356	It is unspecified whether the <i>lhs</i> argument entered for the <i>ex</i> abbreviate and unabbreviate commands is replaced in this fashion. Regardless of whether or not the replacement occurs, the effect of the command shall be as if the replacement had not occurred.
14357	Current line: Unchanged.
14358	Current column: Unchanged.
14359	Append
14360	Synopsis: [1addr] a[ppend][!]
14361 14362	Enter <i>ex</i> text input mode; the input text shall be placed after the specified line. If line zero is specified, the text shall be placed at the beginning of the edit buffer.
14363 14364 14365	This command shall be affected by the number and autoindent edit options; following the command name with '!' shall cause the autoindent edit option setting to be toggled for the duration of this command only.
14366 14367	<i>Current line</i> : Set to the last input line; if no lines were input, set to the specified line, or to the first line of the edit buffer if a line of zero was specified, or zero if the edit buffer is empty.
14368	Current column: Set to non- <blank>.</blank>
14369	Arguments
14370	Synopsis: ar[gs]
14371 14372	Write the current argument list, with the current argument-list entry, if any, between $'$ [$'$ and $'$] $'$ characters.
14373	Current line: Unchanged.
14374	Current column: Unchanged.
14375	Change
14376	Synopsis: [2addr] c[hange][!][count]
14377 14378	Enter <i>ex</i> text input mode; the input text shall replace the specified lines. The specified lines shall be copied into the unnamed buffer, which shall become a line mode buffer.
14379 14380	This command shall be affected by the number and autoindent edit options; following the command name with '!' shall cause the autoindent edit option setting to be toggled for the duration of this command only.
14381	duration of this command only.
14381 14382 14383 14384	Current line: Set to the last input line; if no lines were input, set to the line before the first address, or to the first line of the edit buffer if there are no lines preceding the first address, or to zero if the edit buffer is empty.

14386	Change Directory
14387 14388	Synopsis: chd[ir][!][directory] cd[!][directory]
14389	Change the current working directory to directory.
14390 14391 14392 14393	If no <i>directory</i> argument is specified, and the <i>HOME</i> environment variable is set to a non-null and non-empty value, <i>directory</i> shall default to the value named in the <i>HOME</i> environment variable. If the <i>HOME</i> environment variable is empty or is undefined, the default value of <i>directory</i> is implementation-defined.
14394 14395	If no $'$! ' is appended to the command name, and the edit buffer has been modified since the last complete write, and the current pathname does not begin with a $'$ /', it shall be an error.
14396	Current line: Unchanged.
14397	Current column: Unchanged.
14398	Сору
14399 14400	Synopsis: [2addr] co[py] 1addr [flags] [2addr] t 1addr [flags]
14401 14402	Copy the specified lines after the specified destination line; line zero specifies that the lines shall be placed at the beginning of the edit buffer.
14403	Current line: Set to the last line copied.
14404	Current column: Set to non- <blank>.</blank>
14405	Delete
14406	Synopsis: [2addr] d[elete][buffer][count][flags]
14407 14408	Delete the specified lines into a buffer (defaulting to the unnamed buffer), which shall become a line-mode buffer.
14409 14410	Flags can immediately follow the command name; see Command Line Parsing in ex (on page 362).
14411 14412	<i>Current line</i> : Set to the line following the deleted lines, or to the last line in the edit buffer if that line is past the end of the edit buffer, or to zero if the edit buffer is empty.
14413	Current column: Set to non- <blank>.</blank>
14414	Edit
14415 14416	Synopsis: e[dit][!][+command][file] ex[!][+command][file]
14417 14418	If no '!' is appended to the command name, and the edit buffer has been modified since the last complete write, it shall be an error.
14419 14420 14421 14422	If <i>file</i> is specified, replace the current contents of the edit buffer with the current contents of <i>file</i> , and set the current pathname to <i>file</i> . If <i>file</i> is not specified, replace the current contents of the edit buffer with the current contents of the file named by the current pathname. If for any reason the current contents of the file cannot be accessed, the edit buffer shall be empty.
14423 14424 14425	The $+command$ option shall be $<$ blank $>$ -delimited; $<$ blank $>$ s within $+command$ can be escaped by preceding them with a backslash character. The $+command$ shall be interpreted as an ex command immediately after the contents of the edit buffer have been replaced and the current

Utilities $\mathbf{e}\mathbf{x}$

14426	line and column have been set.
14427	If the edit buffer is empty:
14428	Current line: Set to 0.
14429	Current column: Set to 1.
14430	Otherwise, if executed while in <i>ex</i> command mode or if the + <i>command</i> argument is specified:
14431	Current line: Set to the last line of the edit buffer.
14432	Current column: Set to non- <blank>.</blank>
14433	Otherwise, if <i>file</i> is omitted or results in the current pathname:
14434	Current line: Set to the first line of the edit buffer.
14435	Current column: Set to non- <blank>.</blank>
14436 14437	Otherwise, if <i>file</i> is the same as the last file edited, the line and column shall be set as follows; if the file was previously edited, the line and column may be set as follows:
14438 14439	<i>Current line</i> : Set to the last value held when that file was last edited. If this value is not a valid line in the new edit buffer, set to the first line of the edit buffer.
14440 14441 14442	<i>Current column</i> : If the current line was set to the last value held when the file was last edited, set to the last value held when the file was last edited. Otherwise, or if the last value is not a valid column in the new edit buffer, set to non- Value Column C
14443	Otherwise:
14444	Current line: Set to the first line of the edit buffer.
14445	Current column: Set to non- <blank>.</blank>
14446	File
14447	Synopsis: f[ile][file]
14448 14449	If a <i>file</i> argument is specified, the alternate pathname shall be set to the current pathname, and the current pathname shall be set to <i>file</i> .
14450 14451 14452 14453 14454 14455 14456	Write an informational message. If the file has a current pathname, it shall be included in this message; otherwise, the message shall indicate that there is no current pathname. If the edit buffer contains lines, the current line number and the number of lines in the edit buffer shall be included in this message; otherwise, the message shall indicate that the edit buffer is empty. If the edit buffer has been modified since the last complete write, this fact shall be included in this message. If the readonly edit option is set, this fact shall be included in this message may contain other unspecified information.
14457	Current line: Unchanged.
14458	Current column: Unchanged.

14459 Global

14460 Synopsis: [2addr] g[lobal] /pattern/ [commands]

14461 [2addr] v /pattern/ [commands]

The optional '!' character after the **global** command shall be the same as executing the **v** command.

If *pattern* is empty (for example, "//") or not specified, the last regular expression used in the editor command shall be used as the *pattern*. The *pattern* can be delimited by slashes (shown in the Synopsis), as well as any non-alphanumeric or non-
blank> other than backslash, vertical line, double quote, or <newline>.

If no lines are specified, the lines shall default to the entire file.

The **global** and **v** commands are logically two-pass operations. First, mark the lines within the specified lines for which the line excluding the terminating <newline> matches (**global**) or does not match (**v** or **global**!) the specified pattern. Second, execute the *ex* commands given by *commands*, with the current line (' . ') set to each marked line. If an error occurs during this process, or the contents of the edit buffer are replaced (for example, by the *ex* :edit command) an error message shall be written and no more commands resulting from the execution of this command shall be processed.

Multiple *ex* commands can be specified by entering multiple commands on a single line using a vertical line to delimit them, or one per line, by escaping each <newline> with a backslash.

If no commands are specified:

- 1. If in ex command mode, it shall be as if the **print** command were specified.
- 2. Otherwise, no command shall be executed.

For the **append**, **change**, and **insert** commands, the input text shall be included as part of the command, and the terminating period can be omitted if the command ends the list of commands. The **open** and **visual** commands can be specified as one of the commands, in which case each marked line shall cause the editor to enter open or visual mode. If open or visual mode is exited using the $vi \ Q$ command, the current line shall be set to the next marked line, and open or visual mode reentered, until the list of marked lines is exhausted.

The **global**, **v**, and **undo** commands cannot be used in *commands*. Marked lines may be deleted by commands executed for lines occurring earlier in the file than the marked lines. In this case, no commands shall be executed for the deleted lines.

If the remembered search direction is not set, the **global** and **v** commands shall set it to forward.

The **autoprint** and **autoindent** edit options shall be inhibited for the duration of the **g** or **v** command.

Current line: If no commands executed, set to the last marked line. Otherwise, as specified for the executed *ex* commands.

Current column: If no commands are executed, set to non-
blank>; otherwise, as specified for the individual ex commands.

14497	Insert
14498	Synopsis: [laddr] i[nsert][!]
14499 14500	Enter <i>ex</i> text input mode; the input text shall be placed before the specified line. If the line is zero or 1, the text shall be placed at the beginning of the edit buffer.
14501 14502 14503	This command shall be affected by the number and autoindent edit options; following the command name with '!' shall cause the autoindent edit option setting to be toggled for the duration of this command only.
14504 14505 14506	<i>Current line</i> : Set to the last input line; if no lines were input, set to the line before the specified line, or to the first line of the edit buffer if there are no lines preceding the specified line, or zero if the edit buffer is empty.
14507	Current column: Set to non- <blank>.</blank>
14508	Join
14509	Synopsis: [2addr] j[oin][!][count][flags]
14510	If count is specified:
14511 14512	If no address was specified, the join command shall behave as if $2addr$ were the current line and the current line plus $count(., + count)$.
14513 14514	If one address was specified, the join command shall behave as if $2addr$ were the specified address and the specified address plus $count$ ($addr$, $addr$ + $count$).
14515 14516	If two addresses were specified, the join command shall behave as if an additional address, equal to the last address plus $count - 1$ ($addr1, addr2, addr2 + count - 1$), was specified.
14517 14518	If this would result in a second address greater than the last line of the edit buffer, it shall be corrected to be equal to the last line of the edit buffer.
14519	If no <i>count</i> is specified:
14520 14521	If no address was specified, the join command shall behave as if $2addr$ were the current line and the next line $(.,.+1)$.
14522 14523	If one address was specified, the join command shall behave as if $2addr$ were the specified address and the next line $(addr, addr + 1)$.
14524 14525	Join the text from the specified lines together into a single line, which shall replace the specified lines.
14526 14527	If a '!' character is appended to the command name, the join shall be without modification of any line, independent of the current locale.
14528 14529	Otherwise, in the POSIX locale, set the current line to the first of the specified lines, and then, for each subsequent line, proceed as follows:
14530	1. Discard leading <space>s from the line to be joined.</space>
14531	2. If the line to be joined is now empty, delete it, and skip steps 3 through 5.
14532	3. If the current line ends in a <blank>, or the first character of the line to be joined is a ')' character join the lines without further modification.</blank>

character, join the lines without further modification.

4. If the last character of the current line is a $^\prime$. $^\prime$, join the lines with two <space>s between

1453314534

14535

them.

14536 5. Otherwise, join the lines with a single <space> between them. *Current line*: Set to the first line specified. 14537 Current column: Set to non-

- slank>. 14538 14539 List Synopsis: [2addr] l[ist][count][flags] 14540 This command shall be equivalent to the *ex* command: 14541 14542 [2addr] p[rint][count] l[flags] See **Print** (on page 380). 14543 14544 Map Synopsis: 14545 map[!][lhs rhs] If *lhs* and *rhs* are not specified: 14546 14547 1. If '!' is specified, write the current list of text input mode maps. Otherwise, write the current list of command mode maps. 14548 3. Do nothing more. 14549 Implementations may restrict the set of characters accepted in *lhs* or *rhs*, except that printable 14550 14551 characters and
 shall not be restricted. Additional restrictions shall be implementationdefined. In both *lhs* and *rhs*, any character can be escaped with a <control>-V, in which case the 14552 14553 character shall not be used to delimit lhs from rhs, and the escaping <control>-V shall be discarded. 14554 14555 If the character '!' is appended to the **map** command name, the mapping shall be effective during open or visual text input mode rather than open or visual command mode. This allows 14556 14557 *lhs* to have two different **map** definitions at the same time: one for command mode and one for text input mode. 14558 14559 For command mode mappings: When the *lhs* is entered as any part of a *vi* command in open or visual mode (but not as part 14560 of the arguments to the command), the action shall be as if the corresponding *rhs* had been 14561 entered. 14562 If any character in the command, other than the first, is escaped using a <control>-V 14563 character, that character shall not be part of a match to an *lhs*. 14564 It is unspecified whether implementations shall support map commands where the lhs is 14565 14566 more than a single character in length, where the first character of the *lhs* is printable. If *lhs* contains more than one character and the first character is '#', followed by a sequence 14567 of digits corresponding to a numbered function key, then when this function key is typed it 14568 shall be mapped to rhs. Characters other than digits following a '#' character also represent 14569

For text input mode mappings:

14570

14571

14572

the function key named by the characters in the *lhs* following the '#' and may be mapped to

rhs. It is unspecified how function keys are named or what function keys are supported.

14573 When the *lhs* is entered as any part of text entered in open or visual text input modes, the action shall be as if the corresponding *rhs* had been entered. 14574 14575 If any character in the input text is escaped using a <control>-V character, that character shall not be part of a match to an *lhs*. 14576 14577 It is unspecified whether the *lhs* text entered for subsequent map or unmap commands is replaced with the rhs text for the purposes of the screen display; regardless of whether or not 14578 the display appears as if the corresponding rhs text was entered, the effect of the command 14579 shall be as if the *lhs* text was entered. 14580 If only part of the *lhs* is entered, it is unspecified how long the editor will wait for additional, 14581 possibly matching characters before treating the already entered characters as not matching the 14582 14583 14584 The *rhs* characters shall themselves be subject to remapping, unless otherwise specified by the remap edit option, except that if the characters in *lhs* occur as prefix characters in *rhs*, those 14585 characters shall not be remapped. 14586 On block-mode terminals, the mapping need not occur immediately (for example, it may occur 14587 14588 after the terminal transmits a group of characters to the system), but it shall achieve the same results as if it occurred immediately. 14589 Current line: Unchanged. 14590 Current column: Unchanged. 14591 Mark 14592 Synopsis: [laddr] ma[rk] character 14593 [laddr] k character 14594 14595 Implementations shall support *character* values of a single lowercase letter of the POSIX locale and the characters ''' and '''; support of other characters is implementation-defined. 14596 14597 If executing the vi m command, set the specified mark to the current line and 1-based numbered character referenced by the current column, if any; otherwise, column position 1. 14598 Otherwise, set the specified mark to the specified line and 1-based numbered first non-

values of the specified mark to the specified line and 1-based numbered first non-

values of the specified mark to the specified line and 1-based numbered first non-

values of the specified mark to the specified line and 1-based numbered first non-

values of the specified mark to the specified line and 1-based numbered first non-

values of the specified mark to the specified line and 1-based numbered first non-

values of the specified mark to the specified line and 1-based numbered first non-

values of the specified mark to the specified line and 1-based numbered first non-

values of the specified mark to the specified line and 1-based numbered first non-

values of the specified mark to the specified line and 1-based numbered first non-

values of the specified mark to the specified line and 1-based numbered first non-

values of the specified mark to the specified line and 1-based numbered nu 14599 non-<newline> in the line, if any; otherwise, the last non-<newline> in the line, if any; otherwise, 14600 column position 1. 14601 The mark shall remain associated with the line until the mark is reset or the line is deleted. If a 14602 deleted line is restored by a subsequent **undo** command, any marks previously associated with 14603 the line, which have not been reset, shall be restored as well. Any use of a mark not associated 14604 with a current line in the edit buffer shall be an error. 14605 14606 The marks 'and' shall be set as described previously, immediately before the following events occur in the editor: 14607 1. The use of '\$' as an ex address 14608 The use of a positive decimal number as an ex address 14609

The use of the following open and visual mode commands: <control>-], %, (,), [,], {, }

The use of the following open and visual mode commands: ', G, H, L, M, z if the current

The use of a search command as an ex address

The use of a mark reference as an ex address

line will change as a result of the command

14610

14611

14612 14613

7. The use of the open and visual mode commands: /, ?, N, ', n if the current line or column will change as a result of the command

8. The use of the ex mode commands: z, undo, global, v

For rules 1., 2., 3., and 4., the 'and' marks shall not be set if the *ex* command is parsed as specified by rule 6.a. in **Command Line Parsing in ex** (on page 362).

For rules 5., 6., and 7., the 'and' marks shall not be set if the commands are used as motion commands in open and visual mode.

For rules 1., 2., 3., 4., 5., 6., 7., and 8., the 'and' marks shall not be set if the command fails.

The 'and ' marks shall be set as described previously, each time the contents of the edit buffer are replaced (including the editing of the initial buffer), if in open or visual mode, or if in **ex** mode and the edit buffer is not empty, before any commands or movements (including commands or movements specified by the –c or –t options or the +*command* argument) are executed on the edit buffer. If in open or visual mode, the marks shall be set as if executing the *vi* **m** command; otherwise, as if executing the *ex* **mark** command.

When changing from **ex** mode to open or visual mode, if the 'and 'marks are not already set, the 'and 'marks shall be set as described previously.

Current line: Unchanged.

Current column: Unchanged.

Move

14634 Synopsis: [2addr] m[ove] 1addr [flags]

Move the specified lines after the specified destination line. A destination of line zero specifies that the lines shall be placed at the beginning of the edit buffer. It shall be an error if the destination line is within the range of lines to be moved.

Current line: Set to the last of the moved lines.

14639 Current column: Set to non-

- slank>.

Next

Synopsis: n[ext][!][+command][file ...]

If no '!' is appended to the command name, and the edit buffer has been modified since the last complete write, it shall be an error, unless the file is successfully written as specified by the **autowrite** option.

If one or more files is specified:

- 1. Set the argument list to the specified filenames.
- 2. Set the current argument list reference to be the first entry in the argument list.
- 14648 3. Set the current pathname to the first filename specified.

14649 Otherwise:

- 1. It shall be an error if there are no more filenames in the argument list after the filename currently referenced.
- 2. Set the current pathname and the current argument list reference to the filename after the filename currently referenced in the argument list.

14654 Replace the contents of the edit buffer with the contents of the file named by the current 14655 pathname. If for any reason the contents of the file cannot be accessed, the edit buffer shall be 14656 empty. This command shall be affected by the **autowrite** and **writeany** edit options. 14657 14658 The +command option shall be <blank>-delimited; <blank>s can be escaped by preceding them with a backslash character. The +command shall be interpreted as an ex command immediately 14659 after the contents of the edit buffer have been replaced and the current line and column have 14660 14661 Current line: Set as described for the **edit** command. 14662 Current column: Set as described for the **edit** command. 14663 Number 14664 Synopsis: [2addr] nu[mber][count][flags] 14665 14666 [2addr] #[count][flags] These commands shall be equivalent to the *ex* command: 14667 14668 [2addr] p[rint][count] #[flags] See **Print** (on page 380). 14669 14670 Open [laddr] o[pen] /pattern/ [flags] Synopsis: 14671 14672 This command need not be supported on block-mode terminals or terminals with insufficient capabilities. If standard input, standard output, or standard error are not terminal devices, the 14673 results are unspecified. 14674 14675 Enter open mode. 14676 The trailing delimiter can be omitted from *pattern* at the end of the command line. If *pattern* is empty (for example, "//") or not specified, the last regular expression used in the editor shall be 14677 used as the pattern. The pattern can be delimited by slashes (shown in the Synopsis), as well as 14678 any alphanumeric, or non-

-| other than backslash, vertical line, double quote, or 14679 <newline>. 14680 *Current line*: Set to the specified line. 14681 14682 Current column: Set to non-

- slank>. **Preserve** 14683 14684 Synopsis: pre[serve] Save the edit buffer in a form that can later be recovered by using the **-r** option or by using the *ex* 14685 recover command. After the file has been preserved, a mail message shall be sent to the user. 14686 This message shall be readable by invoking the *mailx* utility. The message shall contain the name 14687 14688 of the file, the time of preservation, and an ex command that could be used to recover the file. Additional information may be included in the mail message. 14689 Current line: Unchanged. 14690 14691 Current column: Unchanged.

Print

14693 Synopsis: [2addr] p[rint] [count] [flags]

Write the addressed lines. The behavior is unspecified if the number of columns on the display is less than the number of columns required to write any single character in the lines being written.

Non-printable characters, except for the <tab>, shall be written as implementation-defined multi-character sequences.

If the # flag is specified or the **number** edit option is set, each line shall be preceded by its line number in the following format:

"%6d $\Delta\Delta$ ", <line number>

If the I flag is specified or the list edit option is set:

- 1. The characters listed in the Base Definitions volume of IEEE Std 1003.1-2001, Table 5-1, Escape Sequences and Associated Actions shall be written as the corresponding escape sequence.
- 2. Non-printable characters not in the Base Definitions volume of IEEE Std 1003.1-2001, Table 5-1, Escape Sequences and Associated Actions shall be written as one three-digit octal number (with a preceding backslash) for each byte in the character (most significant byte first). If the size of a byte on the system is greater than 9 bits, the format used for non-printable characters is implementation-defined.
- 3. The end of each line shall be marked with a '\$', and literal '\$' characters within the line shall be written with a preceding backslash.

Long lines shall be folded; the length at which folding occurs is unspecified, but should be appropriate for the output terminal, considering the number of columns of the terminal.

If a line is folded, and the **l** flag is not specified and the **list** edit option is not set, it is unspecified whether a multi-column character at the folding position is separated; it shall not be discarded.

Current line: Set to the last written line.

Current column: Unchanged if the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise, set to non-

Value of the current line is unchanged; otherwise is unchanged; otherw

Put

14719 Synopsis: [laddr] pu[t][buffer]

Append text from the specified buffer (by default, the unnamed buffer) to the specified line; line zero specifies that the text shall be placed at the beginning of the edit buffer. Each portion of a line in the buffer shall become a new line in the edit buffer, regardless of the mode of the buffer.

Current line: Set to the last line entered into the edit buffer.

14724 Current column: Set to non-

Clark>.

Quit

Synopsis: q[uit][!]

14727 If no '!' is appended to the command name:

- 1. If the edit buffer has been modified since the last complete write, it shall be an error.
- 2. If there are filenames in the argument list after the filename currently referenced, and the last command was not a **quit**, **wq**, **xit**, or **ZZ** (see **Exit** (on page 1022)) command, it shall be an error.

14732 Otherwise, terminate the editing session. 14733 Read [1addr] r[ead][!][file] 14734 Synopsis: 14735 If '!' is not the first non-

-

slank> to follow the command name, a copy of the specified file shall be appended into the edit buffer after the specified line; line zero specifies that the copy shall be 14736 placed at the beginning of the edit buffer. The number of lines and bytes read shall be written. If 14737 no file is named, the current pathname shall be the default. If there is no current pathname, then 14738 14739 file shall become the current pathname. If there is no current pathname or file operand, it shall be 14740 an error. Specifying a *file* that is not of type regular shall have unspecified results. Otherwise, if file is preceded by '!', the rest of the line after the '!' shall have '%', '#', and 14741 14742 '!' characters expanded as described in **Command Line Parsing in ex** (on page 362). The ex utility shall then pass two arguments to the program named by the shell edit option; the 14743 first shall be -c and the second shall be the expanded arguments to the read command as a 14744 14745 single argument. The standard input of the program shall be set to the standard input of the ex program when it was invoked. The standard error and standard output of the program shall be 14746 appended into the edit buffer after the specified line. 14747 Each line in the copied file or program output (as delimited by <newline>s or the end of the file 14748 or output if it is not immediately preceded by a <newline>), shall be a separate line in the edit 14749 buffer. Any occurrences of <carriage-return> and <newline> pairs in the output shall be treated 14750 14751 as single <newline>s. 14752 The special meaning of the '!' following the **read** command can be overridden by escaping it 14753 with a backslash character. Current line: If no lines are added to the edit buffer, unchanged. Otherwise, if in open or visual 14754 14755 mode, set to the first line entered into the edit buffer. Otherwise, set to the last line entered into 14756 the edit buffer. 14757 Current column: Set to non-<blank>. 14758 Recover rec[over][!] file 14759 Synopsis: 14760 If no '!' is appended to the command name, and the edit buffer has been modified since the 14761 last complete write, it shall be an error. 14762 If no file operand is specified, then the current pathname shall be used. If there is no current 14763 pathname or *file* operand, it shall be an error. If no recovery information has previously been saved about file, the recover command shall 14764 behave identically to the edit command, and an informational message to this effect shall be 14765 written. 14766 Otherwise, set the current pathname to file, and replace the current contents of the edit buffer 14767 with the recovered contents of file. If there are multiple instances of the file to be recovered, the 14768 one most recently saved shall be recovered, and an informational message that there are 14769 previous versions of the file that can be recovered shall be written. The editor shall behave as if 14770 the contents of the edit buffer have already been modified. 14771

Current file: Set as described for the **edit** command.

Current column: Set as described for the **edit** command.

14779

14774	Rewind
14775	Synopsis: rew[ind][!]
14776	If no '!' is appended to the command name, and the edit buffer has been modified since the
14777	last complete write, it shall be an error, unless the file is successfully written as specified by the
14778	autowrite option.
14779	If the argument list is empty, it shall be an error.
14780 14781	The current argument list reference and the current pathname shall be set to the first filename in the argument list.
14782	Replace the contents of the edit buffer with the contents of the file named by the current
14782	pathname. If for any reason the contents of the file cannot be accessed, the edit buffer shall be
14784	empty.
14785	This command shall be affected by the autowrite and writeany edit options.
14786	Current line: Set as described for the edit command.
14787	Current column: Set as described for the edit command.
14788	Set
14789	Synopsis: se[t][option[=[value]]][nooption][option?][all]
14790	When no arguments are specified, write the value of the term edit option and those options
14791	whose values have been changed from the default settings; when the argument all is specified,
14792	write all of the option values.
14793	Giving an option name followed by the character '?' shall cause the current value of that
14794	option to be written. The '?' can be separated from the option name by zero or more <blank>s.</blank>
14795	The '?' shall be necessary only for Boolean valued options. Boolean options can be given values
14796	by the form set option to turn them on or set nooption to turn them off; string and numeric
14797	options can be assigned by the form set option=value. Any blank>s in strings can be included as it by proceeding each, which we assert the set of t
14798 14799	as is by preceding each <blank> with an escaping backslash. More than one option can be set or listed by a single set command by specifying multiple arguments, each separated from the next</blank>
14800	by one or more blank>s.
14801	See Edit Options in ex (on page 392) for details about specific options.
14802	Current line: Unchanged.
14803	Current column: Unchanged.
14804	Shell
14805	Synopsis: sh[ell]
14806	Invoke the program named in the shell edit option with the single argument -i (interactive
14807	mode). Editing shall be resumed when the program exits.
14808	Current line: Unchanged.

14809

Current column: Unchanged.

14810	Source
14811	Synopsis: so[urce] file
14812	Read and execute <i>ex</i> commands from <i>file</i> . Lines in the file that are blank lines shall be ignored.
14813	Current line: As specified for the individual ex commands.
14814	Current column: As specified for the individual ex commands.
14815	Substitute
14816 14817	Synopsis: [2addr] s[ubstitute][/pattern/repl/[options][count][flags]] [2addr] &[options][count][flags]]
14818	[2addr] ~ [options] [count] [flags]]
14819	Replace the first instance of the pattern pattern by the string repl on each specified line. (See
14820	Regular Expressions in ex (on page 391) and Replacement Strings in ex (on page 391).) Any
14821	non-alphabetic, non- - slank> delimiter other than $' \setminus '$, $' \mid '$, double quote, or < new - new limiter of the slank of th
14822 14823	used instead of '/'. Backslash characters can be used to escape delimiters, backslash characters, and other special characters.
14824	The trailing delimiter can be omitted from <i>pattern</i> or from <i>repl</i> at the end of the command line. If
14825	both pattern and repl are not specified or are empty (for example, "//"), the last s command
14826	shall be repeated. If only <i>pattern</i> is not specified or is empty, the last regular expression used in
14827	the editor shall be used as the pattern. If only <i>repl</i> is not specified or is empty, the pattern shall be
14828 14829	replaced by nothing. If the entire replacement pattern is ' $%$ ', the last replacement pattern to an s command shall be used.
14830 14831	Entering a <carriage-return> in <i>repl</i> (which requires an escaping backslash in <i>ex</i> mode and an escaping <control>-V in open or <i>vi</i> mode) shall split the line at that point, creating a new line in</control></carriage-return>
14832	the edit buffer. The <carriage-return> shall be discarded.</carriage-return>
14833	If $options$ includes the letter 'g' (global), all non-overlapping instances of the pattern in the line
14834	shall be replaced.
14835	If <i>options</i> includes the letter 'c' (confirm), then before each substitution the line shall be written;
14836	the written line shall reflect all previous substitutions. On the following line, <space>s shall be written beneath the characters from the line that are before the <i>pattern</i> to be replaced, and ' ^ '</space>
14837 14838	characters written beneath the characters included in the <i>pattern</i> to be replaced. The <i>ex</i> utility
14839	shall then wait for a response from the user. An affirmative response shall cause the substitution
14840	to be done, while any other input shall not make the substitution. An affirmative response shall
14841	consist of a line with the affirmative response (as defined by the current locale) at the beginning
14842	of the line. This line shall be subject to editing in the same way as the <i>ex</i> command line.
14843 14844	If interrupted (see the ASYNCHRONOUS EVENTS section), any modifications confirmed by the user shall be preserved in the edit buffer after the interrupt.
14845	If the remembered search direction is not set, the ${\bf s}$ command shall set it to forward.
14846	In the second Synopsis, the & command shall repeat the previous substitution, as if the &
14847	command were replaced by:
14848	s/pattern/repl/
14849	where <i>pattern</i> and <i>repl</i> are as specified in the previous s , &, or ~ command.
14850 14851	In the third Synopsis, the $\tilde{\ }$ command shall repeat the previous substitution, as if the $\tilde{\ }$ $\tilde{\ }$ were replaced by:

14852	s/pattern/repl/
14853 14854	where <i>pattern</i> shall be the last regular expression specified to the editor, and <i>repl</i> shall be from the previous substitution (including & and ~) command.
14855	These commands shall be affected by the $LC_MESSAGES$ environment variable.
14856 14857	Current line: Set to the last line in which a substitution occurred, or, unchanged if no substitution occurred.
14858	Current column: Set to non- - blank>.
14859	Suspend
14860 14861	Synopsis: su[spend][!] st[op][!]
14862 14863 14864	Allow control to return to the invoking process; ex shall suspend itself as if it had received the SIGTSTP signal. The suspension shall occur only if job control is enabled in the invoking shell (see the description of $set-\mathbf{m}$).
14865	These commands shall be affected by the autowrite and writeany edit options.
14866	The current susp character (see <i>stty</i>) shall be equivalent to the suspend command.
14867	Tag
14868	Synopsis: ta[g][!] tagstring
14869 14870	The results are unspecified if the format of a tags file is not as specified by the <i>ctags</i> utility (see <i>ctags</i>) description.
14871 14872 14873 14874 14875 14876	The tag command shall search for <i>tagstring</i> in the tag files referred to by the tag edit option, in the order they are specified, until a reference to <i>tagstring</i> is found. Files shall be searched from beginning to end. If no reference is found, it shall be an error and an error message to this effect shall be written. If the reference is not found, or if an error occurs while processing a file referred to in the tag edit option, it shall be an error, and an error message shall be written at the first occurrence of such an error.
14877 14878	Otherwise, if the tags file contained a pattern, the pattern shall be treated as a regular expression used in the editor; for example, for the purposes of the $\bf s$ command.
14879 14880 14881 14882 14883	If the <i>tagstring</i> is in a file with a different name than the current pathname, set the current pathname to the name of that file, and replace the contents of the edit buffer with the contents of that file. In this case, if no '!' is appended to the command name, and the edit buffer has been modified since the last complete write, it shall be an error, unless the file is successfully written as specified by the autowrite option.
14884	This command shall be affected by the autowrite, tag, taglength, and writeany edit options.
14885 14886 14887	<i>Current line</i> : If the tags file contained a line number, set to that line number. If the line number is larger than the last line in the edit buffer, an error message shall be written and the current line shall be set as specified for the edit command.
14888 14889 14890	If the tags file contained a pattern, set to the first occurrence of the pattern in the file. If no matching pattern is found, an error message shall be written and the current line shall be set as specified for the edit command.
14891 14892 14893	<i>Current column</i> : If the tags file contained a line-number reference and that line-number was not larger than the last line in the edit buffer, or if the tags file contained a pattern and that pattern was found, set to non- - Otherwise, set as specified for the edit command.

14894 Unabbreviate 14895 Synopsis: una[bbrev] 1hs If *lhs* is not an entry in the current list of abbreviations (see **Abbreviate** (on page 370)), it shall be 14896 14897 an error. Otherwise, delete *lhs* from the list of abbreviations. Current line: Unchanged. 14898 14899 Current column: Unchanged. Undo 14900 Synopsis: 14901 u [ndo] 14902 Reverse the changes made by the last command that modified the contents of the edit buffer, 14903 including undo. For this purpose, the global, v, open, and visual commands, and commands resulting from buffer executions and mapped character expansions, are considered single 14904 commands. 14905 If no action that can be undone preceded the **undo** command, it shall be an error. 14906 If the undo command restores lines that were marked, the mark shall also be restored unless it 14907 was reset subsequent to the deletion of the lines. 14908 Current line: 14909 1. If lines are added or changed in the file, set to the first line added or changed. 14910 Set to the line before the first line deleted, if it exists. 14911 14912 3. Set to 1 if the edit buffer is not empty. 4. Set to zero. 14913 Current column: Set to non-<blank>. 14914 14915 **Unmap** 14916 Synopsis: unm[ap][!] 1hs If '!' is appended to the command name, and if *lhs* is not an entry in the list of text input mode 14917 14918 map definitions, it shall be an error. Otherwise, delete lhs from the list of text input mode map 14919 definitions. 14920 If no '!' is appended to the command name, and if *lhs* is not an entry in the list of command 14921 mode map definitions, it shall be an error. Otherwise, delete *lhs* from the list of command mode map definitions. 14922 14923 Current line: Unchanged. Current column: Unchanged. 14924 Version 14925 14926 Synopsis: ve[rsion] Write a message containing version information for the editor. The format of the message is 14927 unspecified. 14928 Current line: Unchanged. 14929

Current column: Unchanged.

14931 Visual

14932 Synopsis: [1addr] vi[sual][type][count][flags]

If *ex* is currently in open or visual mode, the Synopsis and behavior of the visual command shall be the same as the **edit** command, as specified by **Edit** (on page 372).

Otherwise, this command need not be supported on block-mode terminals or terminals with insufficient capabilities. If standard input, standard output, or standard error are not terminal devices, the results are unspecified.

If *count* is specified, the value of the **window** edit option shall be set to *count* (as described in **window** (on page 398)). If the '^' type character was also specified, the **window** edit option shall be set before being used by the type character.

Enter visual mode. If *type* is not specified, it shall be as if a *type* of '+' was specified. The *type* shall cause the following effects:

- Place the beginning of the specified line at the top of the display.
- Place the end of the specified line at the bottom of the display.
- . Place the beginning of the specified line in the middle of the display.
- ^ If the specified line is less than or equal to the value of the **window** edit option, set the line to 1; otherwise, decrement the line by the value of the **window** edit option minus 1. Place the beginning of this line as close to the bottom of the displayed lines as possible, while still displaying the value of the **window** edit option number of lines.

Current line: Set to the specified line.

Current column: Set to non-

- slank>.

14952 Write

```
Synopsis: [2addr] w[rite][!][>>][file]
[2addr] w[rite][!][file]
[2addr] wq[!][>>][file]
```

If no lines are specified, the lines shall default to the entire file.

The command wq shall be equivalent to a write command followed by a quit command; wq! shall be equivalent to write! followed by quit. In both cases, if the write command fails, the quit shall not be attempted.

If the command name is not followed by one or more <blank>s, or *file* is not preceded by a '!' character, the **write** shall be to a file.

- 1. If the >> argument is specified, and the file already exists, the lines shall be appended to the file instead of replacing its contents. If the >> argument is specified, and the file does not already exist, it is unspecified whether the write shall proceed as if the >> argument had not been specified or if the write shall fail.
- 2. If the **readonly** edit option is set (see **readonly** (on page 395)), the **write** shall fail.
- 3. If *file* is specified, and is not the current pathname, and the file exists, the **write** shall fail.
- 4. If *file* is not specified, the current pathname shall be used. If there is no current pathname, the **write** command shall fail.
- 5. If the current pathname is used, and the current pathname has been changed by the **file** or **read** commands, and the file exists, the **write** shall fail. If the **write** is successful,

14972 subsequent writes shall not fail for this reason (unless the current pathname is changed 14973 again). 14974 6. If the whole edit buffer is not being written, and the file to be written exists, the write shall 14975 14976 For rules 1., 2., 4., and 5., the write can be forced by appending the character '!' to the command name. 14977 14978 For rules 2., 4., and 5., the **write** can be forced by setting the **writeany** edit option. 14979 Additional, implementation-defined tests may cause the **write** to fail. If the edit buffer is empty, a file without any contents shall be written. 14980 An informational message shall be written noting the number of lines and bytes written. 14981 Otherwise, if the command is followed by one or more <blank>s, and the file is preceded by 14982 '!', the rest of the line after the '!' shall have '%', '#', and '!' characters expanded as 14983 14984 described in **Command Line Parsing in ex** (on page 362). The ex utility shall then pass two arguments to the program named by the **shell** edit option; the 14985 first shall be -c and the second shall be the expanded arguments to the write command as a 14986 single argument. The specified lines shall be written to the standard input of the command. The 14987 standard error and standard output of the program, if any, shall be written as described for the 14988 print command. If the last character in that output is not a <newline>, a <newline> shall be 14989 written at the end of the output. 14990 The special meaning of the '!' following the write command can be overridden by escaping it 14991 with a backslash character. 14992 Current line: Unchanged. 14993 14994 Current column: Unchanged. Write and Exit 14995 Synopsis: $[2addr] \times [it] [!] [file]$ 14996 If the edit buffer has not been modified since the last complete write, xit shall be equivalent to 14997 the **quit** command, or if a '!' is appended to the command name, to **quit!**. 14998 14999 Otherwise, **xit** shall be equivalent to the **wq** command, or if a '!' is appended to the command name, to wq!. 15000 15001 Current line: Unchanged. Current column: Unchanged. 15002 Yank 15003 Synopsis: [2addr] ya[nk][buffer][count] 15004 Copy the specified lines to the specified buffer (by default, the unnamed buffer), which shall 15005 15006 become a line-mode buffer. Current line: Unchanged. 15007

Current column: Unchanged.

Adjust Window

15009

15014

15015

15016 15017

15018 15019

15020

15021

15022

15023

15024

15025

15026

15027

15028 15029

15030

15031

15032

15033

15034

15035

15036

15037 15038

15039

15040

15041

15042

15043 15044

15045

15046

15047

15010 Synopsis: [laddr] z[!][type ...][count][flags]

If no line is specified, the current line shall be the default; if *type* is omitted as well, the current line value shall first be incremented by 1. If incrementing the current line would cause it to be greater than the last line in the edit buffer, it shall be an error.

If there are <blank>s between the *type* argument and the preceding **z** command name or optional '!' character, it shall be an error.

If *count* is specified, the value of the **window** edit option shall be set to *count* (as described in **window** (on page 398)). If *count* is omitted, it shall default to 2 times the value of the **scroll** edit option, or if! was specified, the number of lines in the display minus 1.

If *type* is omitted, then *count* lines starting with the specified line shall be written. Otherwise, *count* lines starting with the line specified by the *type* argument shall be written.

The *type* argument shall change the lines to be written. The possible values of *type* are as follows:

- The specified line shall be decremented by the following value:

```
(((number of ``-'' characters) x count) -1)
```

If the calculation would result in a number less than 1, it shall be an error. Write lines from the edit buffer, starting at the new value of line, until *count* lines or the last line in the edit buffer has been written.

+ The specified line shall be incremented by the following value:

```
(((number of ''+'' characters) -1) x count) +1
```

If the calculation would result in a number greater than the last line in the edit buffer, it shall be an error. Write lines from the edit buffer, starting at the new value of line, until *count* lines or the last line in the edit buffer has been written.

- =,. If more than a single '.' or '=' is specified, it shall be an error. The following steps shall be taken:
 - 1. If *count* is zero, nothing shall be written.
 - 2. Write as many of the N lines before the current line in the edit buffer as exist. If *count* or '!' was specified, N shall be:

```
(count -1) /2
```

Otherwise, N shall be:

```
(count -3) /2
```

If *N* is a number less than 3, no lines shall be written.

- 3. If '=' was specified as the type character, write a line consisting of the smaller of the number of columns in the display divided by two, or 40'-' characters.
- 4. Write the current line.
- 5. Repeat step 3.
- 6. Write as many of the *N* lines after the current line in the edit buffer as exist. *N* shall be defined as in step 2. If *N* is a number less than 3, no lines shall be written. If *count* is less than 3, no lines shall be written.

15048 The specified line shall be decremented by the following value:

(((number of ``^'' characters) +1) x count) -1 15049

> If the calculation would result in a number less than 1, it shall be an error. Write lines from the edit buffer, starting at the new value of line, until *count* lines or the last line in the edit buffer has been written.

Current line: Set to the last line written, unless the type is =, in which case, set to the specified

Current column: Set to non-

- slank>.

Escape 15056

15050

15051 15052

15053 15054

15055

15059

15060

15061

15062

15063

15064

15065

15066

15067

15068

15069 15070

15071

15072

15073

15074 15075

15076

15077

15078 15079

15080 15081

15082

15083 15084

Synopsis: ! command 15057 15058

[addr]! command

The contents of the line after the '!' shall have '%', '#', and '!' characters expanded as described in **Command Line Parsing in ex** (on page 362). If the expansion causes the text of the line to change, it shall be redisplayed, preceded by a single '!' character.

The ex utility shall execute the program named by the shell edit option. It shall pass two arguments to the program; the first shall be -c, and the second shall be the expanded arguments to the! command as a single argument.

If no lines are specified, the standard input, standard output, and standard error of the program shall be set to the standard input, standard output, and standard error of the ex program when it was invoked. In addition, a warning message shall be written if the edit buffer has been modified since the last complete write, and the warn edit option is set.

If lines are specified, they shall be passed to the program as standard input, and the standard output and standard error of the program shall replace those lines in the edit buffer. Each line in the program output (as delimited by <newline>s or the end of the output if it is not immediately preceded by a <newline>), shall be a separate line in the edit buffer. Any occurrences of <carriage-return> and <newline> pairs in the output shall be treated as single <newline>s. The specified lines shall be copied into the unnamed buffer before they are replaced, and the unnamed buffer shall become a line-mode buffer.

If in ex mode, a single '!' character shall be written when the program completes.

This command shall be affected by the **shell** and **warn** edit options. If no lines are specified, this command shall be affected by the autowrite and writeany edit options. If lines are specified, this command shall be affected by the **autoprint** edit option.

Current line:

- 1. If no lines are specified, unchanged.
- Otherwise, set to the last line read in, if any lines are read in.
- Otherwise, set to the line before the first line of the lines specified, if that line exists.
- Otherwise, set to the first line of the edit buffer if the edit buffer is not empty.
- 15085 5. Otherwise, set to zero.

15086 *Current column*: If no lines are specified, unchanged. Otherwise, set to non-

-blank>.

Utilities $\mathbf{e}\mathbf{x}$

15087	Shift Left
15088	Synopsis: [2addr] <[<] [count] [flags]
15089 15090 15091 15092	Shift the specified lines to the start of the line; the number of column positions to be shifted shall be the number of command characters times the value of the shiftwidth edit option. Only leading blank>s shall be deleted or changed into other blank>s in shifting; other characters shall not be affected.
15093 15094	Lines to be shifted shall be copied into the unnamed buffer, which shall become a line-mode buffer.
15095	This command shall be affected by the autoprint edit option.
15096	Current line: Set to the last line in the lines specified.
15097	Current column: Set to non- <blank>.</blank>
15098	Shift Right
15099	Synopsis: [2addr] >[>] [count] [flags]
15100 15101 15102 15103	Shift the specified lines away from the start of the line; the number of column positions to be shifted shall be the number of command characters times the value of the shiftwidth edit option. The shift shall be accomplished by adding <blank>s as a prefix to the line or changing leading <blank>s into other <blank>s. Empty lines shall not be changed.</blank></blank></blank>
15104 15105	Lines to be shifted shall be copied into the unnamed buffer, which shall become a line-mode buffer.
15106	This command shall be affected by the autoprint edit option.
15107	Current line: Set to the last line in the lines specified.
15108	Current column: Set to non- - slank>.
15109	<control>-D</control>
15110	Synopsis: <control>-D</control>
15111 15112 15113	Write the next n lines, where n is the minimum of the values of the scroll edit option and the number of lines after the current line in the edit buffer. If the current line is the last line of the edit buffer it shall be an error.
15114	Current line: Set to the last line written.
15115	Current column: Set to non- <blank>.</blank>
15116	Write Line Number
15117	Synopsis: [1addr] = [flags]
15118 15119	If <i>line</i> is not specified, it shall default to the last line in the edit buffer. Write the line number of the specified line.
15120	Current line: Unchanged.
15121	Current column: Unchanged.

15122 Execute 15123 Synopsis: [2addr] @ buffer 15124 [2addr] * buffer

If no buffer is specified or is specified as '@' or '*', the last buffer executed shall be used. If no previous buffer has been executed, it shall be an error.

For each line specified by the addresses, set the current line ('.') to the specified line, and execute the contents of the named *buffer* (as they were at the time the @ command was executed) as *ex* commands. For each line of a line-mode buffer, and all but the last line of a character-mode buffer, the *ex* command parser shall behave as if the line was terminated by a <newline>.

If an error occurs during this process, or a line specified by the addresses does not exist when the current line would be set to it, or more than a single line was specified by the addresses, and the contents of the edit buffer are replaced (for example, by the *ex*:*edit* command) an error message shall be written, and no more commands resulting from the execution of this command shall be processed.

Current line: As specified for the individual ex commands.

Current column: As specified for the individual *ex* commands.

Regular Expressions in ex

The *ex* utility shall support regular expressions that are a superset of the basic regular expressions described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.3, Basic Regular Expressions. A null regular expression ("//") shall be equivalent to the last regular expression encountered.

Regular expressions can be used in addresses to specify lines and, in some commands (for example, the **substitute** command), to specify portions of a line to be substituted.

The following constructs can be used to enhance the basic regular expressions:

- \< Match the beginning of a word. (See the definition of word at the beginning of Command Descriptions in ex (on page 368).)</p>
- \> Match the end of a *word*.
 - Match the replacement part of the last **substitute** command. The tilde ('~') character can be escaped in a regular expression to become a normal character with no special meaning. The backslash shall be discarded.

When the editor option **magic** is not set, the only characters with special meanings shall be '^' at the beginning of a pattern, '\$' at the end of a pattern, and '\'. The characters '.', '*', '[', and '~' shall be treated as ordinary characters unless preceded by a '\'; when preceded by a '\' they shall regain their special meaning, or in the case of backslash, be handled as a single backslash. Backslashes used to escape other characters shall be discarded.

Replacement Strings in ex

The character '&' ('\&' if the editor option **magic** is not set) in the replacement string shall stand for the text matched by the pattern to be replaced. The character ' $^{'}$ ' (' $^{'}$ ' if **magic** is not set) shall be replaced by the replacement part of the previous **substitute** command. The sequence ' $^{'}$ n', where *n* is an integer, shall be replaced by the text matched by the pattern enclosed in the *n*th set of parentheses ' $^{'}$ (' and ' $^{'}$)'.

The strings '\l', '\u', '\L', and '\U' can be used to modify the case of elements in the replacement string (using the '\&' or "\"digit) notation. The string '\l' ('\u') shall cause

15165 the character that follows to be converted to lowercase (uppercase). The string '\L' ('\U') shall cause all characters subsequent to it to be converted to lowercase (uppercase) as they are 15166 inserted by the substitution until the string '\e' or '\E', or the end of the replacement string, 15167 is encountered. 15168 Otherwise, any character following a backslash shall be treated as that literal character, and the 15169 escaping backslash shall be discarded. 15170 An example of case conversion with the **s** command is as follows: 15171 15172 **:**p The cat sat on the mat. 15173

15172 :p
15173 The cat sat on the mat.
15174 :s/\<.at\>/\u&/gp
15175 The Cat Sat on the Mat.
15176 :s/S\((.*\))M/S\U\1\eM/p
15177 The Cat SAT ON THE Mat.

Edit Options in ex

The *ex* utility has a number of options that modify its behavior. These options have default settings, which can be changed using the **set** command.

Options are Boolean unless otherwise specified.

autoindent, ai

[Default *unset*]

If **autoindent** is set, each line in input mode shall be indented (using first as many <tab>s as possible, as determined by the editor option **tabstop**, and then using <space>s) to align with another line, as follows:

- 1. If in open or visual mode and the text input is part of a line-oriented command (see the EXTENDED DESCRIPTION in *vi*), align to the first column.
- 2. Otherwise, if in open or visual mode, indentation for each line shall be set as follows:
 - a. If a line was previously inserted as part of this command, it shall be set to the indentation of the last inserted line by default, or as otherwise specified for the <control>-D character in **Input Mode Commands in vi** (on page 1022).
 - b. Otherwise, it shall be set to the indentation of the previous current line, if any; otherwise, to the first column.
- 3. For the ex a, i, and c commands, indentation for each line shall be set as follows:
 - a. If a line was previously inserted as part of this command, it shall be set to the indentation of the last inserted line by default, or as otherwise specified for the *eof* character in **Scroll** (on page 366).
 - b. Otherwise, if the command is the *ex* **a** command, it shall be set to the line appended after, if any; otherwise to the first column.
 - c. Otherwise, if the command is the *ex* **i** command, it shall be set to the line inserted before, if any; otherwise to the first column.
 - d. Otherwise, if the command is the *ex* **c** command, it shall be set to the indentation of the line replaced.

15178

15179 15180

15181

15182

15183

15184

15185

15186

15187

15188

15189

15191

15192

15193

15194

15195

15196

15197

15198

15199

15200

15201

15205 autoprint, ap [Default set] 15206 If autoprint is set, the current line shall be written after each ex command that modifies the 15207 15208 contents of the current edit buffer, and after each tag command for which the tag search pattern 15209 was found or tag line number was valid, unless: The command was executed while in open or visual mode. 15210 The command was executed as part of a **global** or **v** command or @ buffer execution. 15211 15212 The command was the form of the **read** command that reads a file into the edit buffer. The command was the **append**, **change**, or **insert** command. 15213 15214 The command was not terminated by a <newline>. The current line shall be written by a flag specified to the command; for example, **delete** # 15215 15216 shall write the current line as specified for the flag modifier to the **delete** command, and 15217 not as specified by the **autoprint** edit option. autowrite, aw 15218 [Default unset] 15219 If autowrite is set, and the edit buffer has been modified since it was last completely written to 15220 any file, the contents of the edit buffer shall be written as if the ex write command had been 15221 specified without arguments, before each command affected by the autowrite edit option is 15222 executed. Appending the character '!' to the command name of any of the ex commands 15223 15224 except '!' shall prevent the write. If the write fails, it shall be an error and the command shall not be executed. 15225 beautify, bf 15226 15227 XSI [Default *unset*] If **beautify** is set, all non-printable characters, other than <tab>s, <newline>s, and <form-feed>s, 15228 shall be discarded from text read in from files. 15229 15230 directory, dir [Default implementation-defined] 15231 15232 The value of this option specifies the directory in which the editor buffer is to be placed. If this directory is not writable by the user, the editor shall quit. 15233 edcompatible, ed 15234 [Default *unset*] 15235 Causes the presence of g and c suffixes on substitute commands to be remembered, and toggled 15236 by repeating the suffixes. 15237

15238	errorbells, eb
15239	[Default unset]
15240 15241	If the editor is in <i>ex</i> mode, and the terminal does not support a standout mode (such as inverse video), and errorbells is set, error messages shall be preceded by alerting the terminal.
15242	ехгс
15243	[Default unset]
15244	If exrc is set, <i>ex</i> shall access any .exrc file in the current directory, as described in Initialization in
15245	ex and vi (on page 358). If exrc is not set, ex shall ignore any .exrc file in the current directory
15246	during initialization, unless the current directory is that named by the HOME environment
15247	variable.
15248	ignorecase, ic
15249	[Default <i>unset</i>]
15250 15251	If ignorecase is set, characters that have uppercase and lowercase representations shall have those representations considered as equivalent for purposes of regular expression comparison.
15252	The ignorecase edit option shall affect all remembered regular expressions; for example,
15253	unsetting the ignorecase edit option shall cause a subsequent vi n command to search for the
15254	last basic regular expression in a case-sensitive fashion.
15255	list
15256	[Default unset]
15257	If list is set, edit buffer lines written while in ex command mode shall be written as specified for
15258	the print command with the l flag specified. In open or visual mode, each edit buffer line shall
15259	be displayed as specified for the ex print command with the l flag specified. In open or visual
15260	text input mode, when the cursor does not rest on any character in the line, it shall rest on the
15261	'\$' marking the end of the line.
15262	magic
15263	[Default set]
15264	If magic is set, modify the interpretation of characters in regular expressions and substitution
15265	replacement strings (see Regular Expressions in ex (on page 391) and Replacement Strings in
15266	ex (on page 391)).
15267	mesg
15268	[Default set]
15269	If mesg is set, the permission for others to use the write or talk commands to write to the
15270	terminal shall be turned on while in open or visual mode. The shell-level command mesg n shall
15271	take precedence over any setting of the ex mesg option; that is, if mesg y was issued before the
15272	editor started (or in a shell escape), such as:
15273	:!mesg y
15274	the $mesg$ option in ex shall suppress incoming messages, but the $mesg$ option shall not enable
15275	incoming messages if mesg n was issued.

15276 number, nu [Default unset] 15277 If **number** is set, edit buffer lines written while in *ex* command mode shall be written with line 15278 15279 numbers, in the format specified by the **print** command with the # flag specified. In ex text input mode, each line shall be preceded by the line number it will have in the file. 15280 In open or visual mode, each edit buffer line shall be displayed with a preceding line number, in 15281 15282 the format specified by the *ex* **print** command with the # flag specified. This line number shall not be considered part of the line for the purposes of evaluating the current column; that is, 15283 column position 1 shall be the first column position after the format specified by the **print** 15284 command. 15285 paragraphs, para 15286 [Default in the POSIX locale IPLPPPQPP LIpplpipbp] 15287 15288 The paragraphs edit option shall define additional paragraph boundaries for the open and visual mode commands. The paragraphs edit option can be set to a character string consisting of zero 15289 15290 or more character pairs. It shall be an error to set it to an odd number of characters. prompt 15291 15292 [Default set] 15293 If **prompt** is set, *ex* command mode input shall be prompted for with a colon (':'); when unset, no prompt shall be written. 15294 readonly 15295 [Default see text] 15296 If the **readonly** edit option is set, read-only mode shall be enabled (see **Write** (on page 386)). The 15297 **readonly** edit option shall be initialized to set if either of the following conditions are true: 15298 The command-line option –R was specified. 15299 • Performing actions equivalent to the access() function called with the following arguments 15300 indicates that the file lacks write permission: 15301 15302 1. The current pathname is used as the *path* argument. 15303 The constant **W_OK** is used as the *amode* argument. The **readonly** edit option may be initialized to set for other, implementation-defined reasons. 15304 The **readonly** edit option shall not be initialized to unset based on any special privileges of the 15305 user or process. The **readonly** edit option shall be reinitialized each time that the contents of the 15306

edit buffer are replaced (for example, by an edit or next command) unless the user has explicitly

set it, in which case it shall remain set until the user explicitly unsets it. Once unset, it shall again

be reinitialized each time that the contents of the edit buffer are replaced.

15307

15308 15309

15310 redraw [Default unset] 15311 The editor simulates an intelligent terminal on a dumb terminal. (Since this is likely to require a 15312 15313 large amount of output to the terminal, it is useful only at high transmission speeds.) 15314 remap 15315 [Default set] 15316 If **remap** is set, map translation shall allow for maps defined in terms of other maps; translation 15317 shall continue until a final product is obtained. If unset, only a one-step translation shall be done. 15318 report [Default 5] 15319 The value of this **report** edit option specifies what number of lines being added, copied, deleted, 15320 or modified in the edit buffer will cause an informational message to be written to the user. The 15321 following conditions shall cause an informational message. The message shall contain the 15322 number of lines added, copied, deleted, or modified, but is otherwise unspecified. 15323 • An ex or vi editor command, other than open, undo, or visual, that modifies at least the value 15324 of the **report** edit option number of lines, and which is not part of an ex global or v 15325 command, or ex or vi buffer execution, shall cause an informational message to be written. 15326 An ex yank or vi y or Y command, that copies at least the value of the report edit option plus 15327 15328 1 number of lines, and which is not part of an ex global or v command, or ex or vi buffer 15329 execution, shall cause an informational message to be written. • An ex global, v, open, undo, or visual command or ex or vi buffer execution, that adds or 15330 15331 deletes a total of at least the value of the report edit option number of lines, and which is not part of an ex global or v command, or ex or vi buffer execution, shall cause an informational 15332 15333 message to be written. (For example, if 3 lines were added and 8 lines deleted during an ex visual command, 5 would be the number compared against the report edit option after the 15334 15335 command completed.) scroll, scr 15336 15337 [Default (number of lines in the display -1)/2] 15338 The value of the **scroll** edit option shall determine the number of lines scrolled by the ex 15339 <control>-D and z commands. For the vi <control>-D and <control>-U commands, it shall be the initial number of lines to scroll when no previous <control>-D or <control>-U command has 15340 been executed. 15341 sections 15342 [Default in the POSIX locale NHSHH HUnhsh] 15343 The sections edit option shall define additional section boundaries for the open and visual mode 15344

15345

15346

commands. The sections edit option can be set to a character string consisting of zero or more

character pairs; it shall be an error to set it to an odd number of characters.

15347	shell, sh		
15348	[Default from the environment variable SHELL]		
15349 15350 15351	The value of this option shall be a string. The default shall be taken from the $SHELL$ environment variable. If the $SHELL$ environment variable is null or empty, the sh (see sh) utility shall be the default.		
15352	shiftwidth, sw		
15353	[Default 8]		
15354 15355	The value of this option shall give the width in columns of an indentation level used during autoindentation and by the shift commands (< and >).		
15356	showmatch, sm		
15357	[Default <i>unset</i>]		
15358 15359	The functionality described for the showmatch edit option need not be supported on block-mode terminals or terminals with insufficient capabilities.		
15360 15361 15362	If showmatch is set, in open or visual mode, when a ')' or '' is typed, if the matching '(' or '' is currently visible on the display, the matching '(' or '' shall be flagged moving the cursor to its location for an unspecified amount of time.		
15363	showmode		
15364	[Default unset]		
15365 15366 15367 15368	If showmode is set, in open or visual mode, the current mode that the editor is in shall be displayed on the last line of the display. Command mode and text input mode shall be differentiated; other unspecified modes and implementation-defined information may be displayed.		
15369	slowopen		
15370	[Default <i>unset</i>]		
15371 15372 15373	If slowopen is set during open and visual text input modes, the editor shall not update portions of the display other than those display line columns that display the characters entered by the user (see Input Mode Commands in vi (on page 1022)).		
15374	tabstop, ts		
15375	[Default 8]		
15376 15377	The value of this edit option shall specify the column boundary used by a <tab> in the display (see autoprint, ap (on page 393) and Input Mode Commands in vi (on page 1022)).</tab>		
15378	taglength, tl		
15379	[Default zero]		
15380 15381 15382	The value of this edit option shall specify the maximum number of characters that are considered significant in the user-specified tag name and in the tag name from the tags file. If the value is zero, all characters in both tag names shall be significant.		

15383	tags
15384	[Default see text]
15385 15386	The value of this edit option shall be a string of <blank>-delimited pathnames of files used by the tag command. The default value is unspecified.</blank>
15387	term
15388	[Default from the environment variable TERM]
15389 15390 15391 15392	The value of this edit option shall be a string. The default shall be taken from the <i>TERM</i> variable in the environment. If the <i>TERM</i> environment variable is empty or null, the default is unspecified. The editor shall use the value of this edit option to determine the type of the display device.
15393 15394	The results are unspecified if the user changes the value of the term edit option after editor initialization.
15395	terse
15396	[Default unset]
15397 15398 15399	If terse is set, error messages may be less verbose. However, except for this caveat, error messages are unspecified. Furthermore, not all error messages need change for different settings of this option.
15400	warn
15401	[Default set]
15402 15403 15404	If warn is set, and the contents of the edit buffer have been modified since they were last completely written, the editor shall write a warning message before certain! commands (see Escape (on page 389)).
15405	window
15406	[Default see text]
15407 15408	A value used in open and visual mode, by the <control>-B and <control>-F commands, and, in visual mode, to specify the number of lines displayed when the screen is repainted.</control></control>
15409 15410 15411	If the $-\mathbf{w}$ command-line option is not specified, the default value shall be set to the value of the <i>LINES</i> environment variable. If the <i>LINES</i> environment variable is empty or null, the default shall be the number of lines in the display minus 1.
15412 15413 15414	Setting the window edit option to zero or to a value greater than the number of lines in the display minus 1 (either explicitly or based on the –w option or the <i>LINES</i> environment variable) shall cause the window edit option to be set to the number of lines in the display minus 1.
15415	The baud rate of the terminal line may change the default in an implementation-defined manner.

15416 wrapmargin, wm [Default 0] 15417 If the value of this edit option is zero, it shall have no effect. 15418 If not in the POSIX locale, the effect of this edit option is implementation-defined. 15419 Otherwise, it shall specify a number of columns from the ending margin of the terminal. 15420 15421 During open and visual text input modes, for each character for which any part of the character 15422 is displayed in a column that is less than wrapmargin columns from the ending margin of the 15423 display line, the editor shall behave as follows: 1. If the character triggering this event is a <blank>, it, and all immediately preceding 15424 15425
<blank>s on the current line entered during the execution of the current text input 15426 command, shall be discarded, and the editor shall behave as if the user had entered a single <newline> instead. In addition, if the next user-entered character is a <space>, it shall be 15427 discarded as well. 15428 2. Otherwise, if there are one or more

blank>s on the current line immediately preceding the 15429 15430 last group of inserted non-
 -slank>s which was entered during the execution of the current text input command, the <blank>s shall be replaced as if the user had entered a single 15431 <newline> instead. 15432 If the autoindent edit option is set, and the events described in 1. or 2. are performed, any 15433
 15434 The ending margin shall be determined by the system or overridden by the user, as described for 15435 COLUMNS in the ENVIRONMENT VARIABLES section and the Base Definitions volume of 15436 IEEE Std 1003.1-2001, Chapter 8, Environment Variables. 15437 15438 wrapscan, ws [Default set] 15439 If wrapscan is set, searches (the ex / or? addresses, or open and visual mode /, ?, N, and n 15440 15441 commands) shall wrap around the beginning or end of the edit buffer; when unset, searches 15442 shall stop at the beginning or end of the edit buffer. writeany, wa 15443 [Default *unset*] 15444 15445 If writeany is set, some of the checks performed when executing the ex write commands shall be inhibited, as described in editor option autowrite. 15446 15447 EXIT STATUS The following exit values shall be returned: 15448 Successful completion. 15449 >0 An error occurred. 15450 15451 CONSEQUENCES OF ERRORS When any error is encountered and the standard input is not a terminal device file, ex shall not 15452 write the file or return to command or text input mode, and shall terminate with a non-zero exit 15453 15454 status.

asynchronous event.

Otherwise, when an unrecoverable error is encountered, it shall be equivalent to a SIGHUP

15455

15456

Otherwise, when an error is encountered, the editor shall behave as specified in **Command Line**Parsing in ex (on page 362).

15459 APPLICATION USAGE

15460 If a SIGSEGV signal is received while *ex* is saving a file, the file might not be successfully saved.

15461 The **next** command can accept more than one file, so usage such as:

```
15462 next 'ls [abc] * '
```

is valid; it would not be valid for the **edit** or **read** commands, for example, because they expect only one file and unspecified results occur.

15465 EXAMPLES

15466 None.

15467 RATIONALE

The *ex/vi* specification is based on the historical practice found in the 4 BSD and System V implementations of *ex* and *vi*. A freely redistributable implementation of *ex/vi*, which is tracking IEEE Std 1003.1-2001 fairly closely, and demonstrates the intended changes between historical implementations and IEEE Std 1003.1-2001, may be obtained by anonymous FTP from:

```
ftp://ftp.rdg.opengroup.org/pub/mirrors/nvi
```

A *restricted editor* (both the historical *red* utility and modifications to *ex*) were considered and rejected for inclusion. Neither option provided the level of security that users might expect.

It is recognized that *ex* visual mode and related features would be difficult, if not impossible, to implement satisfactorily on a block-mode terminal, or a terminal without any form of cursor addressing; thus, it is not a mandatory requirement that such features should work on all terminals. It is the intention, however, that an *ex* implementation should provide the full set of capabilities on all terminals capable of supporting them.

Options

The -c replacement for +command was inspired by the -e option of sed. Historically, all such commands (see edit and next as well) were executed from the last line of the edit buffer. This meant, for example, that "+/pattern" would fail unless the wrapscan option was set. IEEE Std 1003.1-2001 requires conformance to historical practice. Historically, some implementations restricted the ex commands that could be listed as part of the command line arguments. For consistency, IEEE Std 1003.1-2001 does not permit these restrictions.

In historical implementations of the editor, the **–R** option (and the **readonly** edit option) only prevented overwriting of files; appending to files was still permitted, mapping loosely into the *csh* **noclobber** variable. Some implementations, however, have not followed this semantic, and **readonly** does not permit appending either. IEEE Std 1003.1-2001 follows the latter practice, believing that it is a more obvious and intuitive meaning of **readonly**.

The -s option suppresses all interactive user feedback and is useful for editing scripts in batch jobs. The list of specific effects is historical practice. The terminal type "incapable of supporting open and visual modes" has historically been named "dumb".

The -t option was required because the *ctags* utility appears in IEEE Std 1003.1-2001 and the option is available in all historical implementations of *ex*.

Historically, the ex and vi utilities accepted a -x option, which did encryption based on the algorithm found in the historical crypt utility. The -x option for encryption, and the associated crypt utility, were omitted because the algorithm used was not specifiable and the export control laws of some nations make it difficult to export cryptographic technology. In addition, it did not

historically provide the level of security that users might expect.

Standard Input

An end-of-file condition is not equivalent to an end-of-file character. A common end-of-file character, <control>-D, is historically an *ex* command.

There was no maximum line length in historical implementations of *ex*. Specifically, as it was parsed in chunks, the addresses had a different maximum length than the filenames. Further, the maximum line buffer size was declared as BUFSIZ, which was different lengths on different systems. This version selected the value of {LINE_MAX} to impose a reasonable restriction on portable usage of *ex* and to aid test suite writers in their development of realistic tests that exercise this limit.

Input Files

It was an explicit decision by the standard developers that a <newline> be added to any file lacking one. It was believed that this feature of *ex* and *vi* was relied on by users in order to make text files lacking a trailing <newline> more portable. It is recognized that this will require a user-specified option or extension for implementations that permit *ex* and *vi* to edit files of type other than text if such files are not otherwise identified by the system. It was agreed that the ability to edit files of arbitrary type can be useful, but it was not considered necessary to mandate that an *ex* or *vi* implementation be required to handle files other than text files.

The paragraph in the INPUT FILES section, "By default, . . . ", is intended to close a long-standing security problem in ex and vi; that of the "modeline" or "modelines" edit option. This feature allows any line in the first or last five lines of the file containing the strings "ex:" or "vi:" (and, apparently, "ei:" or "vx:") to be a line containing editor commands, and ex interprets all the text up to the next ':' or <newline> as a command. Consider the consequences, for example, of an unsuspecting user using ex or vi as the editor when replying to a mail message in which a line such as:

```
ex:! rm -rf :
```

appeared in the signature lines. The standard developers believed strongly that an editor should not by default interpret any lines of a file. Vendors are strongly urged to delete this feature from their implementations of *ex* and *vi*.

Asynchronous Events

The intention of the phrase "complete write" is that the entire edit buffer be written to stable storage. The note regarding temporary files is intended for implementations that use temporary files to back edit buffers unnamed by the user.

Historically, SIGQUIT was ignored by ex, but was the equivalent of the \mathbf{Q} command in visual mode; that is, it exited visual mode and entered ex mode. IEEE Std 1003.1-2001 permits, but does not require, this behavior. Historically, SIGINT was often used by vi users to terminate text input mode (<control>-C is often easier to enter than <ESC>). Some implementations of vi alerted the terminal on this event, and some did not. IEEE Std 1003.1-2001 requires that SIGINT behave identically to <ESC>, and that the terminal not be alerted.

Historically, suspending the *ex* editor during text input mode was similar to SIGINT, as completed lines were retained, but any partial line discarded, and the editor returned to command mode. IEEE Std 1003.1-2001 is silent on this issue; implementations are encouraged to follow historical practice, where possible.

Historically, the *vi* editor did not treat SIGTSTP as an asynchronous event, and it was therefore impossible to suspend the editor in visual text input mode. There are two major reasons for this. The first is that SIGTSTP is a broadcast signal on UNIX systems, and the chain of events where the shell *execs* an application that then *execs vi* usually caused confusion for the terminal state if SIGTSTP was delivered to the process group in the default manner. The second was that most implementations of the UNIX *curses* package are not reentrant, and the receipt of SIGTSTP at the wrong time will cause them to crash. IEEE Std 1003.1-2001 is silent on this issue; implementations are encouraged to treat suspension as an asynchronous event if possible.

Historically, modifications to the edit buffer made before SIGINT interrupted an operation were retained; that is, anywhere from zero to all of the lines to be modified might have been modified by the time the SIGINT arrived. These changes were not discarded by the arrival of SIGINT. IEEE Std 1003.1-2001 permits this behavior, noting that the **undo** command is required to be able to undo these partially completed commands.

The action taken for signals other than SIGINT, SIGCONT, SIGHUP, and SIGTERM is unspecified because some implementations attempt to save the edit buffer in a useful state when other signals are received.

Standard Error

For *ex/vi*, diagnostic messages are those messages reported as a result of a failed attempt to invoke *ex* or *vi*, such as invalid options or insufficient resources, or an abnormal termination condition. Diagnostic messages should not be confused with the error messages generated by inappropriate or illegal user commands.

Initialization in ex and vi

If an *ex* command (other than **cd**, **chdir**, or **source**) has a filename argument, one or both of the alternate and current pathnames will be set. Informally, they are set as follows:

- 1. If the *ex* command is one that replaces the contents of the edit buffer, and it succeeds, the current pathname will be set to the filename argument (the first filename argument in the case of the **next** command) and the alternate pathname will be set to the previous current pathname, if there was one.
- 2. In the case of the file read/write forms of the **read** and **write** commands, if there is no current pathname, the current pathname will be set to the filename argument.
- 3. Otherwise, the alternate pathname will be set to the filename argument.

For example, :edit foo and :recover foo, when successful, set the current pathname, and, if there was a previous current pathname, the alternate pathname. The commands :write, !command, and :edit set neither the current or alternate pathnames. If the :edit foo command were to fail for some reason, the alternate pathname would be set. The read and write commands set the alternate pathname to their file argument, unless the current pathname is not set, in which case they set the current pathname to their file arguments. The alternate pathname was not historically set by the :source command. IEEE Std 1003.1-2001 requires conformance to historical practice. Implementations adding commands that take filenames as arguments are encouraged to set the alternate pathname as described here.

Historically, *ex* and *vi* read the .exrc file in the *\$HOME* directory twice, if the editor was executed in the *\$HOME* directory. IEEE Std 1003.1-2001 prohibits this behavior.

Historically, the 4 BSD *ex* and *vi* read the *SHOME* and local .exrc files if they were owned by the real ID of the user, or the **sourceany** option was set, regardless of other considerations. This was a security problem because it is possible to put normal UNIX system commands inside a .exrc

file. IEEE Std 1003.1-2001 does not specify the **sourceany** option, and historical implementations are encouraged to delete it.

 The .exrc files must be owned by the real ID of the user, and not writable by anyone other than the owner. The appropriate privileges exception is intended to permit users to acquire special privileges, but continue to use the .exrc files in their home directories.

System V Release 3.2 and later *vi* implementations added the option [no]exrc. The behavior is that local .exrc files are read-only if the exrc option is set. The default for the exrc option was off, so by default, local .exrc files were not read. The problem this was intended to solve was that System V permitted users to give away files, so there is no possible ownership or writeability test to ensure that the file is safe. This is still a security problem on systems where users can give away files, but there is nothing additional that IEEE Std 1003.1-2001 can do. The implementation-defined exception is intended to permit groups to have local .exrc files that are shared by users, by creating pseudo-users to own the shared files.

IEEE Std 1003.1-2001 does not mention system-wide *ex* and *vi* start-up files. While they exist in several implementations of *ex* and *vi*, they are not present in any implementations considered historical practice by IEEE Std 1003.1-2001. Implementations that have such files should use them only if they are owned by the real user ID or an appropriate user (for example, root on UNIX systems) and if they are not writable by any user other than their owner. System-wide start-up files should be read before the *EXINIT* variable, **\$HOME/.exrc**, or local .exrc files are evaluated.

Historically, any *ex* command could be entered in the *EXINIT* variable or the .exrc file, although ones requiring that the edit buffer already contain lines of text generally caused historical implementations of the editor to drop core. IEEE Std 1003.1-2001 requires that any *ex* command be permitted in the *EXINIT* variable and .exrc files, for simplicity of specification and consistency, although many of them will obviously fail under many circumstances.

The initialization of the contents of the edit buffer uses the phrase "the effect shall be" with regard to various *ex* commands. The intent of this phrase is that edit buffer contents loaded during the initialization phase not be lost; that is, loading the edit buffer should fail if the .exrc file read in the contents of a file and did not subsequently write the edit buffer. An additional intent of this phrase is to specify that the initial current line and column is set as specified for the individual *ex* commands.

Historically, the –t option behaved as if the tag search were a +command; that is, it was executed from the last line of the file specified by the tag. This resulted in the search failing if the pattern was a forward search pattern and the wrapscan edit option was not set. IEEE Std 1003.1-2001 does not permit this behavior, requiring that the search for the tag pattern be performed on the entire file, and, if not found, that the current line be set to a more reasonable location in the file.

Historically, the empty edit buffer presented for editing when a file was not specified by the user was unnamed. This is permitted by IEEE Std 1003.1-2001; however, implementations are encouraged to provide users a temporary filename for this buffer because it permits them the use of *ex* commands that use the current pathname during temporary edit sessions.

Historically, the file specified using the –t option was not part of the current argument list. This practice is permitted by IEEE Std 1003.1-2001; however, implementations are encouraged to include its name in the current argument list for consistency.

Historically, the $-\mathbf{c}$ command was generally not executed until a file that already exists was edited. IEEE Std 1003.1-2001 requires conformance to this historical practice. Commands that could cause the $-\mathbf{c}$ command to be executed include the *ex* commands **edit**, **next**, **recover**, **rewind**, and **tag**, and the *vi* commands <control>-^ and <control>-]. Historically, reading a file into an edit buffer did not cause the $-\mathbf{c}$ command to be executed (even though it might set the

current pathname) with the exception that it did cause the -c command to be executed if: the editor was in ex mode, the edit buffer had no current pathname, the edit buffer was empty, and no read commands had yet been attempted. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior.

Historically, the **-r** option was the same as a normal edit session if there was no recovery information available for the file. This allowed users to enter:

vi -r *.c

and recover whatever files were recoverable. In some implementations, recovery was attempted only on the first file named, and the file was not entered into the argument list; in others, recovery was attempted for each file named. In addition, some historical implementations ignored –**r** if –**t** was specified or did not support command line *file* arguments with the –**t** option. For consistency and simplicity of specification, IEEE Std 1003.1-2001 disallows these special cases, and requires that recovery be attempted the first time each file is edited.

Historically, *vi* initialized the 'and 'marks, but *ex* did not. This meant that if the first command in *ex* mode was **visual** or if an *ex* command was executed first (for example, *vi* +10 *file*), *vi* was entered without the marks being initialized. Because the standard developers believed the marks to be generally useful, and for consistency and simplicity of specification, IEEE Std 1003.1-2001 requires that they always be initialized if in open or visual mode, or if in *ex* mode and the edit buffer is not empty. Not initializing it in *ex* mode if the edit buffer is empty is historical practice; however, it has always been possible to set (and use) marks in empty edit buffers in open and visual mode edit sessions.

Addressing

Historically, ex and vi accepted the additional addressing forms '\/' and '\?'. They were equivalent to "//" and "??", respectively. They are not required by IEEE Std 1003.1-2001, mostly because nobody can remember whether they ever did anything different historically.

Historically, *ex* and *vi* permitted an address of zero for several commands, and permitted the % address in empty files for others. For consistency, IEEE Std 1003.1-2001 requires support for the former in the few commands where it makes sense, and disallows it otherwise. In addition, because IEEE Std 1003.1-2001 requires that % be logically equivalent to "1,\$", it is also supported where it makes sense and disallowed otherwise.

Historically, the % address could not be followed by further addresses. For consistency and simplicity of specification, IEEE Std 1003.1-2001 requires that additional addresses be supported.

All of the following are valid *addresses*:

+++ Three lines after the current line.

/re/- One line before the next occurrence of re.

−2 Two lines before the current line.

15673 3 ---- 2 Line one (note intermediate negative address).

1 2 3 Line six.

Any number of addresses can be provided to commands taking addresses; for example, "1,2,3,4,5p" prints lines 4 and 5, because two is the greatest valid number of addresses accepted by the **print** command. This, in combination with the semicolon delimiter, permits users to create commands based on ordered patterns in the file. For example, the command **3;/foo/;+2print** will display the first line after line 3 that contains the pattern *foo*, plus the next two lines. Note that the address **3;** must be evaluated before being discarded because the search

origin for the /**foo**/ command depends on this.

Historically, values could be added to addresses by including them after one or more
 sign for example, 3-5p wrote the seventh line of the file, and foo/5 was the same as foo/+5.
 However, only absolute values could be added; for example, foo/5 was an error. IEEE Std 1003.1-2001 requires conformance to historical practice. Address offsets are separately specified from addresses because they could historically be provided to visual mode search commands.

Historically, any missing addresses defaulted to the current line. This was true for leading and trailing comma-delimited addresses, and for trailing semicolon-delimited addresses. For consistency, IEEE Std 1003.1-2001 requires it for leading semicolon addresses as well.

Historically, ex and vi accepted the ' $^{\prime}$ ' character as both an address and as a flag offset for commands. In both cases it was identical to the ' $^{\prime}$ ' character. IEEE Std 1003.1-2001 does not require or prohibit this behavior.

Historically, the enhancements to basic regular expressions could be used in addressing; for example, $'\ ^{\prime}$, $'\ ^{\prime}$, and $'\ ^{\prime}$. IEEE Std 1003.1-2001 requires conformance to historical practice; that is, that regular expression usage be consistent, and that regular expression enhancements be supported wherever regular expressions are used.

Command Line Parsing in ex

Historical *ex* command parsing was even more complex than that described here. IEEE Std 1003.1-2001 requires the subset of the command parsing that the standard developers believed was documented and that users could reasonably be expected to use in a portable fashion, and that was historically consistent between implementations. (The discarded functionality is obscure, at best.) Historical implementations will require changes in order to comply with IEEE Std 1003.1-2001; however, users are not expected to notice any of these changes. Most of the complexity in *ex* parsing is to handle three special termination cases:

- 1. The !, global, v, and the filter versions of the read and write commands are delimited by <newline>s (they can contain vertical-line characters that are usually shell pipes).
- 2. The **ex**, **edit**, **next**, and **visual** in open and visual mode commands all take *ex* commands, optionally containing vertical-line characters, as their first arguments.
- 3. The **s** command takes a regular expression as its first argument, and uses the delimiting characters to delimit the command.

Historically, vertical-line characters in the +command argument of the ex, edit, next, vi, and visual commands, and in the pattern and replacement parts of the s command, did not delimit the command, and in the filter cases for read and write, and the !, global, and v commands, they did not delimit the command at all. For example, the following commands are all valid:

Historically, empty or
blank> filled lines in .exrc files and sourced files (as well as *EXINIT* variables and *ex* command scripts) were treated as default commands; that is, **print** commands. IEEE Std 1003.1-2001 specifically requires that they be ignored when encountered in .exrc and sourced files to eliminate a common source of new user error.

Historically, *ex* commands with multiple adjacent (or <blank>-separated) vertical lines were handled oddly when executed from *ex* mode. For example, the command | | < carriage-return>, when the cursor was on line 1, displayed lines 2, 3, and 5 of the file. In addition, the command | would only display the line after the next line, instead of the next two lines. The former worked more logically when executed from *vi* mode, and displayed lines 2, 3, and 4. IEEE Std 1003.1-2001 requires the *vi* behavior; that is, a single default command and line number increment for each command separator, and trailing <newline>s after vertical-line separators are discarded.

Historically, *ex* permitted a single extra colon as a leading command character; for example, *:g/pattern/:p* was a valid command. IEEE Std 1003.1-2001 generalizes this to require that any number of leading colon characters be stripped.

Historically, any prefix of the **delete** command could be followed without intervening <blank>s by a flag character because in the command **d p**, p is interpreted as the buffer p. IEEE Std 1003.1-2001 requires conformance to historical practice.

Historically, the **s** command could be immediately followed by flag and option characters; for example, s/e/E/|s|sgc3p was a valid command. However, flag characters could not stand alone; for example, the commands sp and s l would fail, while the command sp and s gl would succeed. (Obviously, the '#' flag character was used as a delimiter character if it followed the command.) Another issue was that option characters had to precede flag characters even when the command was fully specified; for example, the command s/e/E/pg would fail, while the command s/e/E/pg would succeed. IEEE Std 1003.1-2001 requires conformance to historical practice.

Historically, the first command name that had a prefix matching the input from the user was the executed command; for example, **ve**, **ver**, and **vers** all executed the **version** command. Commands were in a specific order, however, so that **a** matched **append**, not **abbreviate**. IEEE Std 1003.1-2001 requires conformance to historical practice. The restriction on command search order for implementations with extensions is to avoid the addition of commands such that the historical prefixes would fail to work portably.

Historical implementations of *ex* and *vi* did not correctly handle multiple *ex* commands, separated by vertical-line characters, that entered or exited visual mode or the editor. Because implementations of *vi* exist that do not exhibit this failure mode, IEEE Std 1003.1-2001 does not permit it.

The requirement that alphabetic command names consist of all following alphabetic characters up to the next non-alphabetic character means that alphabetic command names must be separated from their arguments by one or more non-alphabetic characters, normally a
blank> or '!' character, except as specified for the exceptions, the **delete**, **k**, and **s** commands.

Historically, the repeated execution of the *ex* default **print** commands (<control>-D, *eof*, <newline>, <carriage-return>) erased any prompting character and displayed the next lines without scrolling the terminal; that is, immediately below any previously displayed lines. This provided a cleaner presentation of the lines in the file for the user. IEEE Std 1003.1-2001 does not require this behavior because it may be impossible in some situations; however, implementations are strongly encouraged to provide this semantic if possible.

Historically, it was possible to change files in the middle of a command, and have the rest of the command executed in the new file; for example:

```
15771 :edit +25 file.c | s/abc/ABC/ | 1
```

was a valid command, and the substitution was attempted in the newly edited file. IEEE Std 1003.1-2001 requires conformance to historical practice. The following commands are examples that exercise the *ex* parser:

```
15775 echo 'foo | bar' > file1; echo 'foo/bar' > file2;
15776 vi
15777 :edit +1 | s/|/PIPE/ | w file1 | e file2 | 1 | s/\//SLASH/ | wq
```

Historically, there was no protection in editor implementations to avoid *ex global*, *v*, @, or * commands changing edit buffers during execution of their associated commands. Because this would almost invariably result in catastrophic failure of the editor, and implementations exist that do exhibit these problems, IEEE Std 1003.1-2001 requires that changing the edit buffer during a *global* or *v* command, or during a @ or * command for which there will be more than a single execution, be an error. Implementations supporting multiple edit buffers simultaneously are strongly encouraged to apply the same semantics to switching between buffers as well.

The *ex* command quoting required by IEEE Std 1003.1-2001 is a superset of the quoting in historical implementations of the editor. For example, it was not historically possible to escape a

<b

Backslash quoting in ex is non-intuitive. Backslash escapes are ignored unless they escape a special character; for example, when performing file argument expansion, the string "\\%" is equivalent to '\%', not "\<current pathname>". This can be confusing for users because backslash is usually one of the characters that causes shell expansion to be performed, and therefore shell quoting rules must be taken into consideration. Generally, quoting characters are only considered if they escape a special character, and a quoting character must be provided for each layer of parsing for which the character is special. As another example, only a single backslash is necessary for the '\1' sequence in substitute replacement patterns, because the character '1' is not special to any parsing layer above it.

<control>-V quoting in ex is slightly different from backslash quoting. In the four commands where <control>-V quoting applies (abbreviate, unabbreviate, map, and unmap), any character may be escaped by a <control>-V whether it would have a special meaning or not. IEEE Std 1003.1-2001 requires conformance to historical practice.

Historical implementations of the editor did not require delimiters within character classes to be escaped; for example, the command :s/[/]// on the string "xxx/yyy" would delete the '/' from the string. IEEE Std 1003.1-2001 disallows this historical practice for consistency and because it places a large burden on implementations by requiring that knowledge of regular expressions be built into the editor parser.

Historically, quoting <newline>s in *ex* commands was handled inconsistently. In most cases, the <newline> always terminated the command, regardless of any preceding escape character, because backslash characters did not escape <newline>s for most *ex* commands. However, some *ex* commands (for example, **s**, **map**, and **abbreviation**) permitted <newline>s to be escaped (although in the case of **map** and **abbreviation**, <control>-V characters escaped them instead of backslashes). This was true in not only the command line, but also .exrc and sourced files. For example, the command:

 would succeed, although it was sometimes difficult to get the <control>-V and the inserted <newline> passed to the *ex* parser. For consistency and simplicity of specification, IEEE Std 1003.1-2001 requires that it be possible to escape <newline>s in *ex* commands at all times, using backslashes for most *ex* commands, and using <control>-V characters for the **map** and **abbreviation** commands. For example, the command **print**<newline>**list** is required to be parsed as the single command **print**<newline>**list**. While this differs from historical practice, IEEE Std 1003.1-2001 developers believed it unlikely that any script or user depended on the historical behavior.

Historically, an error in a command specified using the –c option did not cause the rest of the –c commands to be discarded. IEEE Std 1003.1-2001 disallows this for consistency with mapped keys, the @, global, source, and v commands, the *EXINIT* environment variable, and the .exrc files.

Input Editing in ex

One of the common uses of the historical *ex* editor is over slow network connections. Editors that run in canonical mode can require far less traffic to and from, and far less processing on, the host machine, as well as more easily supporting block-mode terminals. For these reasons, IEEE Std 1003.1-2001 requires that *ex* be implemented using canonical mode input processing, as was done historically.

IEEE Std 1003.1-2001 does not require the historical 4 BSD input editing characters "word erase" or "literal next". For this reason, it is unspecified how they are handled by *ex*, although they must have the required effect. Implementations that resolve them after the line has been ended using a <newline> or <control>-M character, and implementations that rely on the underlying system terminal support for this processing, are both conforming. Implementations are strongly urged to use the underlying system functionality, if at all possible, for compatibility with other system text input interfaces.

Historically, when the *eof* character was used to decrement the **autoindent** level, the cursor moved to display the new end of the **autoindent** characters, but did not move the cursor to a new line, nor did it erase the <control>-D character from the line. IEEE Std 1003.1-2001 does not specify that the cursor remain on the same line or that the rest of the line is erased; however, implementations are strongly encouraged to provide the best possible user interface; that is, the cursor should remain on the same line, and any <control>-D character on the line should be erased

IEEE Std 1003.1-2001 does not require the historical 4 BSD input editing character "reprint", traditionally <control>-R, which redisplayed the current input from the user. For this reason, and because the functionality cannot be implemented after the line has been terminated by the user, IEEE Std 1003.1-2001 makes no requirements about this functionality. Implementations are strongly urged to make this historical functionality available, if possible.

Historically, <control>-Q did not perform a literal next function in *ex*, as it did in *vi*. IEEE Std 1003.1-2001 requires conformance to historical practice to avoid breaking historical *ex* scripts and .exrc files.

eof

Whether the *eof* character immediately modifies the **autoindent** characters in the prompt is left unspecified so that implementations can conform in the presence of systems that do not support this functionality. Implementations are encouraged to modify the line and redisplay it immediately, if possible.

The specification of the handling of the *eof* character differs from historical practice only in that *eof* characters are not discarded if they follow normal characters in the text input. Historically, they were always discarded.

Command Descriptions in ex

Historically, several commands (for example, **global**, **v**, **visual**, **s**, **write**, **wq**, **yank**, **!**, <, >, &, and \rightarrow were executable in empty files (that is, the default address(es) were 0), or permitted explicit addresses of 0 (for example, 0 was a valid address, or 0,0 was a valid range). Addresses of 0, or command execution in an empty file, make sense only for commands that add new text to the edit buffer or write commands (because users may wish to write empty files). IEEE Std 1003.1-2001 requires this behavior for such commands and disallows it otherwise, for consistency and simplicity of specification.

A count to an *ex* command has been historically corrected to be no greater than the last line in a file; for example, in a five-line file, the command **1,6print** would fail, but the command **1print300** would succeed. IEEE Std 1003.1-2001 requires conformance to historical practice.

Historically, the use of flags in *ex* commands could be obscure. General historical practice was as described by IEEE Std 1003.1-2001, but there were some special cases. For instance, the **list**, **number**, and **print** commands ignored trailing address offsets; for example, **3p** +++# would display line 3, and 3 would be the current line after the execution of the command. The **open** and **visual** commands ignored both the trailing offsets and the trailing flags. Also, flags specified to the **open** and **visual** commands interacted badly with the **list** edit option, and setting and then unsetting it during the open/visual session would cause *vi* to stop displaying lines in the specified format. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit any of these exceptions to the general rule.

IEEE Std 1003.1-2001 uses the word *copy* in several places when discussing buffers. This is not intended to imply implementation.

Historically, *ex* users could not specify numeric buffers because of the ambiguity this would cause; for example, in the command **3 delete 2**, it is unclear whether 2 is a buffer name or a *count*. IEEE Std 1003.1-2001 requires conformance to historical practice by default, but does not preclude extensions.

Historically, the contents of the unnamed buffer were frequently discarded after commands that did not explicitly affect it; for example, when using the **edit** command to switch files. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior.

The *ex* utility did not historically have access to the numeric buffers, and, furthermore, deleting lines in *ex* did not modify their contents. For example, if, after doing a delete in *vi*, the user switched to *ex*, did another delete, and then switched back to *vi*, the contents of the numeric buffers would not have changed. IEEE Std 1003.1-2001 requires conformance to historical practice. Numeric buffers are described in the *ex* utility in order to confine the description of buffers to a single location in IEEE Std 1003.1-2001.

The metacharacters that trigger shell expansion in *file* arguments match historical practice, as does the method for doing shell expansion. Implementations wishing to provide users with the flexibility to alter the set of metacharacters are encouraged to provide a **shellmeta** string edit

15906 option.

15907 15908

15909

15910

15911

15912

15913 15914

15915 15916

15917

15918

15919

15920

15921

15922

15923

15924

15925 15926

15927

15928

15929

15930

15931

15935

15938

15939

15940

15941

Historically, ex commands executed from vi refreshed the screen when it did not strictly need to do so; for example, :!date > /dev/null does not require a screen refresh because the output of the UNIX date command requires only a single line of the screen. IEEE Std 1003.1-2001 requires that the screen be refreshed if it has been overwritten, but makes no requirements as to how an implementation should make that determination. Implementations may prompt and refresh the screen regardless.

Abbreviate

Historical practice was that characters that were entered as part of an abbreviation replacement were subject to map expansions, the showmatch edit option, further abbreviation expansions, and so on; that is, they were logically pushed onto the terminal input queue, and were not a simple replacement. IEEE Std 1003.1-2001 requires conformance to historical practice. Historical practice was that whenever a non-word character (that had not been escaped by a <control>-V) was entered after a word character, vi would check for abbreviations. The check was based on the type of the character entered before the word character of the word/non-word pair that triggered the check. The word character of the word/non-word pair that triggered the check and all characters entered before the trigger pair that were of that type were included in the check, with the exception of <blank>s, which always delimited the abbreviation.

This means that, for the abbreviation to work, the *lhs* must end with a word character, there can be no transitions from word to non-word characters (or vice versa) other than between the last and next-to-last characters in the *lhs*, and there can be no <blank>s in the *lhs*. In addition, because of the historical quoting rules, it was impossible to enter a literal <control>-V in the *lhs*. IEEE Std 1003.1-2001 requires conformance to historical practice. Historical implementations did not inform users when abbreviations that could never be used were entered; implementations are strongly encouraged to do so.

For example, the following abbreviations will work:

```
15932
            :ab (p
                     REPLACE
15933
                      REPLACE
            :ab p
15934
            :ab ((p REPLACE
```

The following abbreviations will not work:

```
REPLACE
15936
             :ab (
15937
             :ab (pp REPLACE
```

Historical practice is that words on the *vi* colon command line were subject to abbreviation expansion, including the arguments to the abbrev (and more interestingly) the unabbrev command. Because there are implementations that do not do abbreviation expansion for the first argument to those commands, this is permitted, but not required, by IEEE Std 1003.1-2001. However, the following sequence:

15942

```
:ab foo bar
15943
             :ab foo baz
15944
```

resulted in the addition of an abbreviation of "baz" for the string "bar" in historical ex/vi, and 15945 15946 the sequence:

```
15947
             :ab fool bar
15948
             :ab foo2 bar
```

:unabbreviate foo2 15949

deleted the abbreviation "foo1", not "foo2". These behaviors are not permitted by IEEE Std 1003.1-2001 because they clearly violate the expectations of the user.

It was historical practice that <control>-V, not backslash, characters be interpreted as escaping subsequent characters in the **abbreviate** command. IEEE Std 1003.1-2001 requires conformance to historical practice; however, it should be noted that an abbreviation containing a <blank> will never work.

Append

Historically, any text following a vertical-line command separator after an **append**, **change**, or **insert** command became part of the insert text. For example, in the command:

:g/pattern/append|stuff1

a line containing the text "stuff1" would be appended to each line matching pattern. It was also historically valid to enter:

15962 :append | stuff1

15963 stuff2

15964 .

and the text on the *ex* command line would be appended along with the text inserted after it. There was an historical bug, however, that the user had to enter two terminating lines (the '.' lines) to terminate text input mode in this case. IEEE Std 1003.1-2001 requires conformance to historical practice, but disallows the historical need for multiple terminating lines.

Change

See the RATIONALE for the **append** command. Historical practice for cursor positioning after the change command when no text is input, is as described in IEEE Std 1003.1-2001. However, one System V implementation is known to have been modified such that the cursor is positioned on the first address specified, and not on the line before the first address. IEEE Std 1003.1-2001 disallows this modification for consistency.

Historically, the **change** command did not support buffer arguments, although some implementations allow the specification of an optional buffer. This behavior is neither required nor disallowed by IEEE Std 1003.1-2001.

Change Directory

A common extension in *ex* implementations is to use the elements of a **cdpath** edit option as prefix directories for *path* arguments to **chdir** that are relative pathnames and that do not have '.' or ".." as their first component. Elements in the **cdpath** edit option are colon-separated. The initial value of the **cdpath** edit option is the value of the shell *CDPATH* environment variable. This feature was not included in IEEE Std 1003.1-2001 because it does not exist in any of the implementations considered historical practice.

Сору

Historical implementations of *ex* permitted copies to lines inside of the specified range; for example, **:2,5copy3** was a valid command. IEEE Std 1003.1-2001 requires conformance to historical practice.

Delete 15989 IEEE Std 1003.1-2001 requires support for the historical parsing of a delete command followed 15990 15991 by flags, without any intervening <blank>s. For example: Deletes the first line and prints the line that was second. 15992 **1delep** As for **1dp**. 15993 1d Deletes the first line, saving it in buffer *p*. 15994 **1d p1l** (Pee-one-ell.) Deletes the first line, saving it in buffer p, and listing the line that was 15995 15996 second. **Edit** 15997 Historically, any ex command could be entered as a +command argument to the edit command, 15998 although some (for example, insert and append) were known to confuse historical 15999 implementations. For consistency and simplicity of specification, IEEE Std 1003.1-2001 requires 16000 16001 that any command be supported as an argument to the **edit** command. Historically, the command argument was executed with the current line set to the last line of the 16002 file, regardless of whether the edit command was executed from visual mode or not. 16003 16004 IEEE Std 1003.1-2001 requires conformance to historical practice. Historically, the +command specified to the edit and next commands was delimited by the first 16005
<blank>, and there was no way to quote them. For consistency, IEEE Std 1003.1-2001 requires 16006 that the usual ex backslash quoting be provided. 16007 Historically, specifying the +command argument to the edit command required a filename to be 16008 specified as well; for example, :edit +100 would always fail. For consistency and simplicity of 16009 specification, IEEE Std 1003.1-2001 does not permit this usage to fail for that reason. 16010 Historically, only the cursor position of the last file edited was remembered by the editor. 16011 IEEE Std 1003.1-2001 requires that this be supported; however, implementations are permitted to 16012 16013 remember and restore the cursor position for any file previously edited. File 16014 Historical versions of the ex editor **file** command displayed a current line and number of lines in 16015 the edit buffer of 0 when the file was empty, while the vi <control>-G command displayed a 16016 current line and number of lines in the edit buffer of 1 in the same situation. 16017 IEEE Std 1003.1-2001 does not permit this discrepancy, instead requiring that a message be 16018 16019 displayed indicating that the file is empty. Global 16020 The two-pass operation of the **global** and **v** commands is not intended to imply implementation, 16021 only the required result of the operation. 16022 16023 The current line and column are set as specified for the individual ex commands. This requirement is cumulative; that is, the current line and column must track across all the 16024

commands executed by the **global** or **v** commands.

16025

16026 Insert

See the RATIONALE for the **append** command.

Historically, **insert** could not be used with an address of zero; that is, not when the edit buffer was empty. IEEE Std 1003.1-2001 requires that this command behave consistently with the **append** command.

Join

The action of the **join** command in relation to the special characters is only defined for the POSIX locale because the correct amount of white space after a period varies; in Japanese none is required, in French only a single space, and so on.

List

The historical output of the **list** command was potentially ambiguous. The standard developers believed correcting this to be more important than adhering to historical practice, and IEEE Std 1003.1-2001 requires unambiguous output.

Map

Historically, command mode maps only applied to command names; for example, if the character 'x' was mapped to 'y', the command **fx** searched for the 'x' character, not the 'y' character. IEEE Std 1003.1-2001 requires this behavior. Historically, entering <control>-V as the first character of a *vi* command was an error. Several implementations have extended the semantics of *vi* such that <control>-V means that the subsequent command character is not mapped. This is permitted, but not required, by IEEE Std 1003.1-2001. Regardless, using <control>-V to escape the second or later character in a sequence of characters that might match a **map** command, or any character in text input mode, is historical practice, and stops the entered keys from matching a map. IEEE Std 1003.1-2001 requires conformance to historical practice.

Historical implementations permitted digits to be used as a **map** command *lhs*, but then ignored the map. IEEE Std 1003.1-2001 requires that the mapped digits not be ignored.

The historical implementation of the **map** command did not permit **map** commands that were more than a single character in length if the first character was printable. This behavior is permitted, but not required, by IEEE Std 1003.1-2001.

Historically, mapped characters were remapped unless the **remap** edit option was not set, or the prefix of the mapped characters matched the mapping characters; for example, in the **map**:

16056 :map ab abcd

the characters "ab" were used as is and were not remapped, but the characters "cd" were mapped if appropriate. This can cause infinite loops in the *vi* mapping mechanisms. IEEE Std 1003.1-2001 requires conformance to historical practice, and that such loops be interruptible.

Text input maps had the same problems with expanding the *lhs* for the *ex* **map!** and **unmap!** command as did the *ex* **abbreviate** and **unabbreviate** commands. See the RATIONALE for the *ex* **abbreviate** command. IEEE Std 1003.1-2001 requires similar modification of some historical practice for the **map** and **unmap** commands, as described for the **abbreviate** and **unabbreviate** commands.

Historically, **map**s that were subsets of other **map**s behaved differently depending on the order in which they were defined. For example:

16068	:map!	ab	short
16069	:map!	abc	long

would always translate the characters "ab" to "short", regardless of how fast the characters "abc" were entered. If the entry order was reversed:

```
16072 :map! abc long
16073 :map! ab short
```

the characters "ab" would cause the editor to pause, waiting for the completing 'c' character, and the characters might never be mapped to "short". For consistency and simplicity of specification, IEEE Std 1003.1-2001 requires that the shortest match be used at all times.

The length of time the editor spends waiting for the characters to complete the *lhs* is unspecified because the timing capabilities of systems are often inexact and variable, and it may depend on other factors such as the speed of the connection. The time should be long enough for the user to be able to complete the sequence, but not long enough for the user to have to wait. Some implementations of *vi* have added a **keytime** option, which permits users to set the number of 0,1 seconds the editor waits for the completing characters. Because mapped terminal function and cursor keys tend to start with an <ESC> character, and <ESC> is the key ending *vi* text input mode, **maps** starting with <ESC> characters are generally exempted from this timeout period, or, at least timed out differently.

Mark

Historically, users were able to set the "previous context" marks explicitly. In addition, the *ex* commands" and "and the *vi* commands", ", ", and "all referred to the same mark. In addition, the previous context marks were not set if the command, with which the address setting the mark was associated, failed. IEEE Std 1003.1-2001 requires conformance to historical practice. Historically, if marked lines were deleted, the mark was also deleted, but would reappear if the change was undone. IEEE Std 1003.1-2001 requires conformance to historical practice.

The description of the special events that set the 'and 'marks matches historical practice. For example, historically the command /a/,/b/ did not set the 'and 'marks, but the command /a/,/b/ delete did.

Next

Historically, any ex command could be entered as a +command argument to the **next** command, although some (for example, **insert** and **append**) were known to confuse historical implementations. IEEE Std 1003.1-2001 requires that any command be permitted and that it behave as specified. The **next** command can accept more than one file, so usage such as:

```
16101 next 'ls [abc] '
```

is valid; it need not be valid for the **edit** or **read** commands, for example, because they expect only one filename.

Historically, the **next** command behaved differently from the **:rewind** command in that it ignored the force flag if the **autowrite** flag was set. For consistency, IEEE Std 1003.1-2001 does not permit this behavior.

Historically, the **next** command positioned the cursor as if the file had never been edited before, regardless. IEEE Std 1003.1-2001 does not permit this behavior, for consistency with the **edit** command.

Implementations wanting to provide a counterpart to the **next** command that edited the previous file have used the command **prev[ious]**, which takes no *file* argument.

16112 IEEE Std 1003.1-2001 does not require this command.

Open

 Historically, the **open** command would fail if the **open** edit option was not set. IEEE Std 1003.1-2001 does not mention the **open** edit option and does not require this behavior. Some historical implementations do not permit entering open mode from open or visual mode, only from *ex* mode. For consistency, IEEE Std 1003.1-2001 does not permit this behavior.

Historically, entering open mode from the command line (that is, *vi* +**open**) resulted in anomalous behaviors; for example, the *ex* file and *set* commands, and the *vi* command <control>-G did not work. For consistency, IEEE Std 1003.1-2001 does not permit this behavior.

Historically, the **open** command only permitted '/' characters to be used as the search pattern delimiter. For consistency, IEEE Std 1003.1-2001 requires that the search delimiters used by the **s**, **global**, and **v** commands be accepted as well.

Preserve

The **preserve** command does not historically cause the file to be considered unmodified for the purposes of future commands that may exit the editor. IEEE Std 1003.1-2001 requires conformance to historical practice.

Historical documentation stated that mail was not sent to the user when preserve was executed; however, historical implementations did send mail in this case. IEEE Std 1003.1-2001 requires conformance to the historical implementations.

Print

The writing of NUL by the **print** command is not specified as a special case because the standard developers did not want to require *ex* to support NUL characters. Historically, characters were displayed using the ARPA standard mappings, which are as follows:

- 1. Printable characters are left alone.
- 2. Control characters less than \177 are represented as '\^' followed by the character offset from the '@' character in the ASCII map; for example, \007 is represented as '\^G'.
- 3. \177 is represented as '^' followed by '?'.

The display of characters having their eighth bit set was less standard. Existing implementations use hex (0x00), octal (\setminus 000), and a meta-bit display. (The latter displayed bytes that had their eighth bit set as the two characters "M-" followed by the seven-bit display as described above.) The latter probably has the best claim to historical practice because it was used for the $-\mathbf{v}$ option of 4 BSD and 4 BSD-derived versions of the cat utility since 1980.

No specific display format is required by IEEE Std 1003.1-2001.

Explicit dependence on the ASCII character set has been avoided where possible, hence the use of the phrase an "implementation-defined multi-character sequence" for the display of non-printable characters in preference to the historical usage of, for instance, "^I" for the <tab>. Implementations are encouraged to conform to historical practice in the absence of any strong reason to diverge.

Historically, all ex commands beginning with the letter 'p' could be entered using capitalized versions of the commands; for example, P[rint], Pre[serve], and Pu[t] were all valid command names. IEEE Std 1003.1-2001 permits, but does not require, this historical practice because capital forms of the commands are used by some implementations for other purposes.

Put

Historically, an *ex* **put** command, executed from open or visual mode, was the same as the open or visual mode **P** command, if the buffer was named and was cut in character mode, and the same as the **p** command if the buffer was named and cut in line mode. If the unnamed buffer was the source of the text, the entire line from which the text was taken was usually **put**, and the buffer was handled as if in line mode, but it was possible to get extremely anomalous behavior. In addition, using the **Q** command to switch into *ex* mode, and then doing a **put** often resulted in errors as well, such as appending text that was unrelated to the (supposed) contents of the buffer. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit these behaviors. All *ex* **put** commands are required to operate in line mode, and the contents of the buffers are not altered by changing the mode of the editor.

Read

Historically, an *ex* **read** command executed from open or visual mode, executed in an empty file, left an empty line as the first line of the file. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior. Historically, a **read** in open or visual mode from a program left the cursor at the last line read in, not the first. For consistency, IEEE Std 1003.1-2001 does not permit this behavior.

Historical implementations of *ex* were unable to undo **read** commands that read from the output of a program. For consistency, IEEE Std 1003.1-2001 does not permit this behavior.

Historically, the *ex* and *vi* message after a successful **read** or **write** command specified "characters", not "bytes". IEEE Std 1003.1-2001 requires that the number of bytes be displayed, not the number of characters, because it may be difficult in multi-byte implementations to determine the number of characters read. Implementations are encouraged to clarify the message displayed to the user.

Historically, reads were not permitted on files other than type regular, except that FIFO files could be read (probably only because they did not exist when *ex* and *vi* were originally written). Because the historical *ex* evaluated **read!** and **read!** equivalently, there can be no optional way to force the read. IEEE Std 1003.1-2001 permits, but does not require, this behavior.

Recover

Some historical implementations of the editor permitted users to recover the edit buffer contents from a previous edit session, and then exit without saving those contents (or explicitly discarding them). The intent of IEEE Std 1003.1-2001 in requiring that the edit buffer be treated as already modified is to prevent this user error.

Rewind

Historical implementations supported the **rewind** command when the user was editing the first file in the list; that is, the file that the **rewind** command would edit. IEEE Std 1003.1-2001 requires conformance to historical practice.

16191	Substitute
16192 16193 16194 16195 16196	Historically, <i>ex</i> accepted an r option to the s command. The effect of the r option was to use the last regular expression used in any command as the pattern, the same as the ~command. The r option is not required by IEEE Std 1003.1-2001. Historically, the c and g options were toggled; for example, the command : s / abc / def / was the same as s / abc / def /cccc gggg . For simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior.
16197	The tilde command is often used to replace the last search RE. For example, in the sequence:
16198 16199 16200	s/red/blue/ /green ~
16201	the ~ command is equivalent to:
16202	s/green/blue/
16203	Historically, ex accepted all of the following forms:
16204 16205 16206 16207	s/abc/def/ s/abc/def s/abc/ s/abc
16208	IEEE Std 1003.1-2001 requires conformance to this historical practice.
16209 16210 16211	The s command presumes that the '^' character only occupies a single column in the display. Much of the <i>ex</i> and <i>vi</i> specification presumes that the <space> only occupies a single column in the display. There are no known character sets for which this is not true.</space>
16212 16213 16214 16215 16216 16217	Historically, the final column position for the substitute commands was based on previous column movements; a search for a pattern followed by a substitution would leave the column position unchanged, while a 0 command followed by a substitution would change the column position to the first non- volume to the first non-
16218	Set
16219 16220	Historical implementations redisplayed all of the options for each occurrence of the all keyword. IEEE Std 1003.1-2001 permits, but does not require, this behavior.
16221	Tag
16222 16223 16224 16225 16226 16227	No requirement is made as to where <i>ex</i> and <i>vi</i> shall look for the file referenced by the tag entry. Historical practice has been to look for the path found in the tags file, based on the current directory. A useful extension found in some implementations is to look based on the directory containing the tags file that held the entry, as well. No requirement is made as to which reference for the tag in the tags file is used. This is deliberate, in order to permit extensions such as multiple entries in a tags file for a tag.
16228 16229 16230 16231	Because users often specify many different tags files, some of which need not be relevant or exist at any particular time, IEEE Std 1003.1-2001 requires that error messages about problem tags files be displayed only if the requested tag is not found, and then, only once for each time that the tag edit option is changed.
16232 16233	The requirement that the current edit buffer be unmodified is only necessary if the file indicated by the tag entry is not the same as the current file (as defined by the current pathname).

Historically, the file would be reloaded if the filename had changed, as well as if the filename was different from the current pathname. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior, requiring that the name be the only factor in the decision.

Historically, *vi* only searched for tags in the current file from the current cursor to the end of the file, and therefore, if the **wrapscan** option was not set, tags occurring before the current cursor were not found. IEEE Std 1003.1-2001 considers this a bug, and implementations are required to search for the first occurrence in the file, regardless.

Undo

The **undo** description deliberately uses the word "modified". The **undo** command is not intended to undo commands that replace the contents of the edit buffer, such as **edit**, **next**, **tag**, or **recover**.

Cursor positioning after the **undo** command was inconsistent in the historical *vi*, sometimes attempting to restore the original cursor position (**global**, **undo**, and **v** commands), and sometimes, in the presence of maps, placing the cursor on the last line added or changed instead of the first. IEEE Std 1003.1-2001 requires a simplified behavior for consistency and simplicity of specification.

Version

The **version** command cannot be exactly specified since there is no widely-accepted definition of what the version information should contain. Implementations are encouraged to do something reasonably intelligent.

Write

Historically, the *ex* and *vi* message after a successful **read** or **write** command specified "characters", not "bytes". IEEE Std 1003.1-2001 requires that the number of bytes be displayed, not the number of characters because it may be difficult in multi-byte implementations to determine the number of characters written. Implementations are encouraged to clarify the message displayed to the user.

Implementation-defined tests are permitted so that implementations can make additional checks; for example, for locks or file modification times.

Historically, attempting to append to a nonexistent file caused an error. It has been left unspecified in IEEE Std 1003.1-2001 to permit implementations to let the **write** succeed, so that the append semantics are similar to those of the historical *csh*.

Historical *vi* permitted empty edit buffers to be written. However, since the way *vi* got around dealing with "empty" files was to always have a line in the edit buffer, no matter what, it wrote them as files of a single, empty line. IEEE Std 1003.1-2001 does not permit this behavior.

Historically, ex restored standard output and standard error to their values as of when ex was invoked, before writes to programs were performed. This could disturb the terminal configuration as well as be a security issue for some terminals. IEEE Std 1003.1-2001 does not permit this, requiring that the program output be captured and displayed as if by the ex print command.

Adjust Window

Historically, the line count was set to the value of the **scroll** option if the type character was end-of-file. This feature was broken on most historical implementations long ago, however, and is not documented anywhere. For this reason, IEEE Std 1003.1-2001 is resolutely silent.

Historically, the z command was
 sensitive and z + and z - did different things than z+ and z- because the type could not be distinguished from a flag. (The commands z. and z = were historically invalid.) IEEE Std 1003.1-2001 requires conformance to this historical practice.

Historically, the **z** command was further <blank>-sensitive in that the *count* could not be <blank>-delimited; for example, the commands **z**= **5** and **z**- **5** were also invalid. Because the *count* is not ambiguous with respect to either the type character or the flags, this is not permitted by IEEE Std 1003.1-2001.

Escape

Historically, *ex* filter commands only read the standard output of the commands, letting standard error appear on the terminal as usual. The *vi* utility, however, read both standard output and standard error. IEEE Std 1003.1-2001 requires the latter behavior for both *ex* and *vi*, for consistency.

Shift Left and Shift Right

Historically, it was possible to add shift characters to increase the effect of the command; for example, <<< outdented (or >>> indented) the lines 3 levels of indentation instead of the default 1. IEEE Std 1003.1-2001 requires conformance to historical practice.

<control>-D

Historically, the <control>-D command erased the prompt, providing the user with an unbroken presentation of lines from the edit buffer. This is not required by IEEE Std 1003.1-2001; implementations are encouraged to provide it if possible. Historically, the <control>-D command took, and then ignored, a *count*. IEEE Std 1003.1-2001 does not permit this behavior.

Write Line Number

Historically, the *ex* = command, when executed in *ex* mode in an empty edit buffer, reported 0, and from open or visual mode, reported 1. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior.

Execute

Historically, *ex* did not correctly handle the inclusion of text input commands (that is, **append**, **insert**, and **change**) in executed buffers. IEEE Std 1003.1-2001 does not permit this exclusion for consistency.

Historically, the logical contents of the buffer being executed did not change if the buffer itself were modified by the commands being executed; that is, buffer execution did not support self-modifying code. IEEE Std 1003.1-2001 requires conformance to historical practice.

Historically, the @ command took a range of lines, and the @ buffer was executed once per line, with the current line (' . ') set to each specified line. IEEE Std 1003.1-2001 requires conformance to historical practice.

Some historical implementations did not notice if errors occurred during buffer execution. This, coupled with the ability to specify a range of lines for the *ex* @ command, makes it trivial to cause them to drop **core**. IEEE Std 1003.1-2001 requires that implementations stop buffer

16316 execution if any error occurs, if the specified line doesn't exist, or if the contents of the edit buffer 16317 itself are replaced (for example, the buffer executes the *ex* :edit command). Regular Expressions in ex 16318 Historical practice is that the characters in the replacement part of the last **s** command—that is, 16319 those matched by entering a '~' in the regular expression—were not further expanded by the 16320 regular expression engine. So, if the characters contained the string "a.," they would match 16321 'a' followed by "., " and not 'a' followed by any character. IEEE Std 1003.1-2001 requires 16322 conformance to historical practice. 16323 **Edit Options in ex** 16324 16325 The following paragraphs describe the historical behavior of some edit options that were not, for 16326 whatever reason, included in IEEE Std 1003.1-2001. Implementations are strongly encouraged to only use these names if the functionality described here is fully supported. 16327 extended 16328 The **extended** edit option has been used in some implementations of *vi* to provide extended regular expressions instead of basic regular expressions This option was 16329 16330 omitted from IEEE Std 1003.1-2001 because it is not widespread historical practice. flash The **flash** edit option historically caused the screen to flash instead of beeping on 16331 error. This option was omitted from IEEE Std 1003.1-2001 because it is not found in 16332 some historical implementations. 16333 16334 hardtabs The hardtabs edit option historically defined the number of columns between hardware tab settings. This option was omitted from IEEE Std 1003.1-2001 because 16335 it was believed to no longer be generally useful. 16336 modeline The **modeline** (sometimes named **modelines**) edit option historically caused ex or 16337 *vi* to read the five first and last lines of the file for editor commands. This option is 16338 a security problem, and vendors are strongly encouraged to delete it from 16339 historical implementations. 16340 The **open** edit option historically disallowed the *ex* **open** and **visual** commands. open 16341 16342 This edit option was omitted because these commands are required by 16343 IEEE Std 1003.1-2001. optimize The optimize edit option historically expedited text throughput by setting the 16344 16345 terminal to not do automatic <carriage-return>s when printing more than one logical line of output. This option was omitted from IEEE Std 1003.1-2001 because 16346 it was intended for terminals without addressable cursors, which are rarely, if ever, 16347 still used. 16348 ruler The **ruler** edit option has been used in some implementations of vi to present a 16349 current row/column ruler for the user. This option was omitted from 16350 IEEE Std 1003.1-2001 because it is not widespread historical practice. 16351 sourceany The **sourceany** edit option historically caused ex or vi to source start-up files that 16352 were owned by users other than the user running the editor. This option is a 16353 16354 security problem, and vendors are strongly encouraged to remove it from their 16355 implementations. timeout The **timeout** edit option historically enabled the (now standard) feature of only 16356

16357

16358 16359 waiting for a short period before returning keys that could be part of a macro. This

feature was omitted from IEEE Std 1003.1-2001 because its behavior is now

standard, it is not widely useful, and it was rarely documented.

16360 verbose The **verbose** edit option has been used in some implementations of *vi* to cause *vi* to 16361 output error messages for common errors; for example, attempting to move the cursor past the beginning or end of the line instead of only alerting the screen. (The 16362 historical vi only alerted the terminal and presented no message for such errors. 16363 16364 The historical editor option terse did not select when to present error messages, it only made existing error messages more or less verbose.) This option was omitted 16365 from IEEE Std 1003.1-2001 because it is not widespread historical practice; 16366 however, implementors are encouraged to use it if they wish to provide error 16367 messages for naive users. 16368 wraplen The **wraplen** edit option has been used in some implementations of *vi* to specify an 16369 automatic margin measured from the left margin instead of from the right margin. 16370 This is useful when multiple screen sizes are being used to edit a single file. This 16371 option was omitted from IEEE Std 1003.1-2001 because it is not widespread 16372 historical practice; however, implementors are encouraged to use it if they add this 16373 functionality. 16374 autoindent, ai 16375 16376

Historically, the command 0a did not do any autoindentation, regardless of the current indentation of line 1. IEEE Std 1003.1-2001 requires that any indentation present in line 1 be used.

autoprint, ap

16377

16378

16379

16380

16381

16382

16383

16384

16385 16386

16387

16388

16389

16390 16391

16392 16393

16394

16395

16396 16397

16398

16399

16400

Historically, the autoprint edit option was not completely consistent or based solely on modifications to the edit buffer. Exceptions were the **read** command (when reading from a file, but not from a filter), the append, change, insert, global, and v commands, all of which were not affected by autoprint, and the tag command, which was affected by autoprint. IEEE Std 1003.1-2001 requires conformance to historical practice.

Historically, the autoprint option only applied to the last of multiple commands entered using vertical-bar delimiters; for example, delete <newline> was affected by autoprint, but delete version <newline> was not. IEEE Std 1003.1-2001 requires conformance to historical practice.

autowrite, aw

Appending the '!' character to the ex next command to avoid performing an automatic write was not supported in historical implementations. IEEE Std 1003.1-2001 requires that the behavior match the other *ex* commands for consistency.

ignorecase, ic

Historical implementations of case-insensitive matching (the ignorecase edit option) lead to counterintuitive situations when uppercase characters were used in range expressions. Historically, the process was as follows:

- Take a line of text from the edit buffer.
- Convert uppercase to lowercase in text line.
- 3. Convert uppercase to lowercase in regular expressions, except in character class specifications.
- 4. Match regular expressions against text.
- 16401 This would mean that, with **ignorecase** in effect, the text:

16402 The cat sat on the mat 16403 would be matched by 16404 /^the/ 16405 but not by: 16406 /^[A-Z]he/

For consistency with other commands implementing regular expressions, IEEE Std 1003.1-2001 does not permit this behavior.

paragraphs, para

The ISO POSIX-2: 1993 standard made the default **paragraphs** and **sections** edit options implementation-defined, arguing they were historically oriented to the UNIX system *troff* text formatter, and a "portable user" could use the {, }, [[,]], (, and) commands in open or visual mode and have the cursor stop in unexpected places. IEEE Std 1003.1-2001 specifies their values in the POSIX locale because the unusual grouping (they only work when grouped into two characters at a time) means that they cannot be used for general-purpose movement, regardless.

readonly

Implementations are encouraged to provide the best possible information to the user as to the read-only status of the file, with the exception that they should not consider the current special privileges of the process. This provides users with a safety net because they must force the overwrite of read-only files, even when running with additional privileges.

The **readonly** edit option specification largely conforms to historical practice. The only difference is that historical implementations did not notice that the user had set the **readonly** edit option in cases where the file was already marked read-only for some reason, and would therefore reinitialize the **readonly** edit option the next time the contents of the edit buffer were replaced. This behavior is disallowed by IEEE Std 1003.1-2001.

report

The requirement that lines copied to a buffer interact differently than deleted lines is historical practice. For example, if the **report** edit option is set to 3, deleting 3 lines will cause a report to be written, but 4 lines must be copied before a report is written.

The requirement that the *ex* **global**, **v**, **open**, **undo**, and **visual** commands present reports based on the total number of lines added or deleted during the command execution, and that commands executed by the **global** and **v** commands not present reports, is historical practice. IEEE Std 1003.1-2001 extends historical practice by requiring that buffer execution be treated similarly. The reasons for this are two-fold. Historically, only the report by the last command executed from the buffer would be seen by the user, as each new report would overwrite the last. In addition, the standard developers believed that buffer execution had more in common with **global** and **v** commands than it did with other *ex* commands, and should behave similarly, for consistency and simplicity of specification.

showmatch, sm

The length of time the cursor spends on the matching character is unspecified because the timing capabilities of systems are often inexact and variable. The time should be long enough for the user to notice, but not long enough for the user to become annoyed. Some implementations of *vi* have added a **matchtime** option that permits users to set the number of 0,1 second intervals the cursor pauses on the matching character.

showmode

The **showmode** option has been used in some historical implementations of *ex* and *vi* to display the current editing mode when in open or visual mode. The editing modes have generally included "command" and "input", and sometimes other modes such as "replace" and "change". The string was usually displayed on the bottom line of the screen at the far right-hand corner. In addition, a preceding '*' character often denoted whether the contents of the edit buffer had been modified. The latter display has sometimes been part of the **showmode** option, and sometimes based on another option. This option was not available in the 4 BSD historical implementation of *vi*, but was viewed as generally useful, particularly to novice users, and is required by IEEE Std 1003.1-2001.

The **smd** shorthand for the **showmode** option was not present in all historical implementations of the editor. IEEE Std 1003.1-2001 requires it, for consistency.

Not all historical implementations of the editor displayed a mode string for command mode, differentiating command mode from text input mode by the absence of a mode string. IEEE Std 1003.1-2001 permits this behavior for consistency with historical practice, but implementations are encouraged to provide a display string for both modes.

slowopen

Historically the **slowopen** option was automatically set if the terminal baud rate was less than 1 200 baud, or if the baud rate was 1 200 baud and the **redraw** option was not set. The **slowopen** option had two effects. First, when inserting characters in the middle of a line, characters after the cursor would not be pushed ahead, but would appear to be overwritten. Second, when creating a new line of text, lines after the current line would not be scrolled down, but would appear to be overwritten. In both cases, ending text input mode would cause the screen to be refreshed to match the actual contents of the edit buffer. Finally, terminals that were sufficiently intelligent caused the editor to ignore the **slowopen** option. IEEE Std 1003.1-2001 permits most historical behavior, extending historical practice to require **slowopen** behaviors if the edit option is set by the user.

tags

The default path for tags files is left unspecified as implementations may have their own tags implementations that do not correspond to the historical ones. The default tags option value should probably at least include the file ./tags.

term

Historical implementations of *ex* and *vi* ignored changes to the **term** edit option after the initial terminal information was loaded. This is permitted by IEEE Std 1003.1-2001; however, implementations are encouraged to permit the user to modify their terminal type at any time.

terse

Historically, the **terse** edit option optionally provided a shorter, less descriptive error message, for some error messages. This is permitted, but not required, by IEEE Std 1003.1-2001. Historically, most common visual mode errors (for example, trying to move the cursor past the end of a line) did not result in an error message, but simply alerted the terminal. Implementations wishing to provide messages for novice users are urged to do so based on the **edit** option **verbose**, and not **terse**.

window

In historical implementations, the default for the **window** edit option was based on the baud rate as follows:

1. If the baud rate was less than 1200, the **edit** option **w300** set the window value; for example, the line:

set w300=12

would set the window option to 12 if the baud rate was less than 1 200.

- 2. If the baud rate was equal to 1 200, the **edit** option **w1200** set the window value.
- 3. If the baud rate was greater than 1 200, the **edit** option **w9600** set the window value.

The w300, w1200, and w9600 options do not appear in IEEE Std 1003.1-2001 because of their dependence on specific baud rates.

In historical implementations, the size of the window displayed by various commands was related to, but not necessarily the same as, the **window** edit option. For example, the size of the window was set by the *ex* command **visual 10**, but it did not change the value of the **window** edit option. However, changing the value of the **window** edit option did change the number of lines that were displayed when the screen was repainted. IEEE Std 1003.1-2001 does not permit this behavior in the interests of consistency and simplicity of specification, and requires that all commands that change the number of lines that are displayed do it by setting the value of the **window** edit option.

wrapmargin, wm

Historically, the **wrapmargin** option did not affect maps inserting characters that also had associated *counts*; for example :map K 5aABC DEF. Unfortunately, there are widely used maps that depend on this behavior. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior.

Historically, **wrapmargin** was calculated using the column display width of all characters on the screen. For example, an implementation using "^I" to represent <tab>s when the **list** edit option was set, where '^' and 'I' each took up a single column on the screen, would calculate the **wrapmargin** based on a value of 2 for each <tab>. The **number** edit option similarly changed the effective length of the line as well. IEEE Std 1003.1-2001 requires conformance to historical practice.

16517 FUTURE DIRECTIONS 16518 None. **16519 SEE ALSO** Section 2.9.1.1 (on page 48), ctags, ed, sed, sh, stty, vi, the System Interfaces volume of 16520 IEEE Std 1003.1-2001, access() 16521 16522 CHANGE HISTORY First released in Issue 2. 16523 16524 Issue 5 The FUTURE DIRECTIONS section is added. 16525 16526 Issue 6 This utility is marked as part of the User Portability Utilities option. 16527 The obsolescent SYNOPSIS is removed, removing the +command and – options. 16528 16529 The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification: 16530 • In the **map** command description, the sequence #digit is added. 16531 • The **directory**, **edcompatible**, **redraw**, and **slowopen** edit options are added. 16532 The ex utility is extensively changed for alignment with the IEEE P1003.2b draft standard. This 16533 includes changes as a result of the IEEE PASC Interpretations 1003.2 #31, #38, #49, #50, #51, #52, 16534 16535 #55, #56, #57, #61, #62, #63, #64, #65, and #78. The **–l** option is removed. 16536 IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/23 is applied, correcting a URL. 16537

expand Utilities

16538 NAME16539

expand — convert tabs to spaces

16540 SYNOPSIS

16541 UP expand [-t tablist] [file ...]

16542

16551

16552

16553

16554

16555

16556

16557

16558 16559

16560

16561

16562

16563

16564

16543 **DESCRIPTION**

The *expand* utility shall write files or the standard input to the standard output with <tab>s replaced with one or more <space>s needed to pad to the next tab stop. Any <backspace>s shall be copied to the output and cause the column position count for tab stop calculations to be decremented; the column position count shall not be decremented below zero.

16548 OPTIONS

The *expand* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following option shall be supported:

–t tablist

Specify the tab stops. The application shall ensure that the argument *tablist* consists of either a single positive decimal integer or a list of tabstops. If a single number is given, tabs shall be set that number of column positions apart instead of the default 8.

If a list of tabstops is given, the application shall ensure that it consists of a list of two or more positive decimal integers, separated by $\langle blank \rangle$ s or commas, in ascending order. The tabs shall be set at those specific column positions. Each tab stop N shall be an integer value greater than zero, and the list is in strictly ascending order. This is taken to mean that, from the start of a line of output, tabbing to position N shall cause the next character output to be in the (N+1)th column position on that line.

In the event of *expand* having to process a <tab> at a position beyond the last of those specified in a multiple tab-stop list, the <tab> shall be replaced by a single

16565 <space> in the output.

16566 **OPERANDS**

16567 The following operand shall be supported:

16568 file The pathname of a text file to be used as input.

16569 **STDIN**

See the INPUT FILES section.

16571 INPUT FILES

16572 Input files shall be text files.

16573 ENVIRONMENT VARIABLES

16574 The following environment variables shall affect the execution of *expand*:

16575 LANG Provide a default value for the internationalization variables that are unset or null.
16576 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
16577 Internationalization Variables for the precedence of internationalization variables
16578 used to determine the values of locale categories.)

16579 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

16581 *LC_CTYPE* 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

Utilities expand

16583 arguments and input files), the processing of <tab>s and <space>s, and for the 16584 determination of the width in column positions each character would occupy on an output device. 16585

LC MESSAGES 16586

Determine the locale that should be used to affect the format and contents of 16587 diagnostic messages written to standard error. 16588

NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 16589 XSI

16590 ASYNCHRONOUS EVENTS

16591 Default.

16592 STDOUT

The standard output shall be equivalent to the input files with <tab>s converted into the 16593 appropriate number of <space>s. 16594

16595 STDERR

The standard error shall be used only for diagnostic messages. 16596

16597 OUTPUT FILES

None. 16598

16599 EXTENDED DESCRIPTION

None. 16600

16601 EXIT STATUS

16602 The following exit values shall be returned:

Successful completion 16603

16604 >0 An error occurred.

16605 CONSEQUENCES OF ERRORS

The expand utility shall terminate with an error message and non-zero exit status upon 16606 encountering difficulties accessing one of the *file* operands. 16607

16608 APPLICATION USAGE

None. 16609

16610 EXAMPLES

None. 16611

16612 RATIONALE

The expand utility is useful for preprocessing text files (before sorting, looking at specific 16613 16614 columns, and so on) that contain <tab>s.

See the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.103, Column Position. 16615

The tablist option-argument consists of integers in ascending order. Utility Syntax Guideline 8 16616 mandates that expand shall accept the integers (within the single argument) separated using 16617

either commas or <black>s. 16618

16619 FUTURE DIRECTIONS

None. 16620

16621 SEE ALSO

16622 tabs, unexpand **expand** Utilities

16623 CHANGE HISTORY First released in Issue 4. 16624 16625 **Issue 6** 16626 This utility is marked as part of the User Portability Utilities option. The APPLICATION USAGE section is added. 16627 The obsolescent SYNOPSIS is removed. 16628 The LC_CTYPE environment variable description is updated to align with the IEEE P1003.2b 16629 draft standard. 16630 The normative text is reworded to avoid use of the term "must" for application requirements. 16631

16632 **NAME** 16633 expr — evaluate arguments as an expression 16634 SYNOPSIS 16635 expr operand 16636 DESCRIPTION The *expr* utility shall evaluate an expression and write the result to standard output. 16637 16638 OPTIONS None. 16639 16640 OPERANDS The single expression evaluated by expr shall be formed from the operands, as described in the 16641 EXTENDED DESCRIPTION section. The application shall ensure that each of the expression 16642 16643 operator symbols: 16644 () & >= < <= ! = + and the symbols *integer* and *string* in the table are provided as separate arguments to *expr*. 16645 16646 STDIN Not used. 16647 16648 INPUT FILES None. 16649 16650 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *expr*: 16651 LANG 16652 Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 16653 Internationalization Variables for the precedence of internationalization variables 16654 used to determine the values of locale categories.) 16655 LC_ALL If set to a non-empty string value, override the values of all the other 16656 internationalization variables. 16657 16658 LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multi-16659 character collating elements within regular expressions and by the string 16660 comparison operators. 16661 Determine the locale for the interpretation of sequences of bytes of text data as 16662 LC_CTYPE 16663 characters (for example, single-byte as opposed to multi-byte characters in arguments) and the behavior of character classes within regular expressions. 16664 LC_MESSAGES 16665 Determine the locale that should be used to affect the format and contents of 16666 diagnostic messages written to standard error. 16667 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 16668 XSI 16669 ASYNCHRONOUS EVENTS Default. 16670 16671 STDOUT The expr utility shall evaluate the expression and write the result, followed by a <newline>, to 16672

standard output.

16673

Utilities expr

16674 STDERR

16679

16680

16681

16682 16683

16684

16675 The standard error shall be used only for diagnostic messages.

16676 OUTPUT FILES

16677 None.

16678 EXTENDED DESCRIPTION

The formation of the expression to be evaluated is shown in the following table. The symbols expr, expr1, and expr2 represent expressions formed from integer and string symbols and the expression operator symbols (all separate arguments) by recursive application of the constructs described in the table. The expressions are listed in order of increasing precedence, with equalprecedence operators grouped between horizontal lines. All of the operators shall be leftassociative.

16685	Expression	Description
16686	expr1 expr2	Returns the evaluation of <i>expr1</i> if it is neither null nor zero;
16687		otherwise, returns the evaluation of expr2 if it is not null;
16688		otherwise, zero.
16689	expr1 & expr2	Returns the evaluation of <i>expr1</i> if neither expression evaluates to
16690		null or zero; otherwise, returns zero.
16691		Returns the result of a decimal integer comparison if both
16692		arguments are integers; otherwise, returns the result of a string
16693		comparison using the locale-specific collation sequence. The
16694		result of each comparison is 1 if the specified relationship is true,
16695		or 0 if the relationship is false.
16696	expr1 = expr2	Equal.
16697	expr1 > expr2	Greater than.
16698	<i>expr1</i> >= <i>expr2</i>	Greater than or equal.
16699	expr1 < expr2	Less than.
16700	expr1 <= expr2	Less than or equal.
16701	expr1 != expr2	Not equal.
16702	expr1 + expr2	Addition of decimal integer-valued arguments.
16703	expr1 – expr2	Subtraction of decimal integer-valued arguments.
16704	expr1 * expr2	Multiplication of decimal integer-valued arguments.
16705	expr1 / expr2	Integer division of decimal integer-valued arguments, producing
16706		an integer result.
16707	expr1 % expr2	Remainder of integer division of decimal integer-valued
16708		arguments.
16709	expr1 : expr2	Matching expression; see below.
16710	(expr)	Grouping symbols. Any expression can be placed within
16711		parentheses. Parentheses can be nested to a depth of
16712		{EXPR_NEST_MAX}.
16713	integer	An argument consisting only of an (optional) unary minus
16714		followed by digits.
16715	string	A string argument; see below.

Utilities expr

16716 **Matching Expression**

16717 The ':' matching operator shall compare the string resulting from the evaluation of *expr1* with 16718 the regular expression pattern resulting from the evaluation of expr2. Regular expression syntax shall be that defined in the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.3, Basic 16719 16720 Regular Expressions, except that all patterns are anchored to the beginning of the string (that is, 16721 only sequences starting at the first character of a string are matched by the regular expression) and, therefore, it is unspecified whether ' ^ ' is a special character in that context. Usually, the 16722 16723 matching operator shall return a string representing the number of characters matched ('0' on failure). Alternatively, if the pattern contains at least one regular expression subexpression 16724 " [\setminus (... \setminus)] ", the string corresponding to " \setminus 1" shall be returned. 16725

String Operand

A string argument is an argument that cannot be identified as an *integer* argument or as one of the expression operator symbols shown in the OPERANDS section.

The use of string arguments **length**, **substr**, **index**, or **match** produces unspecified results.

16730 EXIT STATUS

16726

16729

16733

16731 The following exit values shall be returned:

16732 0 The *expression* evaluates to neither null nor zero.

1 The *expression* evaluates to null or zero.

16734 2 Invalid expression.

16735 >2 An error occurred.

16736 CONSEQUENCES OF ERRORS

16737 Default.

16738 APPLICATION USAGE

After argument processing by the shell, *expr* is not required to be able to tell the difference between an operator and an operand except by the value. If "a" is '=', the command:

```
16741 expr $a = ' = '
```

16742 looks like:

16743 expr = = = =

as the arguments are passed to expr (and they all may be taken as the '=' operator). The following works reliably:

```
16746 expr X$a = X=
```

Also note that this volume of IEEE Std 1003.1-2001 permits implementations to extend utilities.

The *expr* utility permits the integer arguments to be preceded with a unary minus. This means that an integer argument could look like an option. Therefore, the conforming application must employ the "--" construct of Guideline 10 of the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines to protect its operands if there is any chance the first operand might be a negative integer (or any string with a leading minus).

16753 EXAMPLES

16754

16755 16756 The *expr* utility has a rather difficult syntax:

 Many of the operators are also shell control operators or reserved words, so they have to be escaped on the command line. **expr** Utilities

• Each part of the expression is composed of separate arguments, so liberal usage of blank>s
is required. For example:

Invalid	Valid
expr 1+2	<i>expr</i> 1 + 2
expr "1 + 2"	<i>expr</i> 1 + 2
expr 1 + (2 * 3)	expr 1 + \(2 * 3 \)

In many cases, the arithmetic and string features provided as part of the shell command language are easier to use than their equivalents in *expr*. Newly written scripts should avoid *expr* in favor of the new features within the shell; see Section 2.5 (on page 33) and Section 2.6.4 (on page 41).

16768 The following command:

```
a=\$(expr \$a + 1)
```

16764

16765

1676616767

16777

16778

16779

16780

adds 1 to the variable *a*.

The following command, for "\$a" equal to either /usr/abc/file or just file:

```
16772 expr a : '.*/(.*)' \mid a
```

returns the last segment of a pathname (that is, **file**). Applications should avoid the character // used alone as an argument; *expr* may interpret it as the division operator.

16775 The following command:

```
16776 \exp r ''/\$a'' : '.*/\(.*\)'
```

is a better representation of the previous example. The addition of the "//" characters eliminates any ambiguity about the division operator and simplifies the whole expression. Also note that pathnames may contain characters contained in the *IFS* variable and should be quoted to avoid having "\$a" expand into multiple arguments.

16781 The following command:

```
16782 expr "$VAR" : '.*'
```

returns the number of characters in *VAR*.

16784 RATIONALE

In an early proposal, EREs were used in the matching expression syntax. This was changed to BREs to avoid breaking historical applications.

The use of a leading circumflex in the BRE is unspecified because many historical implementations have treated it as a special character, despite their system documentation. For example:

```
16790 expr foo : ^foo expr ^foo : ^foo
```

return 3 and 0, respectively, on those systems; their documentation would imply the reverse.
Thus, the anchoring condition is left unspecified to avoid breaking historical scripts relying on this undocumented feature.

16794 FUTURE DIRECTIONS

16795 None.

Utilities **expr**

16796 SEE ALSO

16797 Section 2.5 (on page 33), Section 2.6.4 (on page 41)

16798 CHANGE HISTORY

16799 First released in Issue 2.

16800 **Issue 5**

16801 The FUTURE DIRECTIONS section is added.

16802 **Issue 6**

The *expr* utility is aligned with the IEEE P1003.2b draft standard, to include resolution of IEEE

PASC Interpretation 1003.2 #104.

The normative text is reworded to avoid use of the term "must" for application requirements.

false Utilities

16806 **NAME** 16807 false — return false value 16808 SYNOPSIS 16809 false 16810 **DESCRIPTION** 16811 The false utility shall return with a non-zero exit code. 16812 **OPTIONS** 16813 None. 16814 **OPERANDS** 16815 None. 16816 **STDIN** 16817 Not used. 16818 INPUT FILES 16819 None. 16820 ENVIRONMENT VARIABLES 16821 None. 16822 ASYNCHRONOUS EVENTS 16823 Default. 16824 STDOUT 16825 Not used. 16826 STDERR Not used. 16827 16828 OUTPUT FILES 16829 None. 16830 EXTENDED DESCRIPTION None. 16831 16832 EXIT STATUS The false utility shall always exit with a value other than zero. 16834 CONSEQUENCES OF ERRORS 16835 Default. 16836 APPLICATION USAGE 16837 None. 16838 EXAMPLES 16839 None. 16840 RATIONALE 16841 None. 16842 FUTURE DIRECTIONS None. 16843 16844 SEE ALSO 16845 true

false **Utilities**

16846 CHANGE HISTORY

First released in Issue 2. 16847

16848 **Issue 6**

IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/24 is applied, changing the STDERR section from "None." to "Not used." for alignment with Section 1.11 (on page 20). 16849

16850

fc Utilities

NAME

16852 fc — process the command history list

fc -s[old=new][first]

16853 SYNOPSIS

```
      16854 UP
      fc [-r] [-e editor] [first[last]]

      16855
      fc -l [-nr] [first[last]]
```

16858 DESCRIPTION

The *fc* utility shall list, or shall edit and re-execute, commands previously entered to an interactive *sh*.

The command history list shall reference commands by number. The first number in the list is selected arbitrarily. The relationship of a number to its command shall not change except when the user logs in and no other process is accessing the list, at which time the system may reset the numbering to start the oldest retained command at another number (usually 1). When the number reaches an implementation-defined upper limit, which shall be no smaller than the value in *HISTSIZE* or 32 767 (whichever is greater), the shell may wrap the numbers, starting the next command with a lower number (usually 1). However, despite this optional wrapping of numbers, *fc* shall maintain the time-ordering sequence of the commands. For example, if four commands in sequence are given the numbers 32 766, 32 767, 1 (wrapped), and 2 as they are executed, command 32 767 is considered the command previous to 1, even though its number is higher.

When commands are edited (when the **–l** option is not specified), the resulting lines shall be entered at the end of the history list and then re-executed by *sh*. The *fc* command that caused the editing shall not be entered into the history list. If the editor returns a non-zero exit status, this shall suppress the entry into the history list and the command re-execution. Any command line variable assignments or redirection operators used with *fc* shall affect both the *fc* command itself as well as the command that results; for example:

16878 fc -s -- -1 2 > /dev/null

reinvokes the previous command, suppressing standard error for both fc and the previous command.

16881 OPTIONS

The *fc* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

16884 The following options shall be supported:

16885	− e editor	Use the editor named by <i>editor</i> to edit the commands. The <i>editor</i> string is a utility
16886		name, subject to search via the PATH variable (see the Base Definitions volume of
16887		IEEE Std 1003.1-2001, Chapter 8, Environment Variables). The value in the FCEDIT
16888		variable shall be used as a default when -e is not specified. If FCEDIT is null or
16889		unset, ed shall be used as the editor.
16890	- l	(The letter ell.) List the commands rather than invoking an editor on them. The
16891		commands shall be written in the sequence indicated by the first and last operands,
16892		as affected by $-\mathbf{r}$, with each command preceded by the command number.
16893	−n	Suppress command numbers when listing with -l.
16894	- r	Reverse the order of the commands listed (with -l) or edited (with neither -l nor
16895		−s).

Utilities fc

Re-execute the command without invoking an editor.

16896

 $-\mathbf{s}$

10000	3	ne execute	the command without invoking an eartor.
16897 16898	OPERANDS The followir	ng operands s	hall be supported:
16899 16900 16901	first, last	accessed sh	ommands to list or edit. The number of previous commands that can be all be determined by the value of the <i>HISTSIZE</i> variable. The value of both shall be one of the following:
16902 16903		[+]number	A positive number representing a command number; command numbers can be displayed with the $-\mathbf{l}$ option.
16904 16905 16906		-number	A negative decimal number representing the command that was executed <i>number</i> of commands previously. For example, -1 is the immediately previous command.
16907 16908 16909 16910		string	A string indicating the most recently entered command that begins with that string. If the <i>old=new</i> operand is not also specified with –s, the string form of the <i>first</i> operand cannot contain an embedded equal sign.
16911		When the sy	ynopsis form with − s is used:
16912		• If <i>first</i> is	omitted, the previous command shall be used.
16913		For the sync	opsis forms without –s:
16914 16915			s omitted, $last$ shall default to the previous command when $-\mathbf{l}$ is \mathbf{l} ; otherwise, it shall default to $first$.
16916 16917			nd <i>last</i> are both omitted, the previous 16 commands shall be listed or ious single command shall be edited (based on the –l option).
16918 16919 16920 16921 16922 16923 16924		edited (accompl comman the comm - r . For e	and <i>last</i> are both present, all of the commands from <i>first</i> to <i>last</i> shall be without –I) or listed (with –I). Editing multiple commands shall be ished by presenting to the editor all of the commands at one time, each ad starting on a new line. If <i>first</i> represents a newer command than <i>last</i> , mands shall be listed or edited in reverse sequence, equivalent to using example, the following commands on the first line are equivalent to the ording commands on the second:
16925 16926		fc -r :	10 20 fc 30 40 20 10 fc -r 40 30
16927 16928 16929 16930		values the oldes	range of commands is used, it shall not be an error to specify <i>first</i> or <i>last</i> hat are not in the history list; <i>fc</i> shall substitute the value representing st or newest command in the list, as appropriate. For example, if there ten commands in the history list, numbered 1 to 10:
16931 16932		fc -l fc 1 9	9
16933		shall list	and edit, respectively, all ten commands.
16934 16935	old=new	Replace the string <i>new</i> .	first occurrence of string old in the commands to be re-executed by the

fc Utilities

16936 **STDIN**

16941 16942

16943

16944

16945

16946

16947

16948

16949

16950

16951

16952

16953

16954

16955

16956

16957

16958

16959

16960

16961

16962

16963 16964

16965

16966

16967

16968

16969

16970 16971

16972

16973

16974

16975

16976

16977

16978

16979

16980

16981

16982

16983

16937 Not used.

16938 INPUT FILES

16939 None.

16940 ENVIRONMENT VARIABLES

HISTFILE

HISTSIZE

LANG

The following environment variables shall affect the execution of *fc*:

FCEDIT This variable, when expanded by the shell, shall determine the default value for the -e editor option's editor option-argument. If FCEDIT is null or unset, ed shall be used as the editor.

Determine a pathname naming a command history file. If the HISTFILE variable is not set, the shell may attempt to access or create a file .sh_history in the directory referred to by the HOME environment variable. If the shell cannot obtain both read and write access to, or create, the history file, it shall use an unspecified mechanism that allows the history to operate properly. (References to history "file" in this section shall be understood to mean this unspecified mechanism in such cases.) An implementation may choose to access this variable only when initializing the history file; this initialization shall occur when fc or sh first attempt to retrieve entries from, or add entries to, the file, as the result of commands issued by the user, the file named by the ENV variable, or implementation-defined system start-up files. In some historical shells, the history file is initialized just after the ENV file has been processed. Therefore, it is implementation-defined whether changes made to HISTFILE after the history file has been initialized are effective. Implementations may choose to disable the history list mechanism for users with appropriate privileges who do not set HISTFILE; the specific circumstances under which this occurs are implementation-defined. If more than one instance of the shell is using the same history file, it is unspecified how updates to the history file from those shells interact. As entries are deleted from the history file, they shall be deleted oldest first. It is unspecified when history file entries are physically removed from the history file.

Determine a decimal number representing the limit to the number of previous commands that are accessible. If this variable is unset, an unspecified default greater than or equal to 128 shall be used. The maximum number of commands in the history list is unspecified, but shall be at least 128. An implementation may choose to access this variable only when initializing the history file, as described under *HISTFILE*. Therefore, it is unspecified whether changes made to *HISTSIZE* after the history file has been initialized are effective.

Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)

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

LC_CTYPE 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).

LC_MESSAGES

Determine the locale that should be used to affect the format and contents of

diagnostic messages written to standard error.

Utilities fc

16984 XSI **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 16985 ASYNCHRONOUS EVENTS 16986 Default. 16987 STDOUT When the -l option is used to list commands, the format of each command in the list shall be as 16988 16989 16990 "%d\t%s\n", <line number>, <command> If both the $-\mathbf{l}$ and $-\mathbf{n}$ options are specified, the format of each command shall be: 16991 "\t%s\n", <command> 16992 If the *<command>* consists of more than one line, the lines after the first shall be displayed as: 16993 "\t%s\n", <continued-command> 16994 16995 STDERR The standard error shall be used only for diagnostic messages. 16996 16997 OUTPUT FILES None. 16998 16999 EXTENDED DESCRIPTION None. 17000 17001 EXIT STATUS 17002 The following exit values shall be returned: 17003 0 Successful completion of the listing. >0 An error occurred. 17004 Otherwise, the exit status shall be that of the commands executed by fc. 17005 17006 CONSEQUENCES OF ERRORS 17007 Default. 17008 APPLICATION USAGE Since editors sometimes use file descriptors as integral parts of their editing, redirecting their file 17009 17010 descriptors as part of the fc command can produce unexpected results. For example, if vi is the 17011 *FCEDIT* editor, the command: 17012 fc -s | more does not work correctly on many systems. 17013 17014 Users on windowing systems may want to have separate history files for each window by setting *HISTFILE* as follows: 17015 HISTFILE=\$HOME/.sh hist\$\$ 17016 17017 EXAMPLES None. 17018 17019 RATIONALE This utility is based on the *fc* built-in of the KornShell. 17020 An early proposal specified the -e option as $[-e \ editor \ [old = new \]]$, which is not historical 17021 practice. Historical practice in fc of either [-e editor] or [-e - [old= new]] is acceptable, but not 17022 both together. To clarify this, a new option -s was introduced replacing the [-e-]. This resolves 17023 the conflict and makes *fc* conform to the Utility Syntax Guidelines. 17024

fc Utilities

17025 17026 17027 17028	HISTFILE Some implementations of the KornShell check for the superuser and do not create a history file unless HISTFILE is set. This is done primarily to avoid creating unlinked files in the root file system when logging in during single-user mode. HISTFILE must be set for the superuser to have history.
17029 17030 17031 17032	HISTSIZE Needed to limit the size of history files. It is the intent of the standard developers that when two shells share the same history file, commands that are entered in one shell shall be accessible by the other shell. Because of the difficulties of synchronization over a network, the exact nature of the interaction is unspecified.
17033 17034 17035 17036 17037 17038 17039	The initialization process for the history file can be dependent on the system start-up files, in that they may contain commands that effectively preempt the settings the user has for <i>HISTFILE</i> and <i>HISTSIZE</i> . For example, function definition commands are recorded in the history file. If the system administrator includes function definitions in some system start-up file called before the <i>ENV</i> file, the history file is initialized before the user can influence its characteristics. In some historical shells, the history file is initialized just after the <i>ENV</i> file has been processed. Because of these situations, the text requires the initialization process to be implementation-defined.
17040 17041	Consideration was given to omitting the fc utility in favor of the command line editing feature in sh . For example, in vi editing mode, typing " <esc> \forall" is equivalent to:</esc>
17042	EDITOR=vi fc
17043 17044	However, the $\it fc$ utility allows the user the flexibility to edit multiple commands simultaneously (such as $\it fc$ 10 20) and to use editors other than those supported by $\it sh$ for command line editing.
17045 17046 17047 17048 17049 17050	In the KornShell, the alias \mathbf{r} ("re-do") is preset to \mathbf{fc} – \mathbf{e} – (equivalent to the POSIX \mathbf{fc} – \mathbf{s}). This is probably an easier command name to remember than \mathbf{fc} ("fix command"), but it does not meet the Utility Syntax Guidelines. Renaming \mathbf{fc} to \mathbf{hist} or \mathbf{redo} was considered, but since this description closely matches historical KornShell practice already, such a renaming was seen as gratuitous. Users are free to create aliases whenever odd historical names such as \mathbf{fc} , \mathbf{awk} , \mathbf{cat} , \mathbf{grep} , or \mathbf{yacc} are standardized by POSIX.
17051 17052 17053 17054 17055	Command numbers have no ordering effects; they are like serial numbers. The -r option and <i>-number</i> operand address the sequence of command execution, regardless of serial numbers. So, for example, if the command number wrapped back to 1 at some arbitrary point, there would be no ambiguity associated with traversing the wrap point. For example, if the command history were:
17056 17057 17058	32766: echo 1 32767: echo 2 1: echo 3
17059 17060	the number -2 refers to command 32767 because it is the second previous command, regardless of serial number.
17061 FUTUR 17062	RE DIRECTIONS None.
17063 SEE AI 17064	SO sh
17065 CHAN 17066	GE HISTORY First released in Issue 4.

Utilities fc

17067 **Issue 5**

17068 The FUTURE DIRECTIONS section is added.

17069 Issue 6

17070 This utility is marked as part of the User Portability Utilities option.

17071 In the ENVIRONMENT VARIABLES section, the text "user's home directory" is updated to

"directory referred to by the *HOME* environment variable".

fg Utilities

17073 **NAME**

17074 fg — run jobs in the foreground

17075 SYNOPSIS

17076 UP fg [job id]

17077

17078 DESCRIPTION

If job control is enabled (see the description of $set - \mathbf{m}$), the fg utility shall move a background job from the current environment (see Section 2.12 (on page 61)) into the foreground.

Using *fg* to place a job into the foreground shall remove its process ID from the list of those "known in the current shell execution environment"; see Section 2.9.3.1 (on page 50).

17083 OPTIONS

17084 None.

17085 **OPERANDS**

17086 The following operand shall be supported:

job_id Specify the job to be run as a foreground job. If no job_id operand is given, the job_id for the job that was most recently suspended, placed in the background, or run as a background job shall be used. The format of job_id is described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.203, Job Control Job ID.

17091 **STDIN**

17092 Not used.

17093 INPUT FILES

17094 None.

17095 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of fg:

17097 LANG Provide a default value for the internationalization variables that are unset or null.
17098 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
17099 Internationalization Variables for the precedence of internationalization variables
17100 used to determine the values of locale categories.)

17101 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

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).

17106 LC_MESSAGES

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

17109 XSI NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

17110 ASYNCHRONOUS EVENTS

17111 Default.

17112 STDOUT

The fg utility shall write the command line of the job to standard output in the following format:

17114 "%s\n", <command>

Utilities fg

17115 STDERR

17116 The standard error shall be used only for diagnostic messages.

17117 OUTPUT FILES

17118 None.

17119 EXTENDED DESCRIPTION

17120 None.

17121 EXIT STATUS

17122 The following exit values shall be returned:

17123 0 Successful completion.

17124 >0 An error occurred.

17125 CONSEQUENCES OF ERRORS

If job control is disabled, the *fg* utility shall exit with an error and no job shall be placed in the foreground.

17128 APPLICATION USAGE

The *fg* utility does not work as expected when it is operating in its own utility execution environment because that environment has no applicable jobs to manipulate. See the APPLICATION USAGE section for *bg*. For this reason, *fg* is generally implemented as a shell regular built-in.

17133 EXAMPLES

17134 None.

17135 RATIONALE

The extensions to the shell specified in this volume of IEEE Std 1003.1-2001 have mostly been based on features provided by the KornShell. The job control features provided by bg, fg, and jobs are also based on the KornShell. The standard developers examined the characteristics of the C shell versions of these utilities and found that differences exist. Despite widespread use of the C shell, the KornShell versions were selected for this volume of IEEE Std 1003.1-2001 to maintain a degree of uniformity with the rest of the KornShell features selected (such as the very popular command line editing features).

17143 FUTURE DIRECTIONS

17144 None.

17145 SEE ALSO

17146 Section 2.9.3.1 (on page 50), Section 2.12 (on page 61), bg, kill, jobs, wait

17147 CHANGE HISTORY

17148 First released in Issue 4.

17149 Issue 6

17150 This utility is marked as part of the User Portability Utilities option.

17151 The APPLICATION USAGE section is added.

The JC marking is removed from the SYNOPSIS since job control is mandatory is this issue.

file Utilities

17153 **NAME** file — determine file type 17154 17155 SYNOPSIS file [-dh] [-M file] [-m file] file ... 17156 UP 17157 file -i [-h] file ... 17158 17159 **DESCRIPTION** The file utility shall perform a series of tests in sequence on each specified file in an attempt to 17160 classify it: 17161 17162 1. If *file* does not exist, cannot be read, or its file status could not be determined, the output shall indicate that the file was processed, but that its type could not be determined. 17163 2. If the file is not a regular file, its file type shall be identified. The file types directory, FIFO, 17164 socket, block special, and character special shall be identified as such. Other 17165 implementation-defined file types may also be identified. If file is a symbolic link, by 17166 default the link shall be resolved and file shall test the type of file referenced by the 17167 17168 symbolic link. (See the $-\mathbf{h}$ and $-\mathbf{i}$ options below.) 3. If the length of file is zero, it shall be identified as an empty file. 17169 The *file* utility shall examine an initial segment of *file* and shall make a guess at identifying 17170 its contents based on position-sensitive tests. (The answer is not guaranteed to be correct; 17171 17172 see the $-\mathbf{d}$, $-\mathbf{M}$, and $-\mathbf{m}$ options below.) The file utility shall examine file and make a guess at identifying its contents based on 17173 17174 context-sensitive default system tests. (The answer is not guaranteed to be correct.) The file shall be identified as a data file. 17175 If file does not exist, cannot be read, or its file status could not be determined, the output shall 17176 indicate that the file was processed, but that its type could not be determined. 17177 If file is a symbolic link, by default the link shall be resolved and file shall test the type of file 17178 17179 referenced by the symbolic link. 17180 OPTIONS The file utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, 17181 17182 Utility Syntax Guidelines, except that the order of the -m, -d, and -M options shall be significant. 17183 17184 The following options shall be supported by the implementation: $-\mathbf{d}$ Apply any position-sensitive default system tests and context-sensitive default 17185 system tests to the file. This is the default if no **–M** or **–m** option is specified. 17186 -h When a symbolic link is encountered, identify the file as a symbolic link. If $-\mathbf{h}$ is 17187 not specified and file is a symbolic link that refers to a nonexistent file, file shall 17188 identify the file as a symbolic link, as if **-h** had been specified. 17189 17190 $-\mathbf{i}$ If a file is a regular file, do not attempt to classify the type of the file further, but 17191 identify the file as specified in the STDOUT section. -M file Specify the name of a file containing position-sensitive tests that shall be applied to 17192

17193

17194

17195

a file in order to classify it (see the EXTENDED DESCRIPTION). No positionsensitive default system tests nor context-sensitive default system tests shall be

applied unless the $-\mathbf{d}$ option is also specified.

file **Utilities**

17196 17197	- m file	Specify the name of a file containing position-sensitive tests that shall be applied to a file in order to classify it (see the EXTENDED DESCRIPTION).	
17198 17199 17200 17201 17202 17203	If the -m option is specified without specifying the -d option or the -M option, position-sensitive default system tests shall be applied after the position-sensitive tests specified by the -m option. If the -M option is specified with the -d option, the -m option, or both, or the -m option is specified with the -d option, the concatenation of the position-sensitive tests specified by these options shall be applied in the order specified by the appearance of these options. If a -M or -m file option-argument is - , the results are unspecified.		
17204 OPERA		or an around shall be supposed a	
17205	file	ng operand shall be supported: A pathname of a file to be tested.	
17206 17207 STDIN	me	A patimame of a file to be tested.	
17207 STDIN 17208	Not used.		
17209 INPUT			
17210	The <i>file</i> can l	pe any file type.	
17211 ENVIR 17212	ONMENT VA The followin	ARIABLES ng environment variables shall affect the execution of <i>file</i> :	
17213 17214 17215 17216	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)	
17217 17218	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.	
17219 17220 17221	LC_CTYPE	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).	
17222	LC_MESSA		
17223 17224 17225		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.	
17226 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.	
17227 ASYNC 17228	CHRONOUS I Default.	EVENTS	
17229 STDOU 17230		K locale, the following format shall be used to identify each operand, <i>file</i> specified:	
17231	"%s: %s\n	", <file>, <type></type></file>	
17232 17233 17234 17235	The values for < <i>type</i> > are unspecified, except that in the POSIX locale, if <i>file</i> is identified as one of the types listed in the following table, < <i>type</i> > shall contain (but is not limited to) the corresponding string, unless the file is identified by a position-sensitive test specified by a – M or – m option. Each space shown in the strings shall be exactly one <space>.</space>		

file Utilities

1	7236	
1	7997	

Table 4-8 File Utility Output Strings

17237	
17238	

17257

17258

17259

17261

 $17265 \\ 17266$

17268

17269 17270

If file is:	<pre><type> shall contain the string:</type></pre>	Notes
Nonexistent	cannot open	
Block special	block special	1
Character special	character special	1
Directory	directory	1
FIFO	fifo	1
Socket	socket	1
Symbolic link	symbolic link to	1
Regular file	regular file	1,2
Empty regular file	empty	3
Regular file that cannot be read	cannot open	3
Executable binary	executable	4,6
ar archive library (see ar)	archive	4,6
Extended <i>cpio</i> format (see <i>pax</i>)	cpio archive	4,6
Extended <i>tar</i> format (see ustar in <i>pax</i>)	tar archive	4,6
Shell script	commands text	5,6
C-language source	c program text	5,6
FORTRAN source	fortran program text	5,6
Regular file whose type cannot be determined	data	

Notes:

- 1. This is a file type test.
 - 2. This test is applied only if the -i option is specified.
- 17260 3. This test is applied only if the -i option is not specified.
 - 4. This is a position-sensitive default system test.
- 17262 5. This is a context-sensitive default system test.
- 6. Position-sensitive default system tests and context-sensitive default system tests are not applied if the –**M** option is specified unless the –**d** option is also specified.

In the POSIX locale, if *file* is identified as a symbolic link (see the **-h** option), the following alternative output format shall be used:

17267 "%s: %s \n ", <file>, <type>, <contents of link>"

If the file named by the *file* operand does not exist, cannot be read, or the type of the file named by the *file* operand cannot be determined, this shall not be considered an error that affects the exit status.

17271 STDERR

The standard error shall be used only for diagnostic messages.

17273 OUTPUT FILES

17274 None.

17275 EXTENDED DESCRIPTION

A file specified as an option-argument to the **-m** or **-M** options shall contain one position-sensitive test per line, which shall be applied to the file. If the test succeeds, the message field of the line shall be printed and no further tests shall be applied, with the exception that tests on immediately following lines beginning with a single ' > ' character shall be applied.

Utilities file

17280 Each line shall be composed of the following four
 separated fields: offset An unsigned number (optionally preceded by a single '>' character) specifying 17281 17282 the offset, in bytes, of the value in the file that is to be compared against the value field of the line. If the file is shorter than the specified offset, the test shall fail. 17283 If the *offset* begins with the character '>', the test contained in the line shall not be 17284 applied to the file unless the test on the last line for which the *offset* did not begin 17285 with a '>' was successful. By default, the offset shall be interpreted as an unsigned 17286 decimal number. With a leading 0x or 0X, the offset shall be interpreted as a 17287 hexadecimal number; otherwise, with a leading 0, the offset shall be interpreted as 17288 17289 an octal number. The type of the value in the file to be tested. The type shall consist of the type 17290 type specification characters c, d, f, s, and u, specifying character, signed decimal, 17291 floating point, string, and unsigned decimal, respectively. 17292 The type string shall be interpreted as the bytes from the file starting at the 17293 specified *offset* and including the same number of bytes specified by the *value* field. 17294 If insufficient bytes remain in the file past the *offset* to match the *value* field, the test 17295 shall fail. 17296 The type specification characters d, f, and u can be followed by an optional 17297 unsigned decimal integer that specifies the number of bytes represented by the 17298 type. The type specification character f can be followed by an optional F, D, or L, 17299 17300 indicating that the value is of type float, double, or long double, respectively. The type specification characters d and u can be followed by an optional C, S, I, or L, 17301 indicating that the value is of type **char**, **short**, **int**, or **long**, respectively. 17302 The default number of bytes represented by the type specifiers d, f, and u shall 17303 correspond to their respective C-language types as follows. If the system claims 17304 conformance to the C-Language Development Utilities option, those specifiers 17305 shall correspond to the default sizes used in the c99 utility. Otherwise, the default 17306 17307 sizes shall be implementation-defined. 17308 For the type specifier characters d and u, the default number of bytes shall correspond to the size of a basic integer type of the implementation. For these 17309 specifier characters, the implementation shall support values of the optional 17310 number of bytes to be converted corresponding to the number of bytes in the C-17311 language types **char**, **short**, **int**, or **long**. These numbers can also be specified by an 17312 application as the characters C, S, I, and L, respectively. The byte order used when 17313 17314 interpreting numeric values is implementation-defined, but shall correspond to the order in which a constant of the corresponding type is stored in memory on the 17315 system. 17316 For the type specifier f, the default number of bytes shall correspond to the 17317 number of bytes in the basic double precision floating-point data type of the 17318 underlying implementation. The implementation shall support values of the 17319 optional number of bytes to be converted corresponding to the number of bytes in 17320 17321 the C-language types float, double, and long double. These numbers can also be specified by an application as the characters F, D, and L, respectively. 17322 All type specifiers, except for s, can be followed by a mask specifier of the form 17323 &number. The mask value shall be AND'ed with the value of the input file before 17324 the comparison with the value field of the line is made. By default, the mask shall 17325 be interpreted as an unsigned decimal number. With a leading 0x or 0X, the mask 17326

shall be interpreted as an unsigned hexadecimal number; otherwise, with a leading

17327

file Utilities

17328		0, the mask shall	be interpreted as an unsigned octal number.
17329 17330			e, short , long , and string shall also be supported as type fields, d as dC, dS, dL, and s, respectively.
17331	value	The value to be co	ompared with the value from the file.
17332 17333 17334		Otherwise, inter	om the type field is s or string , then interpret the value as a string. The pret it as a number. If the value is a string, then the test shall then a string value exactly matches the bytes from the file.
17335		If the <i>value</i> is a st	ring, it can contain the following sequences:
17336 17337 17338 17339 17340 17341		∖character	The backslash-escape sequences as specified in the Base Definitions volume of IEEE Std 1003.1-2001, Table 5-1, Escape Sequences and Associated Actions ('\\', '\a', '\b', '\f', '\n', '\r', '\t', '\v'). The results of using any other character, other than an octal digit, following the backslash are unspecified.
17342 17343 17344 17345 17346 17347		∖octal	Octal sequences that can be used to represent characters with specific coded values. An octal sequence shall consist of a backslash followed by the longest sequence of one, two, or three octal-digit characters (01234567). If the size of a byte on the system is greater than 9 bits, the valid escape sequence used to represent a byte is implementation-defined.
17348 17349 17350 17351		number. Any su unsigned hexade	value that is not a string shall be interpreted as a signed decimal uch value, with a leading 0x or 0X, shall be interpreted as an ecimal number; otherwise, with a leading zero, the value shall be a unsigned octal number.
17352 17353 17354			not a string, it can be preceded by a character indicating the performed. Permissible characters and the comparisons they lows:
17355		= The test sha	ll succeed if the value from the file equals the <i>value</i> field.
17356		< The test sha	ll succeed if the value from the file is less than the <i>value</i> field.
17357		> The test sha	ll succeed if the value from the file is greater than the <i>value</i> field.
17358 17359			ll succeed if all of the set bits in the <i>value</i> field are set in the value
17360 17361		^ The test sha the value from	ll succeed if at least one of the set bits in the <i>value</i> field is not set in om the file.
17362 17363			ll succeed if the file is large enough to contain a value of the type arting at the offset specified.
17364 17365 17366 17367	message	using the notati field was a strin	be printed if the test succeeds. The <i>message</i> shall be interpreted on for the <i>printf</i> formatting specification; see <i>printf</i> . If the <i>value</i> g, then the value from the file shall be the argument for the <i>printf</i> fication; otherwise, the value from the file shall be the argument.
17368 EXIT S			
17369	The followin	g exit values shall	l be returned:

17369 The following exit values shall be returned:

17370 0 Successful completion.

Utilities file

17371 >0 An error occurred.

17372 CONSEQUENCES OF ERRORS

Default.

17374 APPLICATION USAGE

The *file* utility can only be required to guess at many of the file types because only exhaustive testing can determine some types with certainty. For example, binary data on some implementations might match the initial segment of an executable or a *tar* archive.

Note that the table indicates that the output contains the stated string. Systems may add text before or after the string. For executables, as an example, the machine architecture and various facts about how the file was link-edited may be included. Note also that on systems that recognize shell script files starting with "#!" as executable files, these may be identified as executable binary files rather than as shell scripts.

17383 EXAMPLES

Determine whether an argument is a binary executable file:

```
17385 file "$1" | grep -Fq executable &&
17386 printf "%s is executable.\n" "$1"
```

17387 RATIONALE

The **–f** option was omitted because the same effect can (and should) be obtained using the *xargs* utility.

Historical versions of the *file* utility attempt to identify the following types of files: symbolic link, directory, character special, block special, socket, *tar* archive, *cpio* archive, SCCS archive, archive library, empty, *compress* output, *pack* output, binary data, C source, FORTRAN source, assembler source, *nroff/troff/eqn/tbl* source *troff* output, shell script, C shell script, English text, ASCII text, various executables, APL workspace, compiled terminfo entries, and CURSES screen images. Only those types that are reasonably well specified in POSIX or are directly related to POSIX utilities are listed in the table.

Historical systems have used a "magic file" named /etc/magic to help identify file types. Because it is generally useful for users and scripts to be able to identify special file types, the -m flag and a portable format for user-created magic files has been specified. No requirement is made that an implementation of *file* use this method of identifying files, only that users be permitted to add their own classifying tests.

In addition, three options have been added to historical practice. The $-\mathbf{d}$ flag has been added to permit users to cause their tests to follow any default system tests. The $-\mathbf{i}$ flag has been added to permit users to test portably for regular files in shell scripts. The $-\mathbf{M}$ flag has been added to permit users to ignore any default system tests.

The IEEE Std 1003.1-2001 description of default system tests and the interaction between the <code>-d</code>, <code>-M</code>, and <code>-m</code> options did not clearly indicate that there were two types of "default system tests". The "position-sensitive tests" determine file types by looking for certain string or binary values at specific offsets in the file being examined. These position-sensitive tests were implemented in historical systems using the magic file described above. Some of these tests are now built into the *file* utility itself on some implementations so the output can provide more detail than can be provided by magic files. For example, a magic file can easily identify a **core** file on most implementations, but cannot name the program file that dropped the core. A magic file could produce output such as:

17415 /home/dwc/core: ELF 32-bit MSB core file SPARC Version 1

file Utilities

but by building the test into the *file* utility, you could get output such as:

/home/dwc/core: ELF 32-bit MSB core file SPARC Version 1, from 'testprog'

These extended built-in tests are still to be treated as position-sensitive default system tests even if they are not listed in /etc/magic or any other magic file.

The context-sensitive default system tests were always built into the *file* utility. These tests looked for language constructs in text files trying to identify shell scripts, C, FORTRAN, and other computer language source files, and even plain text files. With the addition of the $-\mathbf{m}$ and $-\mathbf{M}$ options the distinction between position-sensitive and context-sensitive default system tests became important because the order of testing is important. The context-sensitive system default tests should never be applied before any position-sensitive tests even if the $-\mathbf{d}$ option is specified before a $-\mathbf{m}$ option or $-\mathbf{M}$ option due to the high probability that the context-sensitive system default tests will incorrectly identify arbitrary text files as text files before position-sensitive tests specified by the $-\mathbf{m}$ or $-\mathbf{M}$ option would be applied to give a more accurate identification.

Leaving the meaning of $-\mathbf{M}$ – and $-\mathbf{m}$ – unspecified allows an existing prototype of these options to continue to work in a backwards-compatible manner. (In that implementation, $-\mathbf{M}$ – was roughly equivalent to $-\mathbf{d}$ in IEEE Std 1003.1-2001.)

The historical –c option was omitted as not particularly useful to users or portable shell scripts. In addition, a reasonable implementation of the *file* utility would report any errors found each time the magic file is read.

The historical format of the magic file was the same as that specified by the Rationale in the ISO POSIX-2:1993 standard for the *offset*, *value*, and *message* fields; however, it used less precise type fields than the format specified by the current normative text. The new type field values are a superset of the historical ones.

The following is an example magic file:

```
17441
            0
               short
                           070707
                                                 cpio archive
17442
            0
               short
                           0143561
                                                 Byte-swapped cpio archive
17443
            0
               string
                           070707
                                                 ASCII cpio archive
17444
            0
               long
                           0177555
                                                 Very old archive
17445
            0
               short
                           0177545
                                                 Old archive
                                                 Old packed data
17446
            0
               short
                           017437
            0
               string
                           \037\036
                                                 Packed data
17447
            0
               string
                          \377\037
                                                 Compacted data
17448
               string
                           \037\235
17449
                                                 Compressed data
            >2 byte&0x80 >0
                                                 Block compressed
17450
17451
            >2 byte&0x1f x
                                                 %d bits
            0
               string
                           \032\001
17452
                                                 Compiled Terminfo Entry
17453
            0
               short
                           0433
                                                 Curses screen image
17454
            0
               short
                           0434
                                                 Curses screen image
                                                 System V Release 1 archive
17455
            0
               string
                           <ar>
17456
            0
               string
                           !<arch>\n .SYMDEF
                                                 Archive random library
            0
                           !<arch>
                                                 Archive
17457
               string
            0
               string
                          ARF BEGARF
                                                 PHIGS clear text archive
17458
            0
                           0x137A2950
                                                 Scalable OpenFont binary
17459
               long
                                                 Encrypted scalable OpenFont binary
17460
               long
                           0x137A2951
```

The use of a basic integer data type is intended to allow the implementation to choose a word size commonly used by applications on that architecture.

17461

17462

17418 17419

17420

17421 17422

17423

17424

17425

17426

17427 17428

17429

17430

17431

17432

17433

17434

17435

17436 17437

17438

17439 17440 **Utilities** file

17463 FUTURE DIRECTIONS 17464 None. 17465 SEE ALSO 17466 ar, ls, pax 17467 CHANGE HISTORY 17468 First released in Issue 4. 17469 Issue 6 17470 This utility is marked as part of the User Portability Utilities option. 17471 Options and an EXTENDED DESCRIPTION are added as specified in the IEEE P1003.2b draft 17472 standard. IEEE PASC Interpretations 1003.2 #192 and #178 are applied. 17473 IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/25 is applied, making major changes to 17474 17475 address ambiguities raised in defect reports. IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/26 is applied, making it clear in the 17476 17477 OPTIONS section that the -m, -d, and -M options do not comply with Guideline 11 of the Utility Syntax Guidelines. 17478

find Utilities

17479 N	JAME			
17480	find — find	d files		
17481 S	YNOPSIS			
17482		-L] path [operand_expression]		
17483 L	DESCRIPTION			
17484	The <i>find</i> ut	ility shall recursively descend the directory hierarchy from each file specified by path,		
17485	_	evaluating a Boolean expression composed of the primaries described in the OPERANDS section		
17486	for each file	e encountered.		
17487		tility shall be able to descend to arbitrary depths in a file hierarchy and shall not fail		
17488		th length limitations (unless a <i>path</i> operand specified by the application exceeds		
17489		AX} requirements).		
17490		ility shall detect infinite loops; that is, entering a previously visited directory that is an		
17491 17492		of the last file encountered. When it detects an infinite loop, <i>find</i> shall write a message to standard error and shall either recover its position in the hierarchy or		
17493	terminate.	message to standard error and shan erarer recover his position in the incrurenty of		
17494 (PTIONS			
17495		tility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section		
17496		y Syntax Guidelines.		
17497	The follow	ing options shall be supported by the implementation:		
17498	- H	Cause the file information and file type evaluated for each symbolic link		
17499		encountered on the command line to be those of the file referenced by the link, and		
17500 17501		not the link itself. If the referenced file does not exist, the file information and type shall be for the link itself. File information for all symbolic links not on the		
17502		command line shall be that of the link itself.		
17503 17504	–L	Cause the file information and file type evaluated for each symbolic link to be those of the file referenced by the link, and not the link itself.		
17505 17506		more than one of the mutually-exclusive options –H and –L shall not be considered he last option specified shall determine the behavior of the utility.		
17507 C	PERANDS			
17508	The follow	ing operands shall be supported:		
17509	The <i>path</i> op	perand is a pathname of a starting point in the directory hierarchy.		
17510		rgument that starts with a $'-'$, or is a $'$! ' or a ' (', and all subsequent arguments		
17511		terpreted as an expression made up of the following primaries and operators. In the		
17512 17513		ns, wherever n is used as a primary argument, it shall be interpreted as a decimal ionally preceded by a plus $('+')$ or minus $('-')$ sign, as follows:		
17514	+n More			
17515	n Exactl	v n		
17516	−n Less th			
17517	The follow	ing primaries shall be supported:		
17518	-name pat	tern		
17519	1	The primary shall evaluate as true if the basename of the filename being examined		
17520		matches <i>pattern</i> using the pattern matching notation described in Section 2.13 (on		
17521		page 62).		

Utilities find

17522 17523 17524	-nouser	The primary shall evaluate as true if the file belongs to a user ID for which the <i>getpwuid()</i> function defined in the System Interfaces volume of IEEE Std 1003.1-2001 (or equivalent) returns NULL.
17525 17526 17527	-nogroup	The primary shall evaluate as true if the file belongs to a group ID for which the <code>getgrgid()</code> function defined in the System Interfaces volume of IEEE Std 1003.1-2001 (or equivalent) returns NULL.
17528 17529 17530 17531 17532	-xdev	The primary shall always evaluate as true; it shall cause <i>find</i> not to continue descending past directories that have a different device ID (<i>st_dev</i> , see the <i>stat(</i>) function defined in the System Interfaces volume of IEEE Std 1003.1-2001). If any -xdev primary is specified, it shall apply to the entire expression even if the -xdev primary would not normally be evaluated.
17533 17534 17535	-prune	The primary shall always evaluate as true; it shall cause <i>find</i> not to descend the current pathname if it is a directory. If the –depth primary is specified, the –prune primary shall have no effect.
17536 17537 17538 17539 17540 17541 17542 17543 17544	- perm [-] <i>ma</i>	The <i>mode</i> argument is used to represent file mode bits. It shall be identical in format to the <i>symbolic_mode</i> operand described in <i>chmod</i> , and shall be interpreted as follows. To start, a template shall be assumed with all file mode bits cleared. An <i>op</i> symbol of '+' shall set the appropriate mode bits in the template; '-' shall clear the appropriate bits; '=' shall set the appropriate mode bits, without regard to the contents of process' file mode creation mask. The <i>op</i> symbol of '-' cannot be the first character of <i>mode</i> ; this avoids ambiguity with the optional leading hyphen. Since the initial mode is all bits off, there are not any symbolic modes that need to use '-' as the first character.
17546 17547		If the hyphen is omitted, the primary shall evaluate as true when the file permission bits exactly match the value of the resulting template.
17548 17549		Otherwise, if <i>mode</i> is prefixed by a hyphen, the primary shall evaluate as true if at least all the bits in the resulting template are set in the file permission bits.
17550 17551 17552 17553 17554 17555 17556	-perm [-] <i>on</i>	If the hyphen is omitted, the primary shall evaluate as true when the file permission bits exactly match the value of the octal number <i>onum</i> and only the bits corresponding to the octal mask 07777 shall be compared. (See the description of the octal <i>mode</i> in <i>chmod</i> .) Otherwise, if <i>onum</i> is prefixed by a hyphen, the primary shall evaluate as true if at least all of the bits specified in <i>onum</i> that are also set in the octal mask 07777 are set.
17557 17558 17559	−type c	The primary shall evaluate as true if the type of the file is c , where c is 'b', 'c', 'd', 'l', 'p', 'f', or 's' for block special file, character special file, directory, symbolic link, FIFO, regular file, or socket, respectively.
17560	-links n	The primary shall evaluate as true if the file has n links.
17561 17562 17563	-user uname	The primary shall evaluate as true if the file belongs to the user <i>uname</i> . If <i>uname</i> is a decimal integer and the <i>getpwnam()</i> (or equivalent) function does not return a valid user name, <i>uname</i> shall be interpreted as a user ID.
17564 17565 17566 17567	- group gnan	The primary shall evaluate as true if the file belongs to the group <i>gname</i> . If <i>gname</i> is a decimal integer and the <i>getgrnam()</i> (or equivalent) function does not return a valid group name, <i>gname</i> shall be interpreted as a group ID.

find Utilities

The primary shall evaluate as true if the file size in bytes, divided by 512 and rounded up to the next integer, is n. If n is followed by the character 'c', the size shall be in bytes.

The primary shall evaluate as true if the file access time subtracted from the

-atime *n* The primary shall evaluate as true if the file access time subtracted from the initialization time, divided by 86 400 (with any remainder discarded), is *n*.

–ctime n The primary shall evaluate as true if the time of last change of file status information subtracted from the initialization time, divided by 86 400 (with any remainder discarded), is n.

-mtime *n* The primary shall evaluate as true if the file modification time subtracted from the initialization time, divided by 86 400 (with any remainder discarded), is *n*.

```
-exec utility_name [argument...];
-exec utility_name [argument...] {} +
```

The end of the primary expression shall be punctuated by a semicolon or by a plus sign. Only a plus sign that follows an argument containing the two characters "{}" shall punctuate the end of the primary expression. Other uses of the plus sign shall not be treated as special.

If the primary expression is punctuated by a semicolon, the utility <code>utility_name</code> shall be invoked once for each pathname and the primary shall evaluate as true if the utility returns a zero value as exit status. A <code>utility_name</code> or <code>argument</code> containing only the two characters " { } " shall be replaced by the current pathname.

If the primary expression is punctuated by a plus sign, the primary shall always evaluate as true, and the pathnames for which the primary is evaluated shall be aggregated into sets. The utility <code>utility_name</code> shall be invoked once for each set of aggregated pathnames. Each invocation shall begin after the last pathname in the set is aggregated, and shall be completed before the <code>find</code> utility exits and before the first pathname in the next set (if any) is aggregated for this primary, but it is otherwise unspecified whether the invocation occurs before, during, or after the evaluations of other primaries. If any invocation returns a non-zero value as exit status, the <code>find</code> utility shall return a non-zero exit status. An argument containing only the two characters "{}" shall be replaced by the set of aggregated pathnames, with each pathname passed as a separate argument to the invoked utility in the same order that it was aggregated. The size of any set of two or more pathnames shall be limited such that execution of the utility does not cause the system's {ARG_MAX} limit to be exceeded. If more than one argument containing only the two characters "{}" is present, the behavior is unspecified.

If a *utility_name* or *argument* string contains the two characters "{}", but not just the two characters "{}", it is implementation-defined whether *find* replaces those two characters or uses the string without change. The current directory for the invocation of *utility_name* shall be the same as the current directory when the *find* utility was started. If the *utility_name* names any of the special built-in utilities (see Section 2.14 (on page 64)), the results are undefined.

-**ok** utility_name [argument...];

The $-\mathbf{ok}$ primary shall be equivalent to $-\mathbf{exec}$, except that the use of a plus sign to punctuate the end of the primary expression need not be supported, and *find* shall request affirmation of the invocation of *utility_name* using the current file as an argument by writing to standard error as described in the STDERR section. If the response on standard input is affirmative, the utility shall be invoked. Otherwise, the command shall not be invoked and the value of the $-\mathbf{ok}$ operand shall be false.

Utilities find

17616 17617	–print	The primary shall always evaluate as true; it shall cause the current pathname to be written to standard output.	
17618 17619	-newer file	The primary shall evaluate as true if the modification time of the current file is more recent than the modification time of the file named by the pathname <i>file</i> .	
17620 17621 17622 17623 17624 17625	-depth	The primary shall always evaluate as true; it shall cause descent of the directory hierarchy to be done so that all entries in a directory are acted on before the directory itself. If a -depth primary is not specified, all entries in a directory shall be acted on after the directory itself. If any -depth primary is specified, it shall apply to the entire expression even if the -depth primary would not normally be evaluated.	
17626 17627		The primaries can be combined using the following operators (in order of decreasing precedence):	
17628	(expression)	True if <i>expression</i> is true.	
17629	! expression	Negation of a primary; the unary NOT operator.	
17630 17631 17632 17633	expression [–a	a] expression Conjunction of primaries; the AND operator is implied by the juxtaposition of two primaries or made explicit by the optional —a operator. The second expression shall not be evaluated if the first expression is false.	
17634 17635 17636	expression –o	expression Alternation of primaries; the OR operator. The second expression shall not be evaluated if the first expression is true.	
17637 17638 17639	If no <i>expression</i> is present, –print shall be used as the expression. Otherwise, if the given expression does not contain any of the primaries –exec , –ok , or –print , the given expression shall be effectively replaced by:		
17640	(given_e	(given expression) -print	
17641 17642	The –user , – once.	-group, and -newer primaries each shall evaluate their respective arguments only	
17643 STDIN			
17644 17645		If the $-\mathbf{ok}$ primary is used, the response shall be read from the standard input. An entire line shall be read as the response. Otherwise, the standard input shall not be used.	
17646 INPUT 17647	FILES None.		
17648 ENVIR 17649	ONMENT VARIABLES The following environment variables shall affect the execution of <i>find</i> :		
17650 17651 17652 17653	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)	
17654 17655	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.	
17656 17657 17658 17659	LC_COLLAT	Determine the locale for the behavior of ranges, equivalence classes, and multi-character collating elements used in the pattern matching notation for the –n option and in the extended regular expression defined for the yesexpr locale	

find **Utilities**

17660		keyword in the <i>LC_MESSAGES</i> category.		
17661 17662 17663 17664 17665	LC_CTYPE	This variable determines 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), the behavior of character classes within the pattern matching notation used for the $-\mathbf{n}$ option, and the behavior of character classes within regular expressions used in the extended regular expression defined for the yesexpr locale keyword in the $LC_MESSAGES$ category.		
17667	LC_MESSA			
17668 17669 17670		Determine the locale for the processing of affirmative responses that should be used to affect the format and contents of diagnostic messages written to standard error.		
17671 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .		
17672 17673 17674	PATH	Determine the location of the <i>utility_name</i> for the –exec and –ok primaries, as described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables.		
17675 ASYNCHRONOUS EVENTS 17676 Default.				
17677 STDOU	${f T}$			
17678 17679	The -print primary shall cause the current pathnames to be written to standard output. The format shall be:			
17680	"%s\n", <path></path>			
17681 STDER				
17682 17683 17684	The -ok primary shall write a prompt to standard error containing at least the <i>utility_name</i> to be invoked and the current pathname. In the POSIX locale, the last non- <blank> in the prompt shall be '?'. The exact format used is unspecified.</blank>			
17685	Otherwise, the standard error shall be used only for diagnostic messages.			
17686 OUTPUT FILES 17687 None.				
17688 EXTENDED DESCRIPTION 17689 None.				
17690 EXIT STATUS 17691 The following exit values shall be returned:				
17692	0 All path operands were traversed successfully.			
17693	>0 An erro	r occurred.		

17695

17694 CONSEQUENCES OF ERRORS Default.

Utilities find

17696 APPLICATION USAGE

When used in operands, pattern matching notation, semicolons, opening parentheses, and closing parentheses are special to the shell and must be quoted (see Section 2.2 (on page 30)).

The bit that is traditionally used for sticky (historically 01000) is specified in the **–perm** primary using the octal number argument form. Since this bit is not defined by this volume of IEEE Std 1003.1-2001, applications must not assume that it actually refers to the traditional sticky bit.

17703 EXAMPLES

17699

17700

17701

17702

17704

17714

17715

17716 17717

17718

17719

17720

17722

17723 17724

17725

17726

17727

17728

17729

1. The following commands are equivalent:

```
17705 find .
17706 find . -print
```

17707 They both write out the entire directory hierarchy from the current directory.

17708 2. The following command:

```
17709 find / \( -name tmp -o -name '*.xx' \) -atime +7 -exec rm \{\} \;
```

removes all files named **tmp** or ending in **.xx** that have not been accessed for seven or more 24-hour periods.

17712 3. The following command:

```
17713 find .-perm-o+w,+s
```

prints (**-print** is assumed) the names of all files in or below the current directory, with all of the file permission bits S_ISUID, S_ISGID, and S_IWOTH set.

4. The following command:

```
find . -name SCCS -prune -o -print
```

recursively prints pathnames of all files in the current directory and below, but skips directories named SCCS and files in them.

The following command:

```
17721 find . -print -name SCCS -prune
```

behaves as in the previous example, but prints the names of the SCCS directories.

6. The following command is roughly equivalent to the **-nt** extension to *test*:

```
if [ -n "$(find file1 -prune -newer file2)" ]; then
    printf %s\\n "file1 is newer than file2"
fi
```

7. The descriptions of **–atime**, **–ctime**, and **–mtime** use the terminology n "86 400 second periods (days)". For example, a file accessed at 23:59 is selected by:

```
find . -atime -1 -print
```

at 00:01 the next day (less than 24 hours later, not more than one day ago); the midnight boundary between days has no effect on the 24-hour calculation.

17732 RATIONALE

17733 The **-a** operator was retained as an optional operator for compatibility with historical shell scripts, even though it is redundant with expression concatenation.

find Utilities

The descriptions of the '-' modifier on the *mode* and *onum* arguments to the **-perm** primary agree with historical practice on BSD and System V implementations. System V and BSD documentation both describe it in terms of checking additional bits; in fact, it uses the same bits, but checks for having at least all of the matching bits set instead of having exactly the matching bits set.

The exact format of the interactive prompts is unspecified. Only the general nature of the contents of prompts are specified because:

- Implementations may desire more descriptive prompts than those used on historical implementations.
- Since the historical prompt strings do not terminate with <newline>s, there is no portable way for another program to interact with the prompts of this utility via pipes.

Therefore, an application using this prompting option relies on the system to provide the most suitable dialog directly with the user, based on the general guidelines specified.

The **–name** *file* operand was changed to use the shell pattern matching notation so that *find* is consistent with other utilities using pattern matching.

The -**size** operand refers to the size of a file, rather than the number of blocks it may occupy in the file system. The intent is that the st_size field defined in the System Interfaces volume of IEEE Std 1003.1-2001 should be used, not the st_blocks found in historical implementations. There are at least two reasons for this:

- 1. In both System V and BSD, *find* only uses *st_size* in size calculations for the operands specified by this volume of IEEE Std 1003.1-2001. (BSD uses *st_blocks* only when processing the **-ls** primary.)
- 2. Users usually think of file size in terms of bytes, which is also the unit used by the *ls* utility for the output from the –l option. (In both System V and BSD, *ls* uses *st_size* for the –l option size field and uses *st_blocks* for the *ls* –s calculations. This volume of IEEE Std 1003.1-2001 does not specify *ls* –s.)

The descriptions of -atime, -ctime, and -mtime were changed from the SVID description of n "days" to "24-hour periods". The description is also different in terms of the exact timeframe for the n case (versus the +n or -n), but it matches all known historical implementations. It refers to one 86 400 second period in the past, not any time from the beginning of that period to the current time. For example, -atime 3 is true if the file was accessed any time in the period from 72 hours to 48 hours ago.

Historical implementations do not modify "{}" when it appears as a substring of an **–exec** or **–ok** *utility_name* or argument string. There have been numerous user requests for this extension, so this volume of IEEE Std 1003.1-2001 allows the desired behavior. At least one recent implementation does support this feature, but encountered several problems in managing memory allocation and dealing with multiple occurrences of "{}" in a string while it was being developed, so it is not yet required behavior.

Assuming the presence of **–print** was added to correct a historical pitfall that plagues novice users, it is entirely upwards-compatible from the historical System V *find* utility. In its simplest form (*find directory*), it could be confused with the historical BSD fast *find*. The BSD developers agreed that adding **–print** as a default expression was the correct decision and have added the fast *find* functionality within a new utility called *locate*.

Historically, the **–L** option was implemented using the primary **–follow**. The **–H** and **–L** options were added for two reasons. First, they offer a finer granularity of control and consistency with other programs that walk file hierarchies. Second, the **–follow** primary always evaluated to true.

Utilities find

17781 As they were historically really global variables that took effect before the traversal began, some 17782 valid expressions had unexpected results. An example is the expression -print -o -follow. 17783 Because **-print** always evaluates to true, the standard order of evaluation implies that **-follow** would never be evaluated. This was never the case. Historical practice for the **-follow** primary, 17784 17785 however, is not consistent. Some implementations always follow symbolic links on the command line whether -follow is specified or not. Others follow symbolic links on the 17786 command line only if -follow is specified. Both behaviors are provided by the -H and -L 17787 options, but scripts using the current **-follow** primary would be broken if the **-follow** option is 17788 17789 specified to work either way.

Since the **–L** option resolves all symbolic links and the **–type** *l* primary is true for symbolic links that still exist after symbolic links have been resolved, the command:

17792 find -L . -type 1

17790 17791

17795

17796

17797

17798

17799

17800

17801

17807

17808 17809

prints a list of symbolic links reachable from the current directory that do not resolve to accessible files.

A feature of SVR4's *find* utility was the **–exec** primary's + terminator. This allowed filenames containing special characters (especially <newline>s) to be grouped together without the problems that occur if such filenames are piped to *xargs*. Other implementations have added other ways to get around this problem, notably a **–print0** primary that wrote filenames with a null byte terminator. This was considered here, but not adopted. Using a null terminator meant that any utility that was going to process *find*'s **–print0** output had to add a new option to parse the null terminators it would now be reading.

The "-exec . . . {} +" syntax adopted was a result of IEEE PASC Interpretation 1003.2 #210.

It should be noted that this is an incompatible change to the ISO/IEC 9899: 1999 standard. For example, the following command prints all files with a '-' after their name if they are regular files, and a '+' otherwise:

17806 find / -type f -exec echo {} - ';' -o -exec echo {} + ';'

The change invalidates usage like this. Even though the previous standard stated that this usage would work, in practice many did not support it and the standard developers felt it better to now state that this was not allowable.

17810 FUTURE DIRECTIONS

17811 None.

17812 SEE ALSO

Section 2.2 (on page 30), Section 2.13 (on page 62), Section 2.14 (on page 64), *chmod*, *pax*, *sh*, *test*, the System Interfaces volume of IEEE Std 1003.1-2001, *getgrgid*(), *getpwuid*(), *stat*()

17815 CHANGE HISTORY

17816 First released in Issue 2.

17817 **Issue 5**

17818 The FUTURE DIRECTIONS section is added.

17819 Issue 6

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

• The **-perm** [**-**] *onum* primary is supported.

The *find* utility is aligned with the IEEE P1003.2b draft standard, to include processing of symbolic links and changes to the description of the **atime**, **ctime**, and **mtime** operands.

find Utilities

17825

IEEE PASC Interpretation 1003.2 #210 is applied, extending the **-exec** operand.

fold **Utilities**

17826 **NAME** fold — filter for folding lines 17827 17828 SYNOPSIS 17829 fold [-bs] [-w width] [file...]

17830 **DESCRIPTION**

17831

17832

17833

17834 17835

17836 17837

17840

17841

17842

17851

The fold utility is a filter that shall fold lines from its input files, breaking the lines to have a maximum of width column positions (or bytes, if the -b option is specified). Lines shall be broken by the insertion of a <newline> such that each output line (referred to later in this section as a segment) is the maximum width possible that does not exceed the specified number of column positions (or bytes). A line shall not be broken in the middle of a character. The behavior is undefined if width is less than the number of columns any single character in the input would occupy.

If the <carriage-return>s, <backspace>s, or <tab>s are encountered in the input, and the -b 17838 option is not specified, they shall be treated specially: 17839

 <backspace> The current count of line width shall be decremented by one, although the count never shall become negative. The fold utility shall not insert a <newline> immediately before or after any <backspace>.

<carriage-return> 17843

The current count of line width shall be set to zero. The *fold* utility shall not insert a 17844 <newline> immediately before or after any <carriage-return>. 17845

<tab> Each <tab> encountered shall advance the column position pointer to the next tab 17846 17847 stop. Tab stops shall be at each column position *n* such that *n* modulo 8 equals 1.

17848 OPTIONS

The fold utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 17849 17850 12.2, Utility Syntax Guidelines.

The following options shall be supported:

-b Count *width* in bytes rather than column positions. 17852

If a segment of a line contains a <blank> within the first width column positions (or 17853 -s bytes), break the line after the last such <blank> meeting the width constraints. If 17854 there is no
blank> meeting the requirements, the -s option shall have no effect for 17855 that output segment of the input line. 17856

Specify the maximum line length, in column positions (or bytes if $-\mathbf{b}$ is specified). 17857 -w width 17858 The results are unspecified if width is not a positive decimal number. The default value shall be 80. 17859

17860 OPERANDS

The following operand shall be supported: 17861

file A pathname of a text file to be folded. If no *file* operands are specified, the standard 17862 input shall be used. 17863

17864 **STDIN**

The standard input shall be used only if no file operands are specified. See the INPUT FILES 17865 section. 17866

fold **Utilities**

17867 INPUT FILES

17868 If the -b option is specified, the input files shall be text files except that the lines are not limited to {LINE_MAX} bytes in length. If the -b option is not specified, the input files shall be text files. 17869

17870 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *fold*: 17871

17872 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 17873 Internationalization Variables for the precedence of internationalization variables 17874 used to determine the values of locale categories.) 17875

LC_ALL If set to a non-empty string value, override the values of all the other 17876 17877

internationalization variables.

LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 17878 characters (for example, single-byte as opposed to multi-byte characters in 17879 arguments and input files), and for the determination of the width in column 17880 positions each character would occupy on a constant-width font output device. 17881

LC MESSAGES 17882

Determine the locale that should be used to affect the format and contents of 17883 17884 diagnostic messages written to standard error.

NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 17885 XSI

17886 ASYNCHRONOUS EVENTS

Default. 17887

17888 STDOUT

The standard output shall be a file containing a sequence of characters whose order shall be 17889 preserved from the input files, possibly with inserted <newline>s. 17890

17891 STDERR

The standard error shall be used only for diagnostic messages. 17892

17893 OUTPUT FILES

None. 17894

17895 EXTENDED DESCRIPTION

None. 17896

17897 EXIT STATUS

17898 The following exit values shall be returned:

All input files were processed successfully. 17899

17900 >0 An error occurred.

17901 CONSEQUENCES OF ERRORS

17902 Default. Utilities fold

17903 APPLICATION USAGE

The *cut* and *fold* utilities can be used to create text files out of files with arbitrary line lengths. The *cut* utility should be used when the number of lines (or records) needs to remain constant. The *fold* utility should be used when the contents of long lines need to be kept contiguous.

The *fold* utility is frequently used to send text files to printers that truncate, rather than fold, lines wider than the printer is able to print (usually 80 or 132 column positions).

17909 EXAMPLES

An example invocation that submits a file of possibly long lines to the printer (under the assumption that the user knows the line width of the printer to be assigned by lp):

17912 fold -w 132 bigfile | lp

17913 RATIONALE

17914

17915 17916

17917

17918

17919

17920

17921

17922

17923

17924

17925

17926

17927 17928

17929

17930

17931

Although terminal input in canonical processing mode requires the erase character (frequently set to
backspace>) to erase the previous character (not byte or column position), terminal output is not buffered and is extremely difficult, if not impossible, to parse correctly; the interpretation depends entirely on the physical device that actually displays/prints/stores the output. In all known internationalized implementations, the utilities producing output for mixed column-width output assume that a
backspace> backs up one column position and outputs enough
backspace>s to return to the start of the character when
backspace> is used to provide local line motions to support underlining and emboldening operations. Since *fold* without the -b option is dealing with these same constraints,
backspace> is always treated as backing up one column position rather than backing up one character.

Historical versions of the *fold* utility assumed 1 byte was one character and occupied one column position when written out. This is no longer always true. Since the most common usage of *fold* is believed to be folding long lines for output to limited-length output devices, this capability was preserved as the default case. The $-\mathbf{b}$ option was added so that applications could *fold* files with arbitrary length lines into text files that could then be processed by the standard utilities. Note that although the width for the $-\mathbf{b}$ option is in bytes, a line is never split in the middle of a character. (It is unspecified what happens if a width is specified that is too small to hold a single character found in the input followed by a <newline>.)

The tab stops are hardcoded to be every eighth column to meet historical practice. No new method of specifying other tab stops was invented.

17934 FUTURE DIRECTIONS

17935 None.

17936 **SEE ALSO**

17937 *cut*

17938 CHANGE HISTORY

First released in Issue 4.

17940 Issue 6

The normative text is reworded to avoid use of the term "must" for application requirements.

fort77 Utilities

```
17942 NAME
```

17949

17950 17951

17952

17953

17954

17957

17958

17959

17960

17961

17962

17963

17964

17966

17967

17968 17969

17970

17943 fort77 — FORTRAN compiler (**FORTRAN**)

17944 SYNOPSIS

17945 FD fort77 [-c] [-g] [-L directory]... [-O optlevel] [-o outfile] [-s] [-w]
17946 operand...
17947

17948 **DESCRIPTION**

The *fort77* utility is the interface to the FORTRAN compilation system; it shall accept the full FORTRAN-77 language defined by the ANSI X3.9-1978 standard. The system conceptually consists of a compiler and link editor. The files referenced by *operands* are compiled and linked to produce an executable file. It is unspecified whether the linking occurs entirely within the operation of *fort77*; some implementations may produce objects that are not fully resolved until the file is executed.

17955 If the **-c** option is present, for all pathname operands of the form *file*.**f**, the files:

17956 \$ (basename pathname.f).o

shall be created or overwritten as the result of successful compilation. If the -c option is not specified, it is unspecified whether such .o files are created or deleted for the *file*.f operands.

If there are no options that prevent link editing (such as -c) and all operands compile and link without error, the resulting executable file shall be written into the file named by the -o option (if present) or to the file **a.out**. The executable file shall be created as specified in the System Interfaces volume of IEEE Std 1003.1-2001, except that the file permissions shall be set to:

S IRWXO | S IRWXG | S IRWXU

and that the bits specified by the *umask* of the process shall be cleared.

17965 **OPTIONS**

The *fort77* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines, except that:

- The *l library* operands have the format of options, but their position within a list of operands affects the order in which libraries are searched.
- The order of specifying the multiple –L options is significant.
- Conforming applications shall specify each option separately; that is, grouping option letters (for example, -cg) need not be recognized by all implementations.

17973 The following options shall be supported:

- 17974 c Suppress the link-edit phase of the compilation, and do not remove any object files that are produced.
- 17976 —g Produce symbolic information in the object or executable files; the nature of this information is unspecified, and may be modified by implementation-defined interactions with other options.
- 17979 s Produce object or executable files, or both, from which symbolic and other 17980 information not required for proper execution using the *exec* family of functions defined in the System Interfaces volume of IEEE Std 1003.1-2001 has been removed (stripped). If both g and s options are present, the action taken is unspecified.
- 17983 \mathbf{o} outfile Use the pathname outfile, instead of the default $\mathbf{a.out}$, for the executable file produced. If the $-\mathbf{o}$ option is present with $-\mathbf{c}$, the result is unspecified.

Utilities fort77

17985	–L directory	Change the algorithm of searching for the libraries named in –l operands to look in
17986		the directory named by the <i>directory</i> pathname before looking in the usual places.
17987		Directories named in -L options shall be searched in the specified order. At least
17988		ten instances of this option shall be supported in a single fort77 command
17989		invocation. If a directory specified by a -L option contains a file named libf.a, the
17990		results are unspecified.
17991	-O optlevel	Specify the level of code optimization. If the <i>optlevel</i> option-argument is the digit
17992	_	'0', all special code optimizations shall be disabled. If it is the digit '1', the
17993		nature of the optimization is unspecified. If the –O option is omitted, the nature of
17994		the system's default optimization is unspecified. It is unspecified whether code
17995		generated in the presence of the –O 0 option is the same as that generated when
17996		−O is omitted. Other <i>optlevel</i> values may be supported.
17997	$-\mathbf{w}$	Suppress warnings.
17998	Multiple instances of –L options can be specified.	

17999 OPERANDS

18004

18005

18006

18007

18008 18009

18011

18013

18014

18015

18016

An operand is either in the form of a pathname or the form -1 library. At least one operand of the 18000 pathname form shall be specified. The following operands shall be supported: 18001

file.f The pathname of a FORTRAN source file to be compiled and optionally passed to 18002 the link editor. The filename operand shall be of this form if the -c option is used. 18003

> file.a A library of object files typically produced by ar, and passed directly to the link editor. Implementations may recognize implementation-defined suffixes other than .a as denoting object file libraries.

> > An object file produced by fort77 -c and passed directly to the link editor. Implementations may recognize implementation-defined suffixes other than .o as denoting object files.

The processing of other files is implementation-defined. 18010

> -l library (The letter ell.) Search the library named:

18012 liblibrary.a

file.o

A library is searched when its name is encountered, so the placement of a -l operand is significant. Several standard libraries can be specified in this manner, as described in the EXTENDED DESCRIPTION section. Implementations may recognize implementation-defined suffixes other than .a as denoting libraries.

18017 **STDIN**

Not used. 18018

18019 INPUT FILES

The input file shall be one of the following: a text file containing FORTRAN source code; an 18020 18021 object file in the format produced by fort77 -c; or a library of object files, in the format produced by archiving zero or more object files, using ar. Implementations may supply additional utilities 18022 that produce files in these formats. Additional input files are implementation-defined. 18023

A <tab> encountered within the first six characters on a line of source code shall cause the 18024 compiler to interpret the following character as if it were the seventh character on the line (that 18025 18026 is, in column 7).

fort77 **Utilities**

	18028	The following environment variables shall affect the execution of <i>fort77</i> :			
	18029 18030 18031 18032	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)		
	18033 18034	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.		
	18035 18036 18037	LC_CTYPE	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).		
	18038 18039 18040	LC_MESSAC	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.		
	18041 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .		
	18042 18043	TMPDIR	Determine the pathname that should override the default directory for temporary files, if any.		
18044 ASYNCHRONOUS EVENTS 18045 Default.					
18046 STDOUT 18047 Not used.					
	18048 STDERR 18049 The standard error shall be used only for diagnostic messages. If more than one <i>file</i> operand ending in .f (or possibly other unspecified suffixes) is given, for each such file:				

18052

18027 ENVIRONMENT VARIABLES

may be written to allow identification of the diagnostic message with the appropriate input file.

This utility may produce warning messages about certain conditions that do not warrant 18053 18054

returning an error (non-zero) exit value.

"%s:\n", <file>

18055 OUTPUT FILES

18051

Object files, listing files, and executable files shall be produced in unspecified formats. 18056

18057 EXTENDED DESCRIPTION

18058 **Standard Libraries** The *fort77* utility shall recognize the following –**l** operand for the standard library: 18059 -lfThis library contains all functions referenced in the ANSI X3.9-1978 standard. This 18060 operand shall not be required to be present to cause a search of this library. 18061 18062 In the absence of options that inhibit invocation of the link editor, such as -c, the fort77 utility shall cause the equivalent of a -1 f operand to be passed to the link editor as the last -1 operand, 18063 causing it to be searched after all other object files and libraries are loaded. 18064

It is unspecified whether the library libf.a exists as a regular file. The implementation may 18065 accept as -I operands names of objects that do not exist as regular files. 18066

Utilities fort77

18067 External Symbols

The FORTRAN compiler and link editor shall support the significance of external symbols up to a length of at least 31 bytes; case folding is permitted. The action taken upon encountering symbols exceeding the implementation-defined maximum symbol length is unspecified.

The compiler and link editor shall support a minimum of 511 external symbols per source or object file, and a minimum of 4095 external symbols total. A diagnostic message is written to standard output if the implementation-defined limit is exceeded; other actions are unspecified.

18074 EXIT STATUS

18071

18072

18073

18075 The following exit values shall be returned:

18076 0 Successful compilation or link edit.

18077 >0 An error occurred.

18078 CONSEQUENCES OF ERRORS

When *fort77* encounters a compilation error, it shall write a diagnostic to standard error and continue to compile other source code operands. It shall return a non-zero exit status, but it is implementation-defined whether an object module is created. If the link edit is unsuccessful, a diagnostic message shall be written to standard error, and *fort77* shall exit with a non-zero status.

18084 APPLICATION USAGE

18085 None.

18086 EXAMPLES

The following usage example compiles **xyz.f** and creates the executable file **foo**:

18088 fort77 -o foo xyz.f

The following example compiles **xyz.f** and creates the object file **xyz.o**:

18090 fort77 -c xyz.f

The following example compiles **xyz.f** and creates the executable file **a.out**:

18092 fort77 xyz.f

The following example compiles xyz.f, links it with b.o, and creates the executable a.out:

18094 fort77 xyz.f b.o

18095 RATIONALE

18096

18097

18098

18099

18100

18101

18102

18103

The name of this utility was chosen as *fort77* to parallel the renaming of the C compiler. The name *f77* was not chosen to avoid problems with historical implementations. The ANSI X3.9-1978 standard was selected as a normative reference because the ISO/IEC version of FORTRAN-77 has been superseded by the ISO/IEC 1539: 1990 standard (Fortran-90).

The file inclusion and symbol definition **#define** mechanisms used by the *c99* utility were not included in this volume of IEEE Std 1003.1-2001—even though they are commonly implemented—since there is no requirement that the FORTRAN compiler use the C preprocessor.

The **-onetrip** option was not included in this volume of IEEE Std 1003.1-2001, even though many historical compilers support it, because it is derived from FORTRAN-66; it is an anachronism that should not be perpetuated.

Some implementations produce compilation listings. This aspect of FORTRAN has been left unspecified because there was controversy concerning the various methods proposed for implementing it: a $-\mathbf{V}$ option overlapped with historical vendor practice and a naming

fort77 Utilities

convention of creating files with .l suffixes collided with historical *lex* file naming practice.

There is no –**I** option in this version of this volume of IEEE Std 1003.1-2001 to specify a directory for file inclusion. An INCLUDE directive has been a part of the Fortran-90 discussions, but an interface supporting that standard is not in the current scope.

It is noted that many FORTRAN compilers produce an object module even when compilation errors occur; during a subsequent compilation, the compiler may patch the object module rather than recompiling all the code. Consequently, it is left to the implementor whether or not an object file is created.

A reference to MIL-STD-1753 was removed from an early proposal in response to a request from the POSIX FORTRAN-binding standard developers. It was not the intention of the standard developers to require certification of the FORTRAN compiler, and IEEE Std 1003.9-1992 does not specify the military standard or any special preprocessing requirements. Furthermore, use of that document would have been inappropriate for an international standard.

The specification of optimization has been subject to changes through early proposals. At one time, $-\mathbf{O}$ and $-\mathbf{N}$ were Booleans: optimize and do not optimize (with an unspecified default). Some historical practice led this to be changed to:

 $-\mathbf{O}$ 0 No optimization.

- 18127 —O 1 Some level of optimization.
- $-\mathbf{O}$ *n* Other, unspecified levels of optimization.

It is not always clear whether "good code generation" is the same thing as optimization. Simple optimizations of local actions do not usually affect the semantics of a program. The $-\mathbf{O}$ 0 option has been included to accommodate the very particular nature of scientific calculations in a highly optimized environment; compilers make errors. Some degree of optimization is expected, even if it is not documented here, and the ability to shut it off completely could be important when porting an application. An implementation may treat $-\mathbf{O}$ 0 as "do less than normal" if it wishes, but this is only meaningful if any of the operations it performs can affect the semantics of a program. It is highly dependent on the implementation whether doing less than normal is logical. It is not the intent of the $-\mathbf{O}$ 0 option to ask for inefficient code generation, but rather to assure that any semantically visible optimization is suppressed.

The specification of standard library access is consistent with the C compiler specification. Implementations are not required to have /usr/lib/libf.a, as many historical implementations do, but if not they are required to recognize **f** as a token.

External symbol size limits are in normative text; conforming applications need to know these limits. However, the minimum maximum symbol length should be taken as a constraint on a conforming application, not on an implementation, and consequently the action taken for a symbol exceeding the limit is unspecified. The minimum size for the external symbol table was added for similar reasons.

The CONSEQUENCES OF ERRORS section clearly specifies the behavior of the compiler when compilation or link-edit errors occur. The behavior of several historical implementations was examined, and the choice was made to be silent on the status of the executable, or **a.out**, file in the face of compiler or linker errors. If a linker writes the executable file, then links it on disk with *lseek*()s and *write*()s, the partially linked executable file can be left on disk and its execute bits turned off if the link edit fails. However, if the linker links the image in memory before writing the file to disk, it need not touch the executable file (if it already exists) because the link edit fails. Since both approaches are historical practice, a conforming application shall rely on the exit status of *fort77*, rather than on the existence or mode of the executable file.

Utilities fort77

18156 18157 18158	The $-g$ and $-s$ options are not specified as mutually-exclusive. Historically these two options have been mutually-exclusive, but because both are so loosely specified, it seemed appropriate to leave their interaction unspecified.
18159 18160 18161 18162 18163	The requirement that conforming applications specify compiler options separately is to reserve the multi-character option name space for vendor-specific compiler options, which are known to exist in many historical implementations. Implementations are not required to recognize, for example, $-\mathbf{gc}$ as if it were $-\mathbf{g}$ $-\mathbf{c}$; nor are they forbidden from doing so. The SYNOPSIS shows all of the options separately to highlight this requirement on applications.
18164 18165 18166 18167	Echoing filenames to standard error is considered a diagnostic message because it would otherwise be difficult to associate an error message with the erring file. They are described with "may" to allow implementations to use other methods of identifying files and to parallel the description in <i>c99</i> .
18168 FUTUI 18169 18170	RE DIRECTIONS A compilation system based on the ISO/IEC 1539: 1990 standard (Fortran-90) may be considered for a future version; it may have a different utility name from <i>fort77</i> .
18171 SEE Al 18172	LSO ar, asa, c99, umask, the System Interfaces volume of IEEE Std 1003.1-2001, exec
18173 CHAN 18174	GE HISTORY First released in Issue 4.
18175 Issue 6 18176	This utility is marked as part of the FORTRAN Development Utilities option.

The normative text is reworded to avoid use of the term "must" for application requirements.

18177

fuser Utilities

	NAME	force list a	presses IDs of all pressesses that have one or more files oner			
18179	CVNIOD	•	process IDs of all processes that have one or more files open			
18180	SYNOP XSI		cfu] file			
18182			•			
18183	DESCR	IPTION				
18184			lity shall write to standard output the process IDs of processes running on the local			
18185 18186			have one or more named files open. For block special devices, all processes using nat device are listed.			
18187 18188			ility shall write to standard error additional information about the named files ow the file is being used.			
18189		Any output f	for processes running on remote systems that have a named file open is unspecified.			
18190		A user may r	need appropriate privilege to invoke the <i>fuser</i> utility.			
	OPTIO					
18192 18193			lity shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section Syntax Guidelines.			
18194		The followin	g options shall be supported:			
18195 18196		-с	The file is treated as a mount point and the utility shall report on any files open in the file system.			
18197		$-\mathbf{f}$	The report shall be only for the named files.			
18198 18199		-u	The user name, in parentheses, associated with each process ID written to standard output shall be written to standard error.			
18200	18200 OPERANDS					
18201		The followin	g operand shall be supported:			
18202		file	A pathname on which the file or file system is to be reported.			
	STDIN	NI. a 1				
18204		Not used.				
18205 18206	INPUT	FILES The user data	abase.			
18207	ENVIR	ONMENT VA	RIABLES			
18208		The followin	g environment variables shall affect the execution of <i>fuser</i> :			
18209		LANG	Provide a default value for the internationalization variables that are unset or null.			
18210 18211			(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables			
18212			used to determine the values of locale categories.)			
18213 18214		LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.			
18215		LC_CTYPE	Determine the locale for the interpretation of sequences of bytes of text data as			
18216			characters (for example, single-byte as opposed to multi-byte characters in			
18217			arguments).			
18218		LC_MESSAC	Determine the locale that should be used to affect the format and contents of			
18219 18220			diagnostic messages written to standard error.			

fuser **Utilities**

18221 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 18222 ASYNCHRONOUS EVENTS 18223 Default. 18224 STDOUT 18225 The fuser utility shall write the process ID for each process using each file given as an operand to standard output in the following format: 18226 18227 "%d", cess id> 18228 STDERR 18229 The *fuser* utility shall write diagnostic messages to standard error. The *fuser* utility also shall write the following to standard error: 18230 The pathname of each named file is written followed immediately by a colon. 18231 For each process ID written to standard output, the character 'c' shall be written to 18232 standard error if the process is using the file as its current directory and the character 'r' 18233 shall be written to standard error if the process is using the file as its root directory. 18234 18235 Implementations may write other alphabetic characters to indicate other uses of files. • When the -u option is specified, characters indicating the use of the file shall be followed 18236 immediately by the user name, in parentheses, corresponding to the process' real user ID. If 18237 the user name cannot be resolved from the process' real user ID, the process' real user ID 18238 shall be written instead of the user name. 18239 When standard output and standard error are directed to the same file, the output shall be 18240 interleaved so that the filename appears at the start of each line, followed by the process ID and 18241 characters indicating the use of the file. Then, if the $-\mathbf{u}$ option is specified, the user name or user 18242 ID for each process using that file shall be written. 18243 18244 A <newline> shall be written to standard error after the last output described above for each file 18245 operand. 18246 OUTPUT FILES 18247 None. 18248 EXTENDED DESCRIPTION 18249 None. 18250 EXIT STATUS The following exit values shall be returned: 18251 Successful completion. 18252 >0 An error occurred. 18253

18254 CONSEQUENCES OF ERRORS Default.

18255

fuser **Utilities**

18256 APPLICATION USAGE

18257 None.

18258 EXAMPLES

18259 The command:

fuser -fu . 18260

writes to standard output the process IDs of processes that are using the current directory and 18261 18262

writes to standard error an indication of how those processes are using the directory and the

18263 user names associated with the processes that are using the current directory.

18264 RATIONALE

The definition of the fuser utility follows existing practice. 18265

18266 FUTURE DIRECTIONS

18267 None.

18268 **SEE ALSO**

18269 None.

18270 CHANGE HISTORY

First released in Issue 5. 18271

Utilities gencat

18272 **NAME**

gencat — generate a formatted message catalog 18273

18274 SYNOPSIS

gencat catfile msgfile.. 18275 XSI

18276

18277 **DESCRIPTION**

The gencat utility shall merge the message text source file msgfile into a formatted message 18278 catalog catfile. The file catfile shall be created if it does not already exist. If catfile does exist, its 18279 messages shall be included in the new catfile. If set and message numbers collide, the new 18280 18281 message text defined in *msgfile* shall replace the old message text currently contained in *catfile*.

18282 OPTIONS

18283 None.

18284 OPERANDS

The following operands shall be supported: 18285

catfile A pathname of the formatted message catalog. If '-' is specified, standard output 18286

18287 shall be used. The format of the message catalog produced is unspecified.

msgfile A pathname of a message text source file. If '-' is specified for an instance of 18288

msgfile, standard input shall be used. The format of message text source files is

defined in the EXTENDED DESCRIPTION section.

18291 **STDIN**

18289

18290

18296

18305

The standard input shall not be used unless a *msgfile* operand is specified as '-'. 18292

18293 INPUT FILES

The input files shall be text files. 18294

18295 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *gencat*:

LANG 18297 Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 18298 Internationalization Variables for the precedence of internationalization variables 18299 used to determine the values of locale categories.) 18300

 LC_ALL If set to a non-empty string value, override the values of all the other 18301 internationalization variables. 18302

Determine the locale for the interpretation of sequences of bytes of text data as LC_CTYPE 18303 characters (for example, single-byte as opposed to multi-byte characters in 18304

arguments and input files).

LC_MESSAGES 18306

Determine the locale that should be used to affect the format and contents of 18307 diagnostic messages written to standard error. 18308

NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 18309

18310 ASYNCHRONOUS EVENTS

Default. 18311

18312 STDOUT

18313 The standard output shall not be used unless the *catfile* operand is specified as '-'. **gencat** Utilities

18314 STDERR

18315 The standard error shall be used only for diagnostic messages.

18316 OUTPUT FILES

18317 None.

18318 EXTENDED DESCRIPTION

The content of a message text file shall be in the format defined as follows. Note that the fields of a message text source line are separated by a single <blank>. Any other
 <blank>s are considered to be part of the subsequent field.

Sset n comment

This line specifies the set identifier of the following messages until the next **\$set** or end-of-file appears. The *n* denotes the set identifier, which is defined as a number in the range [1, {NL_SETMAX}] (see the **limits.h>** header defined in the Base Definitions volume of IEEE Std 1003.1-2001). The application shall ensure that set identifiers are presented in ascending order within a single source file, but need not be contiguous. Any string following the set identifier shall be treated as a comment. If no **\$set** directive is specified in a message text source file, all messages shall be located in an implementation-defined default message set NL_SETD (see the **<nl_types.h>** header defined in the Base Definitions volume of IEEE Std 1003.1-2001).

Sdelset n comment

This line deletes message set *n* from an existing message catalog. The *n* denotes the set number [1, {NL_SETMAX}]. Any string following the set number shall be treated as a comment.

\$ comment A line beginning with '\$' followed by a <blank> shall be treated as a comment.

m message-text

The *m* denotes the message identifier, which is defined as a number in the range [1, {NL_MSGMAX}] (see the <**limits.h**> header). The *message-text* shall be stored in the message catalog with the set identifier specified by the last **Sset** directive, and with message identifier *m*. If the *message-text* is empty, and a <blank> field separator is present, an empty string shall be stored in the message catalog. If a message source line has a message number, but neither a field separator nor *message-text*, the existing message with that number (if any) shall be deleted from the catalog. The application shall ensure that message identifiers are in ascending order within a single set, but need not be contiguous. The application shall ensure that the length of *message-text* is in the range [0, {NL_TEXTMAX}] (see the <**limits.h**> header).

Squote n

This line specifies an optional quote character *c*, which can be used to surround *message-text* so that trailing spaces or null (empty) messages are visible in a message source line. By default, or if an empty **Squote** directive is supplied, no quoting of *message-text* shall be recognized.

Empty lines in a message text source file shall be ignored. The effects of lines starting with any character other than those defined above are implementation-defined.

Text strings can contain the special characters and escape sequences defined in the following table:

Utilities gencat

18357 18358	Ī	Description	Symbol	Sequence	
18359		<newline></newline>	NL(LF)	\n	
18360		Horizontal-tab	HT	\t	
18361		<vertical-tab></vertical-tab>	VT	\v	
18362		<backspace></backspace>	BS	\b	
18363		<carriage-return></carriage-return>	CR	\r	
18364		<form-feed></form-feed>	FF	\f	
18365		Backslash	\	\\	
18366		Bit pattern	ddd	\ddd	
18367	The escape sequence "\de	dd" consists of bacl	kslash follo	wed by one,	
18368	which shall be taken to sp				
18369	backslash is not one of tho	se specified, the bac	kslash shall	be ignored.	
18370	Backslash ('∖') followed	by a <newline> is al</newline>	so used to	continue a s	
18371	Thus, the following two lin	nes describe a single	message st	ring:	
18372	1 This line continue	es \			
18373	to the next line				
18374	which shall be equivalent to:				
18375	1 This line continue	es to the next	line		
18376 EXIT	STATUS				
18377	The following exit values s	shall be returned:			
18378	0 Successful completion.				
18379	>0 An error occurred.				
18380 CONS	SEQUENCES OF ERRORS				
18381	Default.				
18382 APPL	ICATION USAGE				
18383	Message catalogs produce				
18384	be guaranteed between				
18385	recompiled for each type of	or machine, so messa	ge catalogs	must be rec	
18386 EXAN 18387	APLES None.				
18388 RATI					
18388 KATT	None.				
18390 FUIC	RE DIRECTIONS None.				
18392 SEE A	iconv, the Base Definitions	volume of IEEE Std	1003.1-2001	, <limits.h></limits.h>	
	NGE HISTORY				
18394 CHAI 18395	First released in Issue 3.				
10000	2 1150 1 010 050 0 111 10500 0.				

The normative text is reworded to avoid use of the term "must" for application requirements.

18396 **Issue 6**

18397

get Utilities

18398 NAME		
18399		version of an SCCS file (DEVELOPMENT)
18400 SYNOF		
18401 XSI 18402	get [-beg	kmnlLpst] [-c cutoff] [-i list] [-r SID] [-x list] file
18403 DESCR	PIPTION	
18404 18405	The <i>get</i> utilit given by its	ty shall generate a text file from each named SCCS <i>file</i> according to the specifications options.
18406 18407		ed text shall normally be written into a file called the g-file whose name is derived CS filename by simply removing the leading "s.".
18408 OPTIO	NS	
18409 18410	The <i>get</i> utilit	ty shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, ax Guidelines.
18411	The following	ng options shall be supported:
18412	-r SID	Indicate the SCCS Identification String (SID) of the version (delta) of an SCCS file
18413		to be retrieved. The table shows, for the most useful cases, what version of an
18414		SCCS file is retrieved (as well as the SID of the version to be eventually created by
18415		delta if the $-\mathbf{e}$ option is also used), as a function of the SID specified.
18416	− c cutoff	Indicate the <i>cutoff</i> date-time, in the form:
18417		YY[MM[DD[HH[MM[SS]]]]]
18418 18419		For the YY component, values in the range [69,99] shall refer to years 1969 to 1999 inclusive, and values in the range [00,68] shall refer to years 2000 to 2068 inclusive.
18420 18421		Note: It is expected that in a future version of IEEE Std 1003.1-2001 the default century inferred from a 2-digit year will change. (This would apply to all
18422		commands accepting a 2-digit year as input.)
18423		No changes (deltas) to the SCCS file that were created after the specified cutoff
18424		date-time shall be included in the generated text file. Units omitted from the date-
18425 18426		time default to their maximum possible values; for example, $-c$ 7502 is equivalent to $-c$ 750228235959.
18427		Any number of non-numeric characters may separate the various 2-digit pieces of
18428		the <i>cutoff</i> date-time. This feature allows the user to specify a <i>cutoff</i> date in the form:
18429		-c "77/2/2 9:22:25".
18430	−e	Indicate that the <i>get</i> is for the purpose of editing or making a change (delta) to the
18431		SCCS file via a subsequent use of <i>delta</i> . The –e option used in a <i>get</i> for a particular
18432		version (SID) of the SCCS file shall prevent further <i>get</i> commands from editing on
18433		the same SID until <i>delta</i> is executed or the \mathbf{j} (joint edit) flag is set in the SCCS file.
18434		Concurrent use of get –e for different SIDs is always allowed.
18435		If the g-file generated by <i>get</i> with a –e option is accidentally ruined in the process
18436 18437		of editing, it may be regenerated by re-executing the <i>get</i> command with the $-\mathbf{k}$ option in place of the $-\mathbf{e}$ option.
18438		SCCS file protection specified via the ceiling, floor, and authorized user list stored in the SCCS file shall be enforced when the a option is used
18439		in the SCCS file shall be enforced when the $-\mathbf{e}$ option is used.
18440	-b	Use with the –e option to indicate that the new delta should have an SID in a new
18441		branch as shown in the table below. This option shall be ignored if the b flag is not
18442		present in the file or if the retrieved delta is not a leaf delta. (A leaf delta is one that

Utilities get

18443		has no successors on the SCCS file tree.)
18444		Note: A branch delta may always be created from a non-leaf delta.
18445 18446	− i list	Indicate a <i>list</i> of deltas to be included (forced to be applied) in the creation of the generated file. The <i>list</i> has the following syntax:
18447 18448		<pre><list> ::= <range> <list> , <range> <range> ::= SID SID - SID</range></range></list></range></list></pre>
18449 18450 18451 18452 18453		SID, the SCCS Identification of a delta, may be in any form shown in the "SID Specified" column of the table in the EXTENDED DESCRIPTION section, except that the result of supplying a partial SID is unspecified. A diagnostic message shall be written if the first SID in the range is not an ancestor of the second SID in the range.
18454 18455	−x list	Indicate a <i>list</i> of deltas to be excluded (forced not to be applied) in the creation of the generated file. See the $-\mathbf{i}$ option for the <i>list</i> format.
18456 18457	-k	Suppress replacement of identification keywords (see below) in the retrieved text by their value. The $-\mathbf{k}$ option shall be implied by the $-\mathbf{e}$ option.
18458	-l	Write a delta summary into an l-file .
18459 18460 18461	–L	Write a delta summary to standard output. All informative output that normally is written to standard output shall be written to standard error instead, unless the $-s$ option is used, in which case it shall be suppressed.
18462 18463 18464 18465	- p	Write the text retrieved from the SCCS file to the standard output. No g-file shall be created. All informative output that normally goes to the standard output shall go to standard error instead, unless the $-\mathbf{s}$ option is used, in which case it shall disappear.
18466 18467 18468	−s	Suppress all informative output normally written to standard output. However, fatal error messages (which shall always be written to the standard error) shall remain unaffected.
18469 18470	- m	Precede each text line retrieved from the SCCS file by the SID of the delta that inserted the text line in the SCCS file. The format shall be:
18471		"%s\t%s", <sid>, <text line=""></text></sid>
18472 18473	-n	Precede each generated text line with the $\% M\%$ identification keyword value (see below). The format shall be:
18474		"%s\t%s", <%M% value>, <text line=""></text>
18475 18476		When both the $-\mathbf{m}$ and $-\mathbf{n}$ options are used, the $<$ text line $>$ shall be replaced by the $-\mathbf{m}$ option-generated format.
18477 18478	- g	Suppress the actual retrieval of text from the SCCS file. It is primarily used to generate an l-file , or to verify the existence of a particular SID.
18479 18480	-t	Use to access the most recently created (top) delta in a given release (for example, $-\mathbf{r} 1$), or release and level (for example, $-\mathbf{r} 1.2$).
18481 OPER 18482		ng operands shall be supported:
18483 18484 18485	file	A pathname of an existing SCCS file or a directory. If <i>file</i> is a directory, the <i>get</i> utility shall behave as though each file in the directory were specified as a named file, except that non-SCCS files (last component of the pathname does not begin

get Utilities

18486		with s.) and unreadable files shall be silently ignored.		
18487 18488 18489		If exactly one <i>file</i> operand appears, and it is $'-'$, the standard input shall be read; each line of the standard input is taken to be the name of an SCCS file to be processed. Non-SCCS files and unreadable files shall be silently ignored.		
18490 18491 18492		rd input shall be a text file used only if the <i>file</i> operand is specified as $'-'$. Each line ile shall be interpreted as an SCCS pathname.		
18493 18494	INPUT FILES The SCCS fi	iles shall be files of an unspecified format.		
18495	ENVIRONMENT V	ARIABLES		
18496	The followi	ng environment variables shall affect the execution of <i>get</i> :		
18497 18498 18499 18500	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)		
18501 18502	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.		
18503 18504 18505	LC_CTYPE	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).		
18506	LC_MESSA	GES		
18507 18508 18509		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error, and informative messages written to standard output (or standard error, if the $-\mathbf{p}$ option is used).		
18510	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .		
18511 18512 18513	TZ	Determine the timezone in which the times and dates written in the SCCS file are evaluated. If the TZ variable is unset or NULL, an unspecified system default timezone is used.		
	ASYNCHRONOUS	EVENTS		
18515	Default.			
18516 18517	STDOUT For each file	e processed, <i>get</i> shall write to standard output the SID being accessed and the number		
18518		of lines retrieved from the SCCS file, in the following format:		
18519	"%s\n%d]	ines\n", <sid>, <number lines="" of=""></number></sid>		
18520 18521		tion is used, the SID of the delta to be made shall appear after the SID accessed and number of lines generated, in the POSIX locale:		
18522 18523		<pre>delta %s\n%d lines\n", <sid accessed="">, to be made>, <number lines="" of=""></number></sid></pre>		

- shall be written before each of the lines shown in one of the preceding formats:

 "\n%s:\n", <pathname>
- If the **–L** option is used, a delta summary shall be written following the format specified below for **l-files**.

If there is more than one named file or if a directory or standard input is named, each pathname

18524

get **Utilities**

18529	If the – i opti	on is used, included deltas shall be listed following the notation, in the POSIX locale:
18530	"Included	:\n"
18531 18532	If the -x op locale:	tion is used, excluded deltas shall be listed following the notation, in the POSIX
18533	"Excluded	:\n"
18534 18535	If the -p or the SCCS file	-L options are specified, the standard output shall consist of the text retrieved from e.
18536 STDER		
18537 18538		d error shall be used only for diagnostic messages, except if the $-\mathbf{p}$ or $-\mathbf{L}$ options are shall include all informative messages normally sent to standard output.
18539 OUTPU 18540 18541 18542 18543 18544 18545 18546	Several auxi file, p-file, a be formed fi SCCS filena leading s w removing th	liary files may be created by <i>get</i> . These files are known generically as the g-file , l -md z-file . The letter before the hyphen is called the <i>tag</i> . An auxiliary filename shall rom the SCCS filename: the application shall ensure that the last component of all mes is of the form s. <i>module-name</i> ; the auxiliary files shall be named by replacing the ith the tag. The g-file shall be an exception to this scheme: the g-file is named by the s. prefix. For example, for s.xyz.c , the auxiliary filenames would be xyz.c , l.xyz.c , z.xyz.c , respectively.
18547 18548 18549 18550 18551	-p option is generated by g-file shall	which contains the generated text, shall be created in the current directory (unless the used). A g-file shall be created in all cases, whether or not any lines of text were by the <i>get</i> . It shall be owned by the real user. If the -k option is used or implied, the be writable by the owner only (read-only for everyone else); otherwise, it shall be only the real user need have write permission in the current directory.
18552 18553 18554 18555	text. The l-f	nall contain a table showing which deltas were applied in generating the retrieved ile shall be created in the current directory if the $-l$ option is used; it shall be readis owned by the real user. Only the real user need have write permission in the ctory.
18556	Lines in the	l-file shall have the following format:
18557 18558		$s\t%s\Delta%s\n"$, $<$ $code1>$, $<$ $code2>$, $<$ $code3>$, , $<$ $date-time>$, $<$ $login>$
18559	where the en	ntries are:
18560	<code1></code1>	A < space > if the delta was applied; $'*'$ otherwise.
18561 18562	< <i>code2</i> >	A <space> if the delta was applied or was not applied and ignored; $'*'$ if the delta was not applied and was not ignored.</space>
18563	<code3></code3>	A character indicating a special reason why the delta was or was not applied:
18564		I Included.
18565		X Excluded.
18566		C Cut off (by a $-c$ option).
18567 18568	<date-time></date-time>	Date and time (using the format of the <i>date</i> utility's $y/\mbox{m}/\mbox{d}\ \T$ conversion specification format) of creation.
18569	<login></login>	Login name of person who created <i>delta</i> .

get Utilities

18570 The comments and MR data shall follow on subsequent lines, indented one <tab>. A blank line 18571 shall terminate each entry. 18572 The **p-file** shall be used to pass information resulting from a get with a -e option along to delta. Its contents shall also be used to prevent a subsequent execution of get with a -e option for the 18573 same SID until delta is executed or the joint edit flag, j, is set in the SCCS file. The p-file shall be 18574 created in the directory containing the SCCS file and the application shall ensure that the 18575 effective user has write permission in that directory. It shall be writable by owner only, and 18576 owned by the effective user. Each line in the **p-file** shall have the following format: 18577 "% $s\Delta$ % $s\Delta$ % $s\Delta$ %s%s%s%s%n", <g-file SID>, 18578 18579 <SID of new delta>, <login-name of real user>, <date-time>, <i-value>, <x-value> 18580 where *<i-value*> uses the format " " if no –i option was specified, and shall use the format: 18581 " Δ -i%s", <-i option option-argument> 18582 if a -i option was specified and <x-value> uses the format " " if no -x option was specified, and 18583 shall use the format: 18584 " Δ -x%s", <-x option option-argument> 18585 if a -x option was specified. There can be an arbitrary number of lines in the **p-file** at any time; 18586 no two lines shall have the same new delta SID. 18587 The z-file shall serve as a lock-out mechanism against simultaneous updates. Its contents shall 18588 be the binary process ID of the command (that is, get) that created it. The z-file shall be created 18589 in the directory containing the SCCS file for the duration of get. The same protection restrictions 18590 18591 as those for the **p-file** shall apply for the **z-file**. The **z-file** shall be created read-only.

Utilities get

18592 EXTENDED DESCRIPTION

	Det	ermination of SCCS Identification	on String	
SID* Specified	–b Keyletter Used†	Other Conditions	SID Retrieved	SID of Delta to be Created
none‡	no	R defaults to mR	mR.mL	mR.(mL+1)
none‡	yes	R defaults to mR	mR.mL	mR.mL.(mB+1).
R	no	R > mR	mR.mL	R.1***
R	no	R = mR	mR.mL	mR.(mL+1)
R	yes	R > mR	mR.mL	mR.mL.(mB+1).
R	yes	R = mR	mR.mL	mR.mL.(mB+1).
R	_	R < mR and R does not exist	hR.mL**	hR.mL.(mB+1).1
R	_	$\label{eq:Trunk successor} Trunk \ successor \ in \ release > R \\ and \ R \ exists$	R.mL	R.mL.(mB+1).1
R.L	no	No trunk successor	R.L	R.(L+1)
R.L	yes	No trunk successor	R.L	R.L.(mB+1).1
R.L	_	Trunk successor in release ≥ R	R.L	R.L.(mB+1).1
R.L.B	no	No branch successor	R.L.B.mS	R.L.B.(mS+1)
R.L.B	yes	No branch successor	R.L.B.mS	R.L.(mB+1).1
R.L.B.S	no	No branch successor	R.L.B.S	R.L.B.(S+1)
R.L.B.S	yes	No branch successor	R.L.B.S	R.L.(mB+1).1
R.L.B.S	_	Branch successor	R.L.B.S	R.L.(mB+1).1

- R, L, B, and S are the release, level, branch, and sequence components of the SID, respectively; m means maximum. Thus, for example, R.mL means "the maximum level number within release R"; R.L.(mB+1).1 means "the first sequence number on the new branch (that is, maximum branch number plus one) of level L within release R". Note that if the SID specified is of the form R.L, R.L.B, or R.L.B.S, each of the specified components shall exist.
- ** hR is the highest existing release that is lower than the specified, nonexistent, release R.
- *** This is used to force creation of the first delta in a new release.
- † The $-\mathbf{b}$ option is effective only if the \mathbf{b} flag is present in the file. An entry of '-' means "irrelevant".
- This case applies if the **d** (default SID) flag is not present in the file. If the **d** flag is present in the file, then the SID obtained from the **d** flag is interpreted as if it had been specified on the command line. Thus, one of the other cases in this table applies.

get Utilities

18628	System Dat	System Date and Time		
18629 18630	When a g-file is generated, the creation time of deltas in the SCCS file may be taken into account. If any of these times are apparently in the future, the behavior is unspecified.			
18631	Identificati	on Keywords		
18632 18633 18634	identificatio	Identifying information shall be inserted into the text retrieved from the SCCS file by replacing identification keywords with their value wherever they occur. The following keywords may be used in the text stored in an SCCS file:		
18635 18636	% M %	Module name: either the value of the ${\bf m}$ flag in the file, or if absent, the name of the SCCS file with the leading ${\bf s}$. removed.		
18637 18638	% I %	SCCS identification (SID) (% \mathbf{R} %.% \mathbf{L} % or % \mathbf{R} %.% \mathbf{L} %.% \mathbf{B} %.% \mathbf{S} %) of the retrieved text.		
18639	% R %	Release.		
18640	% L %	Level.		
18641	% B %	Branch.		
18642	% S %	Sequence.		
18643	% D %	Current date (YY/MM/DD).		
18644	% H %	Current date (MM/DD/YY).		
18645	% T %	Current time (HH:MM:SS).		
18646	% E %	Date newest applied delta was created (YY/MM/DD).		
18647	% G %	Date newest applied delta was created (MM/DD/YY).		
18648	% U %	Time newest applied delta was created (HH:MM:SS).		
18649	% Y %	Module type: value of the t flag in the SCCS file.		
18650	% F %	SCCS filename.		
18651	% P %	SCCS absolute pathname.		
18652	$%\mathbf{Q}\%$	The value of the $\bf q$ flag in the file.		
18653 18654 18655	% C %	Current line number. This keyword is intended for identifying messages output by the program, such as "this should not have happened" type errors. It is not intended to be used on every line to provide sequence numbers.		
18656	% Z %	The four-character string "@(#)" recognizable by what.		
18657	% W %	A shorthand notation for constructing what strings:		
18658		%W%=%Z%%M% <tab>%I%</tab>		
18659	% A %	Another shorthand notation for constructing what strings:		
18660		%A%=%Z%%Y%%M%%I%%Z%		
18661 EXIT STATUS 18662 The following exit values shall be returned:				
18663	0 Successful completion.			
18664	>0 An erro	or occurred.		

Utilities get

18665 CONSEQUENCES OF ERRORS

18666 Default.

18667 APPLICATION USAGE

Problems can arise if the system date and time have been modified (for example, put forward and then back again, or unsynchronized clocks across a network) and can also arise when different values of the *TZ* environment variable are used.

Problems of a similar nature can also arise for the operation of the *delta* utility, which compares the previous file body against the working file as part of its normal operation.

18673 EXAMPLES

18674 None.

18675 RATIONALE

18676 None.

18677 FUTURE DIRECTIONS

The -lp option may be withdrawn in a future version.

18679 SEE ALSO

18680 admin, delta, prs, what

18681 CHANGE HISTORY

18682 First released in Issue 2.

18683 Issue 5

18684 A correction is made to the first format string in STDOUT.

The interpretation of the YY component of the -c *cutoff* argument is noted.

18686 **Issue 6**

18693 18694

18695

18696

18697

The obsolescent SYNOPSIS is removed, removing the -lp option.

The normative text is reworded to avoid use of the term "must" for application requirements.

The Open Group Corrigendum U025/5 is applied, correcting text in the OPTIONS section.

The Open Group Corrigendum U048/1 is applied.

The Open Group Interpretation PIN4C.00014 is applied.

The Open Group Base Resolution bwg2001-007 is applied as follows:

- The EXTENDED DESCRIPTION section is updated to make partial SID handling unspecified, reflecting common usage, and to clarify SID ranges.
- New text is added to the EXTENDED DESCRIPTION and APPLICATION USAGE sections regarding how the system date and time may be taken into account.
 - The TZ environment variable is added to the ENVIRONMENT VARIABLES section.

getconf Utilities

18698 NAME		
18699	getconf — g	et configuration values
18700 SYNOP	PSIS	
18701	getconf [-v specification] system_var
18702	getconf [-v specification] path_var pathname
18703 DESCR	IPTION	
18704 18705		synopsis form, the <i>getconf</i> utility shall write to the standard output the value of the cified by the <i>system_var</i> operand.
18706 18707		d synopsis form, the <i>getconf</i> utility shall write to the standard output the value of the cified by the <i>path_var</i> operand for the path specified by the <i>pathname</i> operand.
18708 18709 18710 18711	function from System Inter	f each configuration variable shall be determined as if it were obtained by calling the m which it is defined to be available by this volume of IEEE Std 1003.1-2001 or by the rfaces volume of IEEE Std 1003.1-2001 (see the OPERANDS section). The value shall itions in the current operating environment.
18712 OPTIO	NS	
18713 18714		utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section Syntax Guidelines.
18715	The following	ng option shall be supported:
18716	–v specificati	ion
18717	•	Indicate a specific specification and version for which configuration variables shall
18718 18719		be determined. If this option is not specified, the values returned correspond to an implementation default conforming compilation environment.
18720		If the command:
18721		getconf POSIX V6 ILP32 OFF32
18722		does not write "-1\n" or "undefined\n" to standard output, then commands of
18723		the form:
18724		getconf -v POSIX_V6_ILP32_OFF32
18725		determine values for configuration variables corresponding to the
18726 18727		POSIX_V6_ILP32_OFF32 compilation environment specified in <i>c99</i> , the EXTENDED DESCRIPTION.
18728		If the command:
18729		getconf _POSIX_V6_ILP32_OFFBIG
18730		does not write "-1\n" or "undefined\n" to standard output, then commands of
18731		the form:
18732		getconf -v POSIX_V6_ILP32_OFFBIG
18733		determine values for configuration variables corresponding to the
18734 18735		POSIX_V6_ILP32_OFFBIG compilation environment specified in <i>c99</i> , the EXTENDED DESCRIPTION.
18736		If the command:
18737		getconf _POSIX_V6_LP64_OFF64
18738 18739		does not write "-1\n" or "undefined\n" to standard output, then commands of the form:

Utilities getconf

18740		getconf -v POSIX_V6_LP64_OFF64
18741		determine values for configuration variables corresponding to the
18742		POSIX_V6_LP64_OFF64 compilation environment specified in c99, the
18743		EXTENDED DESCRIPTION.
18744		If the command:
18745		getconf _POSIX_V6_LPBIG_OFFBIG
18746		does not write "-1\n" or "undefined\n" to standard output, then commands of
18747		the form:
18748		getconf -v POSIX_V6_LPBIG_OFFBIG
18749		determine values for configuration variables corresponding to the
18750		POSIX_V6_LPBIG_OFFBIG compilation environment specified in c99, the
18751		EXTENDED DESCRIPTION.
18752 OPERA	NDS	
18753	The following	ng operands shall be supported:
18754	path_var	A name of a configuration variable. All of the variables in the Variable column of
18755		the table in the DESCRIPTION of the fpathconf() function defined in the System
18756		Interfaces volume of IEEE Std 1003.1-2001, without the enclosing braces, shall be
18757		supported. The implementation may add other local variables.
18758	pathname	A pathname for which the variable specified by <i>path_var</i> is to be determined.
18759	system_var	A name of a configuration variable. All of the following variables shall be
18760		supported:
18761		• The names in the Variable column of the table in the DESCRIPTION of the
18762		sysconf() function in the System Interfaces volume of IEEE Std 1003.1-2001,
18763		except for the entries corresponding to _SC_CLK_TCK,
18764		_SC_GETGR_R_SIZE_MAX, and _SC_GETPW_R_SIZE_MAX, without the
18765		enclosing braces.
18766		For compatibility with earlier versions, the following variable names shall also
18767		be supported:
18768		POSIX2_C_BIND
18769		POSIX2_C_DEV
18770		POSIX2_CHAR_TERM
18771		POSIX2_FORT_DEV
18772		POSIX2_FORT_RUN
18773		POSIX2_LOCALEDEF
18774		POSIX2_SW_DEV
18775		POSIX2_UPE
18776		POSIX2_VERSION
18777		and shall be equivalent to the same name prefixed with an underscore. This
18778		requirement may be removed in a future version.
18779		• The names of the symbolic constants used as the <i>name</i> argument of the <i>confstr()</i>
18780		function in the System Interfaces volume of IEEE Std 1003.1-2001, without the
18781		_CS_ prefix.
18782		• The names of the symbolic constants listed under the headings "Maximum
18783		Values" and "Minimum Values" in the description of the limits.h> header in

getconf Utilities

18784 18785		the Base Definitions volume of IEEE Std 1003.1-2001, without the enclosing braces.			
18786 18787		For compatibility with earlier versions, the following variable names shall also be supported:			
18788 18789 18790 18791 18792 18793 18794 18795		POSIX2_BC_BASE_MAX POSIX2_BC_DIM_MAX POSIX2_BC_SCALE_MAX POSIX2_BC_STRING_MAX POSIX2_COLL_WEIGHTS_MAX POSIX2_EXPR_NEST_MAX POSIX2_LINE_MAX POSIX2_RE_DUP_MAX			
18796 18797		and shall be equivalent to the same name prefixed with an underscore. This requirement may be removed in a future version.			
18798		The implementation may add other local values.			
18799 STDIN 18800	Not used.				
18801 INPUT 18802	FILES None.				
18803 ENVIR 0 18804		ARIABLES ag environment variables shall affect the execution of <i>getconf</i> :			
18805 18806 18807 18808	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)			
18809 18810	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.			
18811 18812 18813	LC_CTYPE	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).			
18814	LC_MESSA				
18815 18816		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.			
18817 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .			
18818 ASYNC 18819	18818 ASYNCHRONOUS EVENTS 18819 Default.				
18820 STDOU					
18821 18822 18823	If the specified variable is defined on the system and its value is described to be available from the <i>confstr()</i> function defined in the System Interfaces volume of IEEE Std 1003.1-2001, its value shall be written in the following format:				
18824	"%s\n", <	value>			
18825 18826	Otherwise, if the specified variable is defined on the system, its value shall be written in the following format:				

Utilities getconf

```
18827
             "%d\n", <value>
             If the specified variable is valid, but is undefined on the system, getconf shall write using the
18828
18829
             following format:
18830
             "undefined\n"
             If the variable name is invalid or an error occurs, nothing shall be written to standard output.
18831
18832 STDERR
             The standard error shall be used only for diagnostic messages.
18833
18834 OUTPUT FILES
             None.
18835
18836 EXTENDED DESCRIPTION
18837
             None.
18838 EXIT STATUS
18839
             The following exit values shall be returned:
                 The specified variable is valid and information about its current state was written
18840
                 successfully.
18841
             >0 An error occurred.
18842
18843 CONSEQUENCES OF ERRORS
             Default.
18844
18845 APPLICATION USAGE
18846
             None.
18847 EXAMPLES
             The following example illustrates the value of {NGROUPS_MAX}:
18848
18849
             getconf NGROUPS MAX
18850
             The following example illustrates the value of {NAME_MAX} for a specific directory:
18851
             getconf NAME MAX /usr
             The following example shows how to deal more carefully with results that might be unspecified:
18852
             if value=$(getconf PATH MAX /usr); then
18853
                  if [ "$value" = "undefined" ]; then
18854
18855
                       echo PATH MAX in /usr is infinite.
18856
                  else
18857
                       echo PATH MAX in /usr is $value.
                  fi
18858
             else
18859
18860
                  echo Error in getconf.
             fi
18861
             Note that:
18862
             sysconf ( SC POSIX C BIND);
18863
             and:
18864
18865
             system("getconf POSIX2 C BIND");
18866
             in a C program could give different answers. The sysconf() call supplies a value that corresponds
             to the conditions when the program was either compiled or executed, depending on the
18867
```

getconf Utilities

implementation; the *system()* call to *getconf* always supplies a value corresponding to conditions when the program is executed.

18870 RATIONALE

The original need for this utility, and for the *confstr*() function, was to provide a way of finding the configuration-defined default value for the *PATH* environment variable. Since *PATH* can be modified by the user to include directories that could contain utilities replacing the standard utilities, shell scripts need a way to determine the system-supplied *PATH* environment variable value that contains the correct search path for the standard utilities. It was later suggested that access to the other variables described in this volume of IEEE Std 1003.1-2001 could also be useful to applications.

This functionality of *getconf* would not be adequately subsumed by another command such as:

18879 grep var /etc/conf

because such a strategy would provide correct values for neither those variables that can vary at runtime, nor those that can vary depending on the path.

Early proposal versions of *getconf* specified exit status 1 when the specified variable was valid, but not defined on the system. The output string "undefined" is now used to specify this case with exit code 0 because so many things depend on an exit code of zero when an invoked utility is successful.

18886 FUTURE DIRECTIONS

18887 None.

18888 SEE ALSO

c99, the Base Definitions volume of IEEE Std 1003.1-2001, limits.h>, the System Interfaces | volume of IEEE Std 1003.1-2001, confstr(), pathconf(), system()

18891 CHANGE HISTORY

18892 First released in Issue 4.

18893 Issue 5

18894 In the OPERANDS section:

• {NL_MAX} is changed to {NL_NMAX}.

Entries beginning NL_ are deleted from the list of standard configuration variables.

The list of variables previously marked UX is merged with the list marked EX.

Operands are added to support new Option Groups.

• Operands are added so that *getconf* can determine supported programming environments.

18900 Issue 6

18898

18901

18902

18905

The Open Group Corrigendum U029/4 is applied, correcting the example command in the last paragraph of the OPTIONS section.

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

• Operands are added to determine supported programming environments.

This reference page is updated for alignment with the ISO/IEC 9899: 1999 standard. Specifically, new macros for *c99* programming environments are introduced.

18908 XSI marked *system_var* (XBS5_*) values are marked LEGACY.

Utilities getconf

18909 18910 IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/27 is applied, correcting the descriptions of *path_var* and *system_var* in the OPERANDS section.

getopts Utilities

```
    18911 NAME
    18912 getopts — parse utility options
    18913 SYNOPSIS
    18914 getopts optstring name [arg...]
```

DESCRIPTION

 The *getopts* utility shall retrieve options and option-arguments from a list of parameters. It shall support the Utility Syntax Guidelines 3 to 10, inclusive, described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

Each time it is invoked, the *getopts* utility shall place the value of the next option in the shell variable specified by the *name* operand and the index of the next argument to be processed in the shell variable *OPTIND*. Whenever the shell is invoked, *OPTIND* shall be initialized to 1.

When the option requires an option-argument, the *getopts* utility shall place it in the shell variable *OPTARG*. If no option was found, or if the option that was found does not have an option-argument, *OPTARG* shall be unset.

If an option character not contained in the *optstring* operand is found where an option character is expected, the shell variable specified by *name* shall be set to the question-mark ('?') character. In this case, if the first character in *optstring* is a colon (':'), the shell variable *OPTARG* shall be set to the option character found, but no output shall be written to standard error; otherwise, the shell variable *OPTARG* shall be unset and a diagnostic message shall be written to standard error. This condition shall be considered to be an error detected in the way arguments were presented to the invoking application, but shall not be an error in *getopts* processing.

If an option-argument is missing:

- If the first character of *optstring* is a colon, the shell variable specified by *name* shall be set to the colon character and the shell variable *OPTARG* shall be set to the option character found.
- Otherwise, the shell variable specified by *name* shall be set to the question-mark character, the shell variable *OPTARG* shall be unset, and a diagnostic message shall be written to standard error. This condition shall be considered to be an error detected in the way arguments were presented to the invoking application, but shall not be an error in *getopts* processing; a diagnostic message shall be written as stated, but the exit status shall be zero.

When the end of options is encountered, the *getopts* utility shall exit with a return value greater than zero; the shell variable *OPTIND* shall be set to the index of the first non-option-argument, where the first "--" argument is considered to be an option-argument if there are no other non-option-arguments appearing before it, or the value "\$#"+1 if there are no non-option-arguments; the *name* variable shall be set to the question-mark character. Any of the following shall identify the end of options: the special option "--", finding an argument that does not begin with a '-', or encountering an error.

The shell variables *OPTIND* and *OPTARG* shall be local to the caller of *getopts* and shall not be exported by default.

The shell variable specified by the *name* operand, *OPTIND*, and *OPTARG* shall affect the current shell execution environment; see Section 2.12 (on page 61).

If the application sets *OPTIND* to the value 1, a new set of parameters can be used: either the current positional parameters or new *arg* values. Any other attempt to invoke *getopts* multiple times in a single shell execution environment with parameters (positional parameters or *arg* operands) that are not the same in all invocations, or with an *OPTIND* value modified to be a value other than 1, produces unspecified results.

getopts **Utilities**

18956 OPTIONS

18957 None.

18958 OPERANDS

The following operands shall be supported: 18959

optstring A string containing the option characters recognized by the utility invoking *getopts*. 18960 If a character is followed by a colon, the option shall be expected to have an 18961 argument, which should be supplied as a separate argument. Applications should 18962 specify an option character and its option-argument as separate arguments, but 18963 getopts shall interpret the characters following an option character requiring 18964 18965 arguments as an argument whether or not this is done. An explicit null optionargument need not be recognized if it is not supplied as a separate argument when 18966 getopts is invoked. (See also the getopt() function defined in the System Interfaces 18967 volume of IEEE Std 1003.1-2001.) The characters question-mark and colon shall not 18968 be used as option characters by an application. The use of other option characters 18969 that are not alphanumeric produces unspecified results. If the option-argument is 18970 not supplied as a separate argument from the option character, the value in 18971 *OPTARG* shall be stripped of the option character and the '-'. The first character 18972 in optstring determines how getopts behaves if an option character is not known or 18973

an option-argument is missing. 18974

name The name of a shell variable that shall be set by the *getopts* utility to the option 18975 18976

character that was found.

The getopts utility by default shall parse positional parameters passed to the invoking shell 18977 procedure. If args are given, they shall be parsed instead of the positional parameters. 18978

18979 **STDIN**

18984

18992

18993

Not used. 18980

18981 INPUT FILES

18982 None.

18983 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *getopts*:

LANG Provide a default value for the internationalization variables that are unset or null. 18985 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 18986 Internationalization Variables for the precedence of internationalization variables 18987 18988

used to determine the values of locale categories.)

18989 LC_ALL If set to a non-empty string value, override the values of all the other 18990

internationalization variables.

LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 18991

characters (for example, single-byte as opposed to multi-byte characters in

arguments and input files).

LC_MESSAGES 18994

Determine the locale that should be used to affect the format and contents of 18995

diagnostic messages written to standard error. 18996

NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 18997 XSI

OPTIND 18998 This variable shall be used by the *getopts* utility as the index of the next argument

to be processed. 18999

getopts Utilities

19000 ASYNCHRONOUS EVENTS

19001 Default.

19002 STDOUT

19003 Not used.

19004 STDERR

19008

19009

19010

19015

19016 19017

19025

Whenever an error is detected and the first character in the *optstring* operand is not a colon (':'), a diagnostic message shall be written to standard error with the following information in an unspecified format:

• The invoking program name shall be identified in the message. The invoking program name shall be the value of the shell special parameter 0 (see Section 2.5.2 (on page 34)) at the time the *getopts* utility is invoked. A name equivalent to:

19011 basename "\$0"

may be used.

• If an option is found that was not specified in *optstring*, this error is identified and the invalid option character shall be identified in the message.

 If an option requiring an option-argument is found, but an option-argument is not found, this error shall be identified and the invalid option character shall be identified in the message.

19018 OUTPUT FILES

19019 None.

19020 EXTENDED DESCRIPTION

19021 None.

19022 EXIT STATUS

19023 The following exit values shall be returned:

19024 0 An option, specified or unspecified by *optstring*, was found.

>0 The end of options was encountered or an error occurred.

19026 CONSEQUENCES OF ERRORS

19027 Default.

19028 APPLICATION USAGE

Since *getopts* affects the current shell execution environment, it is generally provided as a shell regular built-in. If it is called in a subshell or separate utility execution environment, such as one of the following:

```
19032 (getopts abc value "$@")
19033 nohup getopts ...
19034 find . -exec getopts ... \;
```

it does not affect the shell variables in the caller's environment.

Note that shell functions share *OPTIND* with the calling shell even though the positional parameters are changed. If the calling shell and any of its functions uses *getopts* to parse arguments, the results are unspecified.

19039 EXAMPLES

19040 The following example script parses and displays its arguments:

19041 aflag= 19042 bflag= Utilities getopts

```
19043
            while getopts ab: name
19044
19045
                case $name in
                       aflag=1;;
                a)
19046
19047
                b)
                       bflaq=1
                       bval="$OPTARG";;
19048
                      printf "Usage: %s: [-a] [-b value] args\n" $0
19049
                 ?)
                       exit 2;;
19050
19051
                 esac
19052
            done
19053
            if [ ! -z "\$aflag" ]; then
19054
                printf "Option -a specified\n"
            fi
19055
            if [ ! -z "$bflag" ]; then
19056
                printf 'Option -b "%s" specified\n' "$bval"
19057
            fi
19058
            shift $(($OPTIND - 1))
19059
            printf "Remaining arguments are: %s\n" "$*"
19060
```

19061 RATIONALE

19062

19063

19064

19065 19066

19067 19068

19069

19070

19071

19072 19073

19074

19075

19076

19077

19078

19079

19080 19081

19082

19086

19087

The *getopts* utility was chosen in preference to the System V *getopt* utility because *getopts* handles option-arguments containing <blank>s.

The *OPTARG* variable is not mentioned in the ENVIRONMENT VARIABLES section because it does not affect the execution of *getopts*; it is one of the few "output-only" variables used by the standard utilities.

The colon is not allowed as an option character because that is not historical behavior, and it violates the Utility Syntax Guidelines. The colon is now specified to behave as in the KornShell version of the *getopts* utility; when used as the first character in the *optstring* operand, it disables diagnostics concerning missing option-arguments and unexpected option characters. This replaces the use of the *OPTERR* variable that was specified in an early proposal.

The formats of the diagnostic messages produced by the *getopts* utility and the *getopt()* function are not fully specified because implementations with superior ("friendlier") formats objected to the formats used by some historical implementations. The standard developers considered it important that the information in the messages used be uniform between *getopts* and getopt(). Exact duplication of the messages might not be possible, particularly if a utility is built on another system that has a different getopt() function, but the messages must have specific information included so that the program name, invalid option character, and type of error can be distinguished by a user.

Only a rare application program intercepts a *getopts* standard error message and wants to parse it. Therefore, implementations are free to choose the most usable messages they can devise. The following formats are used by many historical implementations:

Historical shells with built-in versions of *getopt()* or *getopts* have used different formats, frequently not even indicating the option character found in error.

getopts Utilities

19088 FUTURE DIRECTIONS

19089 None.

19090 SEE ALSO

Section 2.5.2 (on page 34), the System Interfaces volume of IEEE Std 1003.1-2001, getopt()

19092 CHANGE HISTORY

19093 First released in Issue 4.

19094 **Issue 6**

The normative text is reworded to avoid use of the term "must" for application requirements.

Utilities grep

```
19096 NAME
19097 grep — search a file for a pattern
19098 SYNOPSIS
19099 grep [-E| -F] [-c| -1| -q] [-insvx] -e pattern_list...
19100 [-f pattern_file] ... [file...]
19101 grep [-E| -F] [-c| -1| -q] [-insvx] [-e pattern_list] ...
19102 -f pattern_file... [file...]
19103 grep [-E| -F] [-c| -1| -q] [-insvx] pattern_list[file...]

DESCRIPTION
```

19104 DESCRIPTION

The *grep* utility shall search the input files, selecting lines matching one or more patterns; the types of patterns are controlled by the options specified. The patterns are specified by the –e option, –f option, or the *pattern_list* operand. The *pattern_list*'s value shall consist of one or more patterns separated by <newline>s; the *pattern_file*'s contents shall consist of one or more patterns terminated by <newline>. By default, an input line shall be selected if any pattern, treated as an entire basic regular expression (BRE) as described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.3, Basic Regular Expressions, matches any part of the line excluding the terminating <newline>; a null BRE shall match every line. By default, each selected input line shall be written to the standard output.

Regular expression matching shall be based on text lines. Since a <newline> separates or terminates patterns (see the -e and -f options below), regular expressions cannot contain a <newline>. Similarly, since patterns are matched against individual lines (excluding the terminating <newline>s) of the input, there is no way for a pattern to match a <newline> found in the input.

19119 OPTIONS

The *grep* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported:

Match using extended regular expressions. Treat each pattern specified as an ERE, as described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.4, Extended Regular Expressions. If any entire ERE pattern matches some part of an input line excluding the terminating <newline>, the line shall be matched. A null ERE shall match every line.

Match using fixed strings. Treat each pattern specified as a string instead of a regular expression. If an input line contains any of the patterns as a contiguous sequence of bytes, the line shall be matched. A null string shall match every line.

19131 —c Write only a count of selected lines to standard output.

–e *pattern_list*

 $-\mathbf{F}$

Specify one or more patterns to be used during the search for input. The application shall ensure that patterns in *pattern_list* are separated by a <newline>. A null pattern can be specified by two adjacent <newline>s in *pattern_list*. Unless the –E or –F option is also specified, each pattern shall be treated as a BRE, as described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.3, Basic Regular Expressions. Multiple –e and –f options shall be accepted by the *grep* utility. All of the specified patterns shall be used when matching lines, but the order of evaluation is unspecified.

grep Utilities

19141	- f pattern_fi	le	
19142	- F	Read one or more patterns from the file named by the pathname pattern_file.	
19143		Patterns in <i>pattern_file</i> shall be terminated by a <newline>. A null pattern can be</newline>	
19144		specified by an empty line in pattern_file. Unless the -E or -F option is also	
19145		specified, each pattern shall be treated as a BRE, as described in the Base	
19146		Definitions volume of IEEE Std 1003.1-2001, Section 9.3, Basic Regular Expressions.	
19147	−i	Perform pattern matching in searches without regard to case; see the Base	
19148		Definitions volume of IEEE Std 1003.1-2001, Section 9.2, Regular Expression	
19149		General Requirements.	
19150	-l	(The letter ell.) Write only the names of files containing selected lines to standard	
19151		output. Pathnames shall be written once per file searched. If the standard input is	
19152		searched, a pathname of "(standard input)" shall be written, in the POSIX locale. In other locales, "standard input" may be replaced by something more	
19153 19154		appropriate in those locales.	
19155	-n	Precede each output line by its relative line number in the file, each file starting at	
19156		line 1. The line number counter shall be reset for each file processed.	
19157 19158	-q	Quiet. Nothing shall be written to the standard output, regardless of matching lines. Exit with zero status if an input line is selected.	
19159 19160	-s	Suppress the error messages ordinarily written for nonexistent or unreadable files. Other error messages shall not be suppressed.	
19161 19162	- v	Select lines not matching any of the specified patterns. If the -v option is not specified, selected lines shall be those that match any of the specified patterns.	
19163	- x	Consider only input lines that use all characters in the line excluding the	
19164	A	terminating <newline> to match an entire fixed string or regular expression to be</newline>	
19165		matching lines.	
19166 OPERA	NDS		
19167	The following operands shall be supported:		
19168 19169	pattern_list	Specify one or more patterns to be used during the search for input. This operand shall be treated as if it were specified as –e <i>pattern_list</i> .	
19170 19171	file	A pathname of a file to be searched for the patterns. If no <i>file</i> operands are specified, the standard input shall be used.	
19172 STDIN			
19172 STD1 11		d input shall be used only if no file operands are specified. See the INPUT FILES	
19174	section.	The state of the s	
19175 INPUT	FILES		
19176		les shall be text files.	
	-		
19177 EINVIR 19178	19177 ENVIRONMENT VARIABLES 19178 The following environment variables shall affect the execution of <i>grep</i> :		
19179	LANG	Provide a default value for the internationalization variables that are unset or null.	
19180		(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,	
19181		Internationalization Variables for the precedence of internationalization variables	
19182		used to determine the values of locale categories.)	
19183	LC_ALL	If set to a non-empty string value, override the values of all the other	
19184		internationalization variables.	

Utilities grep

19185 19186 19187	LC_COLLAT	TE Determine the locale for the behavior of ranges, equivalence classes, and multi- character collating elements within regular expressions.		
19188 19189 19190 19191	LC_CTYPE	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) and the behavior of character classes within regular expressions.		
19192 19193 19194	LC_MESSA	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.		
19195 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.		
19196 ASYNCHRONOUS EVENTS 19197 Default.				
19198 STDOU	J T			
19199 19200	If the $-\mathbf{l}$ option is in effect, and the $-\mathbf{q}$ option is not, the following shall be written for each file containing at least one selected input line:			
19201	"%s\n", <file></file>			
19202 19203	Otherwise, if more than one <i>file</i> argument appears, and $-\mathbf{q}$ is not specified, the <i>grep</i> utility shall prefix each output line by:			
19204	"%s:", <file></file>			
19205	The remainder of each output line shall depend on the other options specified:			
19206	• If the $-c$ option is in effect, the remainder of each output line shall contain:			
19207	"%d\n",	, <count></count>		
19208 19209	• Otherwise, if $-c$ is not in effect and the $-n$ option is in effect, the following shall be written to standard output:			
19210	"%d:",	eline number>		
19211	• Finally, t	he following shall be written to standard output:		
19212	"%s",	selected-line contents>		
19213 STDER	R			
19214	The standar	d error shall be used only for diagnostic messages.		
19215 OUTPUT FILES 19216 None.				
19217 EXTENDED DESCRIPTION 19218 None.				
19219 EXIT S				
19220	o de la companya de			
19221		more lines were selected.		
19222	1 No line	s were selected.		
19223	>1 An erro	r occurred.		

grep Utilities

19224 CONSEQUENCES OF ERRORS

If the $-\mathbf{q}$ option is specified, the exit status shall be zero if an input line is selected, even if an error was detected. Otherwise, default actions shall be performed.

19227 APPLICATION USAGE

Care should be taken when using characters in *pattern_list* that may also be meaningful to the command interpreter. It is safest to enclose the entire *pattern_list* argument in single quotes:

19230 ' . . . '

19228

19229

19234

19235

19236

19237

19238 19239

19240

19241 19242

19244

19245

19246 19247

19248

19249 19250

19251

19254

19260

19261

19262

The **–e** *pattern_list* option has the same effect as the *pattern_list* operand, but is useful when pattern_list begins with the hyphen delimiter. It is also useful when it is more convenient to provide multiple patterns as separate arguments.

Multiple –e and –f options are accepted and *grep* uses all of the patterns it is given while matching input text lines. (Note that the order of evaluation is not specified. If an implementation finds a null string as a pattern, it is allowed to use that pattern first, matching every line, and effectively ignore any other patterns.)

The $-\mathbf{q}$ option provides a means of easily determining whether or not a pattern (or string) exists in a group of files. When searching several files, it provides a performance improvement (because it can quit as soon as it finds the first match) and requires less care by the user in choosing the set of files to supply as arguments (because it exits zero if it finds a match even if *grep* detected an access or read error on earlier *file* operands).

19243 EXAMPLES

 To find all uses of the word "Posix" (in any case) in file text.mm and write with line numbers:

```
grep -i -n posix text.mm
```

2. To find all empty lines in the standard input:

```
grep ^$
or:
grep -v .
```

3. Both of the following commands print all lines containing strings "abc" or "def" or both:

```
19252 grep -E 'abc|def'
19253 grep -F 'abc
```

def'

19255 4. Both of the following commands print all lines matching exactly "abc" or "def":

```
19256 grep -E '^abc$|^def$'
19257 grep -F -x 'abc
19258 def'
```

19259 RATIONALE

This *grep* has been enhanced in an upwards-compatible way to provide the exact functionality of the historical *egrep* and *fgrep* commands as well. It was the clear intention of the standard developers to consolidate the three *greps* into a single command.

The old *egrep* and *fgrep* commands are likely to be supported for many years to come as implementation extensions, allowing historical applications to operate unmodified.

Utilities grep

	19265 19266	Historical implementations usually silently ignored all but one of multiply-specified —e and —f options, but were not consistent as to which specification was actually used.				
	19267 19268	The $-\mathbf{b}$ option was omitted from the OPTIONS section because block numbers are implementation-defined.				
	19269	The System V restriction on using – to mean standard input was omitted.				
	19270 19271	A definition of action taken when given a null BRE or ERE is specified. This is an error condition in some historical implementations.				
	19272 19273 19274	The –l option previously indicated that its use was undefined when no files were explicitly named. This behavior was historical and placed an unnecessary restriction on future implementations. It has been removed.				
	19275 19276	The historical BSD $grep$ –s option practice is easily duplicated by redirecting standard output to /dev/null. The –s option required here is from System V.				
	19277 19278	The $-\mathbf{x}$ option, historically available only with \textit{fgrep} , is available here for all of the non-obsolescent versions.				
	19279 FUTUR 19280	RE DIRECTIONS None.				
19281 SEE ALSO						
	19282	sed				
	19283 CHAN (19284	GE HISTORY First released in Issue 2.				
	19285 Issue 6					
	19286	The Open Group Corrigendum U029/5 is applied, correcting the SYNOPSIS.				
	19287	The normative text is reworded to avoid use of the term "must" for application requirements.				
	19288 19289	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/28 is applied, correcting the examples using the $grep$ – \mathbf{F} option which did not match the normative description of the – \mathbf{F} option.				

hash Utilities

19290 NAME 19291 hash — remember or report utility locations 19292 SYNOPSIS hash [utility...] 19293 XSI 19294 hash -r 19295 19296 **DESCRIPTION** The hash utility shall affect the way the current shell environment remembers the locations of 19297 utilities found as described in Section 2.9.1.1 (on page 48). Depending on the arguments 19298 specified, it shall add utility locations to its list of remembered locations or it shall purge the 19299 contents of the list. When no arguments are specified, it shall report on the contents of the list. 19300 Utilities provided as built-ins to the shell shall not be reported by *hash*. 19301 19302 OPTIONS The hash utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 19303 12.2, Utility Syntax Guidelines. 19304 The following option shall be supported: 19305 Forget all previously remembered utility locations. 19306 $-\mathbf{r}$ 19307 OPERANDS The following operand shall be supported: 19308 The name of a utility to be searched for and added to the list of remembered utility 19309 locations. If *utility* contains one or more slashes, the results are unspecified. 19310 19311 **STDIN** 19312 Not used. 19313 INPUT FILES 19314 None. 19315 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *hash*: 19316 LANG Provide a default value for the internationalization variables that are unset or null. 19317 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 19318 Internationalization Variables for the precedence of internationalization variables 19319 19320 used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other 19321 internationalization variables. 19322 Determine the locale for the interpretation of sequences of bytes of text data as 19323 LC_CTYPE characters (for example, single-byte as opposed to multi-byte characters in 19324 19325 arguments). LC_MESSAGES 19326 19327 Determine the locale that should be used to affect the format and contents of 19328 diagnostic messages written to standard error.

19329

19330 19331 **NLSPATH**

PATH

Determine the location of message catalogs for the processing of *LC_MESSAGES*.

IEEE Std 1003.1-2001, Chapter 8, Environment Variables.

Determine the location of *utility*, as described in the Base Definitions volume of

Utilities hash

19332 ASYNCHRONOUS EVENTS

19333 Default.

19334 **STDOUT**

The standard output of *hash* shall be used when no arguments are specified. Its format is unspecified, but includes the pathname of each utility in the list of remembered locations for the current shell environment. This list shall consist of those utilities named in previous *hash* invocations that have been invoked, and may contain those invoked and found through the normal command search process.

19340 STDERR

19341 The standard error shall be used only for diagnostic messages.

19342 OUTPUT FILES

19343 None.

19344 EXTENDED DESCRIPTION

19345 None.

19346 EXIT STATUS

19347 The following exit values shall be returned:

19348 0 Successful completion.

19349 >0 An error occurred.

19350 CONSEQUENCES OF ERRORS

19351 Default.

19352 APPLICATION USAGE

Since hash affects the current shell execution environment, it is always provided as a shell regular built-in. If it is called in a separate utility execution environment, such as one of the

19355 following:

19356 nohup hash -r

19357 find . -type f | xargs hash

it does not affect the command search process of the caller's environment.

The *hash* utility may be implemented as an alias—for example, *alias* –**t** –, in which case utilities found through normal command search are not listed by the *hash* command.

The effects of hash –**r** can also be achieved portably by resetting the value of PATH; in the simplest form, this can be:

19363 PATH="\$PATH"

The use of *hash* with *utility* names is unnecessary for most applications, but may provide a performance improvement on a few implementations; normally, the hashing process is included by default

19366 by default.

19367 EXAMPLES

19368 None.

19369 RATIONALE

19370 None.

19371 FUTURE DIRECTIONS

19372 None.

hash Utilities

19373 **SEE ALSO**

19374 Section 2.9.1.1 (on page 48)

19375 CHANGE HISTORY

19376 First released in Issue 2.

Utilities head

	NAME								
19378		head — copy the first part of files							
	SYNOPS								
19380			number][file]						
19381 19382 19383			ity shall copy its input files to the standard output, ending the output for each file at point.						
19384 19385			ll end at the point in each input file indicated by the $-\mathbf{n}$ <i>number</i> option. The option- <i>mber</i> shall be counted in units of lines.						
19386	OPTION	I S							
19387 19388			ity shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section Syntax Guidelines.						
19389		The following	g option shall be supported:						
19390 19391 19392		– n number	The first <i>number</i> lines of each input file shall be copied to standard output. The application shall ensure that the <i>number</i> option-argument is a positive decimal integer.						
19393 19394		When a file o This shall not	contains less than <i>number</i> lines, it shall be copied to standard output in its entirety. t be an error.						
19395		If no options	are specified, <i>head</i> shall act as if $-\mathbf{n}$ 10 had been specified.						
19396 19397	OPERAN		g operand shall be supported:						
19398 19399		file	A pathname of an input file. If no <i>file</i> operands are specified, the standard input shall be used.						
19400	STDIN								
19401 19402		The standard section.	d input shall be used only if no file operands are specified. See the INPUT FILES						
19403	INPUT F								
19404		Input files sh	all be text files, but the line length is not restricted to {LINE_MAX} bytes.						
19405 19406		NMENT VA The following	RIABLES g environment variables shall affect the execution of <i>head</i> :						
19407 19408 19409 19410		LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)						
19411 19412		LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.						
19413 19414 19415		LC_CTYPE	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).						

diagnostic messages written to standard error.

Determine the locale that should be used to affect the format and contents of

LC_MESSAGES

19416

19417

19418

head Utilities

19419 XSI **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 19420 ASYNCHRONOUS EVENTS Default. 19422 STDOUT 19423 The standard output shall contain designated portions of the input files. If multiple file operands are specified, head shall precede the output for each with the header: 19424 $"\n==> %s <==\n", <pathname>$ 19425 except that the first header written shall not include the initial <newline>. 19426 19427 STDERR 19428 The standard error shall be used only for diagnostic messages. 19429 OUTPUT FILES None. 19430 19431 EXTENDED DESCRIPTION None. 19433 EXIT STATUS The following exit values shall be returned: 19434 Successful completion. 19435 19436 >0 An error occurred. 19437 CONSEQUENCES OF ERRORS 19438 Default. 19439 APPLICATION USAGE 19440 The obsolescent -number form is withdrawn in this version. Applications should use the -n19441 number option. 19442 EXAMPLES To write the first ten lines of all files (except those with a leading period) in the directory: 19443 19444 head * 19445 RATIONALE 19446 Although it is possible to simulate head with sed 10q for a single file, the standard developers decided that the popularity of *head* on historical BSD systems warranted its inclusion alongside 19447 tail. 19448 This standard version of head follows the Utility Syntax Guidelines. The -n option was added to 19449 this new interface so that *head* and *tail* would be more logically related. 19450 There is no -c option (as there is in tail) because it is not historical practice and because other 19451 utilities in this volume of IEEE Std 1003.1-2001 provide similar functionality. 19452 19453 FUTURE DIRECTIONS None. 19454 19455 SEE ALSO sed. tail 19456

Utilities head

19457 CHANGE HISTORY

19458 First released in Issue 4.

19459 **Issue 6**

19460 The obsolescent **–number** form is withdrawn.

The normative text is reworded to avoid use of the term "must" for application requirements.

The DESCRIPTION is updated to clarify that when a file contains less than the number of lines

requested, the entire file is copied to standard output.

iconv Utilities

19464 NAME 19465	icony cod	eset conversi	on	
		eset conversi	JII	
19466 SYNOI 19467		s] -f from	map -t tomap [file]	1
19468	iconv -f	fromcode [-cs] [-t tocode [file]	ı
19469	iconv -t	tocode [-c	s] [-f fromcode] [file]	İ
19470	iconv -l			1
19471 DESCR	RIPTION			
19472 19473		ility shall con ults to standa	vert the encoding of characters in <i>file</i> from one codeset to another and ord output.	
19474 19475 19476 19477	the codeset character na	conversion s mes in the tw	te that charmap files are used to specify the codesets (see OPTIONS), shall be accomplished by performing a logical join on the symbolic o charmaps. The implementation need not support the use of charmap on unless the POSIX2_LOCALEDEF symbol is defined on the system.	
19478 OPTIO	NS			
19479 19480		ility shall con Syntax Guide	nform to the Base Definitions volume of IEEE Std 1003.1-2001, Section lines.	
19481	The following	ng options sha	all be supported:	
19482 19483 19484 19485 19486 19487	-с	output. Who input stream that have n specified in	haracters that are invalid in the codeset of the input file from the en $-c$ is not used, the results of encountering invalid characters in the either those that are not characters in the codeset of the input file or corresponding character in the codeset of the output file) shall be the system documentation. The presence or absence of $-c$ shall not it status of <i>iconv</i> .	
19488	- f fromcodese	et		-
19489 19490		•	e codeset of the input file. The implementation shall recognize the vo forms of the <i>fromcodeset</i> option-argument:	
19491 19492 19493 19494		fromcode	The <i>fromcode</i> option-argument must not contain a slash character. It shall be interpreted as the name of one of the codeset descriptions provided by the implementation in an unspecified format. Valid values of <i>fromcode</i> are implementation-defined.	
19495 19496 19497 19498 19499		frommap	The <i>frommap</i> option-argument must contain a slash character. It shall be interpreted as the pathname of a charmap file as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.4, Character Set Description File. If the pathname does not represent a valid, readable charmap file, the results are undefined.	
19500		If this option	n is omitted, the codeset of the current locale shall be used.	
19501 19502	- l	Write all sup format.	oported fromcode and tocode values to standard output in an unspecified	
19503 19504 19505 19506 19507	−s	When -s is stream (eith that have n specified in	ny messages written to standard error concerning invalid characters. not used, the results of encountering invalid characters in the input er those that are not valid characters in the codeset of the input file or o corresponding character in the codeset of the output file) shall be the system documentation. The presence or absence of —s shall not it status of input	

affect the exit status of *iconv*.

19508

Utilities iconv

19509 19510	-t tocodeset	Identify the codeset to be used for the output file. The implementation shall recognize the following two forms of the <i>tocodeset</i> option-argument:
19511		tocode The semantics shall be equivalent to the -f fromcode option.
19512		tomap The semantics shall be equivalent to the tomap option.
19513		If this option is omitted, the codeset of the current locale shall be used.
19514	If either − f o	or -t represents a charmap file, but the other does not (or is omitted), or both -f and
19515	−t are omitte	ed, the results are undefined.
19516 OPER 19517		ng operand shall be supported:
19518 19519	file	A pathname of an input file. If no <i>file</i> operands are specified, or if a <i>file</i> operand is $'-'$, the standard input shall be used.
19520 STDI	N	
19521	The standar	d input shall be used only if no <i>file</i> operands are specified, or if a <i>file</i> operand is $'-'$.
19522 INPU 19523		le shall be a text file.
19524 ENVI	RONMENT VA	
19525		ng environment variables shall affect the execution of <i>iconv</i> :
19526 19527	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
19528		Internationalization Variables for the precedence of internationalization variables
19529		used to determine the values of locale categories.)
19530 19531	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
19532	LC_CTYPE	Determine the locale for the interpretation of sequences of bytes of text data as
19533 19534		characters (for example, single-byte as opposed to multi-byte characters in arguments). During translation of the file, this variable is superseded by the use of
19535		the fromcode option-argument.
19536	LC_MESSA	GES
19537 19538		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
19539 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .
	CHRONOUS	
19541	Default.	EVENIS
19542 STDO 19543		l option is used, the standard output shall contain all supported fromcode and tocode
19545		ten in an unspecified format.
19545	When the –	l option is not used, the standard output shall contain the sequence of characters
19546	read from th	ne input files, translated to the specified codeset. Nothing else shall be written to the
19547	standard ou	tput.
19548 STDE	KK	damen al all harvard and of Conditions and a

507

The standard error shall be used only for diagnostic messages.

19549

iconv Utilities

19550 OUTPUT FILES

19551 None.

19552 EXTENDED DESCRIPTION

19553 None.

19554 EXIT STATUS

19555 The following exit values shall be returned:

19556 0 Successful completion.

19557 >0 An error occurred.

19558 CONSEQUENCES OF ERRORS

19559 Default.

19560 APPLICATION USAGE

The user must ensure that both charmap files use the same symbolic names for characters the two codesets have in common.

19563 EXAMPLES

The following example converts the contents of file mail.x400 from the ISO/IEC 6937:1994 standard codeset to the ISO/IEC 8859-1:1998 standard codeset, and stores the results in file mail.local:

19567 iconv -f IS6937 -t IS8859 mail.x400 > mail.local

19568 RATIONALE

19569

19570

19571

19572 19573

19574

19575

19576 19577 The *iconv* utility can be used portably only when the user provides two charmap files as optionarguments. This is because a single charmap provided by the user cannot reliably be joined with the names in a system-provided character set description. The valid values for *fromcode* and *tocode* are implementation-defined and do not have to have any relation to the charmap mechanisms. As an aid to interactive users, the –l option was adopted from the Plan 9 operating system. It writes information concerning these implementation-defined values. The format is unspecified because there are many possible useful formats that could be chosen, such as a matrix of valid combinations of *fromcode* and *tocode*. The –l option is not intended for shell script usage; conforming applications will have to use charmaps.

19578 FUTURE DIRECTIONS

19579 None.

19580 SEE ALSO

19581 gencat

19582 CHANGE HISTORY

19583 First released in Issue 3.

19584 Issue 6

This utility has been rewritten to align with the IEEE P1003.2b draft standard. Specifically, the ability to use charmap files for conversion has been added.

IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/29 is applied, making changes to address inconsistencies with the *iconv*() function in the System Interfaces volume of IEEE Std 1003.1-2001.

id **Utilities**

19590 **NAME** 19591 id — return user identity 19592 SYNOPSIS id [user] 19593 19594 id -G[-n] [user] id -g[-nr] [user] 19595 id -u[-nr] [user] 19596 19597 DESCRIPTION If no user operand is provided, the id utility shall write the user and group IDs and the 19598 corresponding user and group names of the invoking process to standard output. If the effective 19599 and real IDs do not match, both shall be written. If multiple groups are supported by the 19600 underlying system (see the description of {NGROUPS MAX} in the System Interfaces volume of 19601 IEEE Std 1003.1-2001), the supplementary group affiliations of the invoking process shall also be 19602 19603 written. If a user operand is provided and the process has the appropriate privileges, the user and group 19604 IDs of the selected user shall be written. In this case, effective IDs shall be assumed to be 19605 identical to real IDs. If the selected user has more than one allowable group membership listed 19606 in the group database, these shall be written in the same manner as the supplementary groups 19607 described in the preceding paragraph. 19608 19609 OPTIONS The *id* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, 19610 Utility Syntax Guidelines. 19611 19612 The following options shall be supported: 19613 $-\mathbf{G}$ Output all different group IDs (effective, real, and supplementary) only, using the format "%u\n". If there is more than one distinct group affiliation, output each 19614 such affiliation, using the format " %u", before the <newline> is output. 19615 Output only the effective group ID, using the format " $u\n$ ". 19616 -g 19617 -n Output the name in the format "%s" instead of the numeric ID using the format "%u". 19618 19619 Output the real ID instead of the effective ID. $-\mathbf{r}$ 19620 -u Output only the effective user ID, using the format " $u\n$ ". 19621 OPERANDS The following operand shall be supported: 19622 19623 user The login name for which information is to be written. 19624 STDIN Not used. 19625 19626 INPUT FILES None. 19627 19628 ENVIRONMENT VARIABLES 19629

The following environment variables shall affect the execution of *id*:

19630	LANG	Provi	de a c	lefault	value for the	internatio	naliz	ation variables that are	unset or	null.
19631		(See	the	Base	Definitions	volume	of	IEEE Std 1003.1-2001,	Section	8.2,
19632		Interr	ation	alizatio	on Variables i	for the pre	eced	ence of internationaliza	tion varia	bles

id Utilities

19633		used to determine the values of locale categories.)						
19634 19635	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.						
19636 19637 19638	LC_CTYPE	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).						
19639 19640 19641 19642	LC_MESSAC	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.						
19643 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.						
19644 ASYNC 19645	HRONOUS I Default.	EVENTS						
19646 STDOU 19647 19648 19649	The following POSIX locale	ng formats shall be used when the <i>LC_MESSAGES</i> locale category specifies the e. In other locales, the strings <i>uid</i> , <i>gid</i> , <i>euid</i> , <i>egid</i> , and <i>groups</i> may be replaced with briate strings corresponding to the locale.						
19650 19651		s) gid=%u(%s)\n", <real id="" user="">, <user-name>, group ID>, <group-name></group-name></user-name></real>						
19652 19653		ve and real user IDs do not match, the following shall be inserted immediately n' character in the previous format:						
19654	" euid=%u	(%s)"						
19655	with the follo	owing arguments added at the end of the argument list:						
19656	<effective< td=""><td colspan="7"><effective id="" user="">, <effective user-name=""></effective></effective></td></effective<>	<effective id="" user="">, <effective user-name=""></effective></effective>						
19657 19658 19659	the '\n' cha	If the effective and real group IDs do not match, the following shall be inserted directly before the ' \n' character in the format string (and after any addition resulting from the effective and real user IDs not matching):						
19660	" egid=%u	(%s)"						
19661	with the follo	owing arguments added at the end of the argument list:						
19662	<effective< td=""><td>e group-ID>, <effective group="" name=""></effective></td></effective<>	e group-ID>, <effective group="" name=""></effective>						
19663 19664		s has supplementary group affiliations or the selected user is allowed to belong to ups, the first shall be added directly before the <newline> in the format string:</newline>						
19665	" groups=	%u (%s) "						
19666	with the follo	owing arguments added at the end of the argument list:						
19667	<supplement< td=""><td>ntary group ID>, <supplementary group="" name=""></supplementary></td></supplement<>	ntary group ID>, <supplementary group="" name=""></supplementary>						
19668 19669	and the nece group IDs:	essary number of the following added after that for any remaining supplementary						
19670	",%u(%s)"							
19671	and the nece	ssary number of the following arguments added at the end of the argument list:						
19672	<supplement< td=""><td>ntary group ID>, <supplementary group="" name=""></supplementary></td></supplement<>	ntary group ID>, <supplementary group="" name=""></supplementary>						

Utilities id

If any of the user ID, group ID, effective user ID, effective group ID, or supplementary/multiple group IDs cannot be mapped by the system into printable user or group names, the corresponding "(%s)" and *name* argument shall be omitted from the corresponding format string.

When any of the options are specified, the output format shall be as described in the OPTIONS section.

19679 STDERR

19680 The standard error shall be used only for diagnostic messages.

19681 **OUTPUT FILES** 19682 None.

19683 EXTENDED DESCRIPTION

19684 None.

19685 EXIT STATUS

19688

19692

19693

19694

19695 19696

19697

19698

19699

19686 The following exit values shall be returned:

>0 An error occurred.

19687 0 Successful completion.

19689 CONSEQUENCES OF ERRORS

19690 Default.

19691 APPLICATION USAGE

Output produced by the $-\mathbf{G}$ option and by the default case could potentially produce very long lines on systems that support large numbers of supplementary groups. (On systems with user and group IDs that are 32-bit integers and with group names with a maximum of 8 bytes per name, 93 supplementary groups plus distinct effective and real group and user IDs could theoretically overflow the 2 048-byte {LINE_MAX} text file line limit on the default output case. It would take about 186 supplementary groups to overflow the 2 048-byte barrier using id $-\mathbf{G}$). This is not expected to be a problem in practice, but in cases where it is a concern, applications should consider using fold $-\mathbf{s}$ before postprocessing the output of id.

19700 EXAMPLES

19701 None.

19702 RATIONALE

The functionality provided by the 4 BSD *groups* utility can be simulated using:

```
19704 id -Gn [ user ]
```

The 4 BSD command *groups* was considered, but it was not included because it did not provide the functionality of the *id* utility of the SVID. Also, it was thought that it would be easier to modify *id* to provide the additional functionality necessary to systems with multiple groups than to invent another command.

The options **–u**, **–g**, **–n**, and **–r** were added to ease the use of *id* with shell commands substitution. Without these options it is necessary to use some preprocessor such as *sed* to select the desired piece of information. Since output such as that produced by:

19712 id -u -n

is frequently wanted, it seemed desirable to add the options.

id Utilities

19714 **FUTURE DIRECTIONS**

19715 None.

19716 **SEE ALSO**

19717 fold, logname, who, the System Interfaces volume of IEEE Std 1003.1-2001, getgid(), getgroups(),

19718 *getuid()*

19719 CHANGE HISTORY

First released in Issue 2.

Utilities ipcrm

19721 **NAME** 19722 ipcrm — remove an XSI message queue, semaphore set, or shared memory segment identifier 19723 SYNOPSIS ipcrm [-q msgid | -Q msgkey | -s semid | -S semkey | 19724 XSI 19725 -m shmid | -M shmkey] ... 19726 19727 **DESCRIPTION** The *ipcrm* utility shall remove zero or more message queues, semaphore sets, or shared memory 19728 19729 segments. The interprocess communication facilities to be removed are specified by the options. 19730 Only a user with appropriate privilege shall be allowed to remove an interprocess communication facility that was not created by or owned by the user invoking *ipcrm*. 19731 19732 OPTIONS The *ipcrm* facility supports the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, 19733 19734 Utility Syntax Guidelines. 19735 The following options shall be supported: Remove the message queue identifier msgid from the system and destroy the -q msgid 19736 message queue and data structure associated with it. 19737 -m shmid Remove the shared memory identifier *shmid* from the system. The shared memory 19738 segment and data structure associated with it shall be destroyed after the last 19739 19740 detach. 19741 −s semid Remove the semaphore identifier *semid* from the system and destroy the set of 19742 semaphores and data structure associated with it. Remove the message queue identifier, created with key msgkey, from the system 19743 −**Q** msgkey 19744 and destroy the message queue and data structure associated with it. Remove the shared memory identifier, created with key *shmkey*, from the system. 19745 -M shmkey 19746 The shared memory segment and data structure associated with it shall be destroyed after the last detach. 19747 19748 -S semkey Remove the semaphore identifier, created with key semkey, from the system and destroy the set of semaphores and data structure associated with it. 19749 19750 **OPERANDS** None. 19751 19752 **STDIN** Not used. 19753 19754 INPUT FILES None. 19755 19756 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *ipcrm*: 19757 LANG 19758 Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 19759 Internationalization Variables for the precedence of internationalization variables 19760 used to determine the values of locale categories.) 19761 LC ALL If set to a non-empty string value, override the values of all the other 19762

internationalization variables.

19763

ipcrm Utilities

19764 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 19765 characters (for example, single-byte as opposed to multi-byte characters in arguments). 19766 LC_MESSAGES 19767 Determine the locale that should be used to affect the format and contents of 19768 diagnostic messages written to standard error. 19769 **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 19770 19771 ASYNCHRONOUS EVENTS Default. 19772 19773 STDOUT Not used. 19774 19775 STDERR The standard error shall be used only for diagnostic messages. 19776 19777 OUTPUT FILES None. 19778 19779 EXTENDED DESCRIPTION None. 19780 19781 EXIT STATUS The following exit values shall be returned: 19782 0 Successful completion. 19783 >0 An error occurred. 19784 19785 CONSEQUENCES OF ERRORS 19786 Default. 19787 APPLICATION USAGE 19788 None. 19789 EXAMPLES 19790 None. 19791 RATIONALE None. 19792 19793 FUTURE DIRECTIONS 19794 None. 19795 SEE ALSO ipcs, the System Interfaces volume of IEEE Std 1003.1-2001, msgctl(), semctl(), shmctl() 19796 19797 CHANGE HISTORY

19798

First released in Issue 5.

Utilities ipcs

19799 **NAME** ipcs — report XSI interprocess communication facilities status 19800 19801 SYNOPSIS 19802 XSI ipcs [-qms] [-a | -bcopt] 19803 19804 DESCRIPTION The *ipcs* utility shall write information about active interprocess communication facilities. 19805 Without options, information shall be written in short format for message queues, shared 19806 memory segments, and semaphore sets that are currently active in the system. Otherwise, the 19807 information that is displayed is controlled by the options specified. 19808 19809 OPTIONS 19810 The *ipcs* facility supports the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines. 19811 19812 The *ipcs* utility accepts the following options: Write information about active message queues. 19813 -q 19814 $-\mathbf{m}$ Write information about active shared memory segments. Write information about active semaphore sets. 19815 -sIf $-\mathbf{q}$, $-\mathbf{m}$, or $-\mathbf{s}$ are specified, only information about those facilities shall be written. If none of 19816 19817 these three are specified, information about all three shall be written subject to the following options: 19818 19819 Use all print options. (This is a shorthand notation for $-\mathbf{b}$, $-\mathbf{c}$, $-\mathbf{o}$, $-\mathbf{p}$, and $-\mathbf{t}$.) -a 19820 -b Write information on maximum allowable size. (Maximum number of bytes in 19821 messages on queue for message queues, size of segments for shared memory, and number of semaphores in each set for semaphores.) 19822 19823 **-с** Write creator's user name and group name; see below. 19824 -0 Write information on outstanding usage. (Number of messages on queue and total number of bytes in messages on queue for message queues, and number of 19825 processes attached to shared memory segments.) 19826 19827 Write process number information. (Process ID of the last process to send a -p message and process ID of the last process to receive a message on message 19828 queues, process ID of the creating process, and process ID of the last process to 19829 attach or detach on shared memory segments.) 19830 $-\mathbf{t}$ Write time information. (Time of the last control operation that changed the access 19831 permissions for all facilities, time of the last *msgsnd()* and *msgrcv()* operations on 19832 message queues, time of the last shmat() and shmdt() operations on shared 19833 memory, and time of the last *semop()* operation on semaphores.) 19834

19835 **OPERANDS**

19836 None.

19837 **STDIN**

19838 Not used.

ipcs Utilities

19839 INPUT FILES

The group databaseThe user database

19842 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *ipcs*:

Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)

19848 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

19850 *LC_CTYPE* 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).

19853 LC_MESSAGES

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

19856 NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

Determine the timezone for the date and time strings written by *ipcs*. If *TZ* is unset or null, an unspecified default timezone shall be used.

19859 ASYNCHRONOUS EVENTS

19860 Default.

19861 **STDOUT**

19862 An introductory line shall be written with the format:

19863 "IPC status from %s as of %s\n", <source>, <date>

where *<source>* indicates the source used to gather the statistics and *<date>* is the information that would be produced by the *date* command when invoked in the POSIX locale.

The *ipcs* utility then shall create up to three reports depending upon the **-q**, **-m**, and **-s** options.

The first report shall indicate the status of message queues, the second report shall indicate the status of shared memory segments, and the third report shall indicate the status of semaphore sets.

19870 If the corresponding facility is not installed or has not been used since the last reboot, then the report shall be written out in the format:

19872 "%s facility not in system.\n", < facility>

where < facility> is Message Queue, Shared Memory, or Semaphore, as appropriate. If the facility has been installed and has been used since the last reboot, column headings separated by one or more spaces and followed by a <newline> shall be written as indicated below followed by the facility name written out using the format:

19877 "%s:\n", < facility>

where *<facility>* is *Message Queues*, *Shared Memory*, or *Semaphores*, as appropriate. On the second and third reports the column headings need not be written if the last column headings written already provide column headings for all information in that report.

Utilities ipcs

198 198 198 198 198	82 83 84 85	The column headings provided in the first column below and the meaning of the information in those columns shall be given in order below; the letters in parentheses indicate the options that shall cause the corresponding column to appear; "all" means that the column shall always appear. Each column is separated by one or more <space>s. Note that these options only determine what information is provided for each report; they do not determine which reports are written.</space>					
198	887	T	(all)	Type of	facility:		
198	888			q	Message queue.		
198	89			m	Shared memory segment.		
198	90			S	Semaphore.		
198	91			This fiel	d is a single character written using the format %c.		
198 198		ID	(all)	The ider	ntifier for the facility entry. This field shall be written using the format		
198 198		KEY	(all)	The key facility e	used as an argument to $msgget()$, $semget()$, or $shmget()$ to create the entry.		
198 198 198	97			Note:	The key of a shared memory segment is changed to IPC_PRIVATE when the segment has been removed until all processes attached to the segment detach it.		
198	99			This fiel	d shall be written using the format 0x%x.		
199 199		MODE	(all)		lity access modes and flags. The mode shall consist of 11 characters interpreted as follows.		
199	02			The first	character shall be:		
199	03			S	If a process is waiting on a <i>msgsnd()</i> operation.		
199	04			_	If the above is not true.		
199	05			The seco	ond character shall be:		
199	06			R	If a process is waiting on a <i>msgrcv()</i> operation.		
199 199				Cor-	If the associated shared memory segment is to be cleared when the first attach operation is executed.		
199	09			_	If none of the above is true.		
199 199 199 199 199	11 12 13 14			The firs others in each set indicates	t nine characters shall be interpreted as three sets of three bits each. It set refers to the owner's permissions; the next to permissions of a the usergroup of the facility entry; and the last to all others. Within the first character indicates permission to read, the second character is permission to write or alter the facility entry, and the last character is sign $('-')$.		
199	16			The peri	missions shall be indicated as follows:		
199	17			r	If read permission is granted.		
199	18			W	If write permission is granted.		
199	19			a	If alter permission is granted.		
199	20			_	If the indicated permission is not granted.		

ipcs Utilities

19921 19922 19923 19924 19925			The first character following the permissions specifies if there is an alternate or additional access control method associated with the facility. If there is no alternate or additional access control method associated with the facility, a single <space> shall be written; otherwise, another printable character is written.</space>
19926 19927 19928 19929	OWNER	(all)	The user name of the owner of the facility entry. If the user name of the owner is found in the user database, at least the first eight column positions of the name shall be written using the format %s. Otherwise, the user ID of the owner shall be written using the format %d.
19930 19931 19932 19933	GROUP	(all)	The group name of the owner of the facility entry. If the group name of the owner is found in the group database, at least the first eight column positions of the name shall be written using the format %s. Otherwise, the group ID of the owner shall be written using the format %d.
19934	The follow	ving niı	ne columns shall be only written out for message queues:
19935 19936 19937 19938	CREATOR	? (a,c)	The user name of the creator of the facility entry. If the user name of the creator is found in the user database, at least the first eight column positions of the name shall be written using the format %s. Otherwise, the user ID of the creator shall be written using the format %d.
19939 19940 19941 19942	CGROUP	(a,c)	The group name of the creator of the facility entry. If the group name of the creator is found in the group database, at least the first eight column positions of the name shall be written using the format %s. Otherwise, the group ID of the creator shall be written using the format %d.
19943 19944	CBYTES	(a,o)	The number of bytes in messages currently outstanding on the associated message queue. This field shall be written using the format %d.
19945 19946	QNUM	(a,o)	The number of messages currently outstanding on the associated message queue. This field shall be written using the format %d.
19947 19948	QBYTES	(a,b)	The maximum number of bytes allowed in messages outstanding on the associated message queue. This field shall be written using the format %d.
19949 19950	LSPID	(a,p)	The process ID of the last process to send a message to the associated queue. This field shall be written using the format:
19951			"%d", <pid></pid>
19952 19953 19954			where $< pid >$ is 0 if no message has been sent to the corresponding message queue; otherwise, $< pid >$ shall be the process ID of the last process to send a message to the queue.
19955 19956	LRPID	(a,p)	The process ID of the last process to receive a message from the associated queue. This field shall be written using the format:
19957			"%d", <pid></pid>
19958 19959 19960			where $<\!pid\!>$ is 0 if no message has been received from the corresponding message queue; otherwise, $<\!pid\!>$ shall be the process ID of the last process to receive a message from the queue.
19961 19962 19963 19964 19965	STIME	(a,t)	The time the last message was sent to the associated queue. If a message has been sent to the corresponding message queue, the hour, minute, and second of the last time a message was sent to the queue shall be written using the format %d:%2.2d:%2.2d. Otherwise, the format "no-entry" shall be written.

Utilities ipcs

19966 19967 19968 19969 19970	RTIME	(a,t)	The time the last message was received from the associated queue. If a message has been received from the corresponding message queue, the hour, minute, and second of the last time a message was received from the queue shall be written using the format %d:%2.2d:%2.2d. Otherwise, the format "no-entry" shall be written.
19971	The follow	ving eig	tht columns shall be only written out for shared memory segments.
19972 19973 19974 19975	CREATOR	₹ (a,c)	The user of the creator of the facility entry. If the user name of the creator is found in the user database, at least the first eight column positions of the name shall be written using the format %s. Otherwise, the user ID of the creator shall be written using the format %d.
19976 19977 19978 19979	CGROUP	(a,c)	The group name of the creator of the facility entry. If the group name of the creator is found in the group database, at least the first eight column positions of the name shall be written using the format %s. Otherwise, the group ID of the creator shall be written using the format %d.
19980 19981	NATTCH	(a,o)	The number of processes attached to the associated shared memory segment. This field shall be written using the format $\&d$.
19982 19983	SEGSZ	(a , b)	The size of the associated shared memory segment. This field shall be written using the format $\$ \mbox{d}.$
19984 19985	CPID	(a,p)	The process ID of the creator of the shared memory entry. This field shall be written using the format %d.
19986 19987	LPID	(a,p)	The process ID of the last process to attach or detach the shared memory segment. This field shall be written using the format:
19988			"%d", <pid></pid>
19989 19990 19991			where $<\!pid\!>$ is 0 if no process has attached the corresponding shared memory segment; otherwise, $<\!pid\!>$ shall be the process ID of the last process to attach or detach the segment.
19992 19993 19994 19995 19996	ATIME	(a,t)	The time the last attach on the associated shared memory segment was completed. If the corresponding shared memory segment has ever been attached, the hour, minute, and second of the last time the segment was attached shall be written using the format %d:%2.2d:%2.2d. Otherwise, the format " no-entry" shall be written.
19997 19998 19999 20000 20001	DTIME	(a,t)	The time the last detach on the associated shared memory segment was completed. If the corresponding shared memory segment has ever been detached, the hour, minute, and second of the last time the segment was detached shall be written using the format %d:%2.2d:%2.2d. Otherwise, the format " no-entry" shall be written.
20002	The follow	ving fou	ur columns shall be only written out for semaphore sets:
20003 20004 20005 20006	CREATOR	R (a,c)	The user of the creator of the facility entry. If the user name of the creator is found in the user database, at least the first eight column positions of the name shall be written using the format %s. Otherwise, the user ID of the creator shall be written using the format %d.
20007 20008 20009 20010	CGROUP	(a,c)	The group name of the creator of the facility entry. If the group name of the creator is found in the group database, at least the first eight column positions of the name shall be written using the format $\$s$. Otherwise, the group ID of the creator shall be written using the format $\$d$.

ipcs Utilities

20011 20012	NSEMS	(a , b)	The number of semaphores in the set associated with the semaphore entry. This field shall be written using the format %d.
20013 20014 20015 20016 20017 20018	OTIME	(a,t)	The time the last semaphore operation on the set associated with the semaphore entry was completed. If a semaphore operation has ever been performed on the corresponding semaphore set, the hour, minute, and second of the last semaphore operation on the semaphore set shall be written using the format %d:%2.2d:%2.2d. Otherwise, the format "no-entry" shall be written.
20019	The follow	wing co	lumn shall be written for all three reports when it is requested:
20020 20021 20022	CTIME	(a,t)	The time the associated entry was created or changed. The hour, minute, and second of the time when the associated entry was created shall be written using the format %d:%2.2d:%2.2d.
20023 STDE I			
20024		lard err	or shall be used only for diagnostic messages.
20025 OUTP 20026	UT FILES None.		
20027 EXTEN 20028		CRIPTI	ON
20029 EXIT S	TATUS		
20030	The follow	wing ex	it values shall be returned:
20031	0 Succ	essful c	ompletion.
20032	>0 An e	rror occ	urred.
20033 CONS 20034	EQUENCES Default.	S OF EF	RRORS
20035 APPLI 20036 20037	Things ca	ın chan	ge while <i>ipcs</i> is running; the information it gives is guaranteed to be accurate retrieved.
20038 EXAM			
20039	None.		
20040 RATIC 20041	NALE None.		
20042 FUTU l 20043	RE DIRECT None.	IONS	
20044 SEE A l 20045 20046			rfaces volume of IEEE Std 1003.1-2001, msgrcv(), msgsnd(), semget(), semop(), shmget()
20047 CHAN 20048	GE HISTO First relea		ssue 5.
20049 Issue 6 20050		ı Group	Corrigendum U020/1 is applied, correcting the SYNOPSIS.
20051	The Oper	Group	Corrigenda $U032/1$ and $U032/2$ are applied, clarifying the output format.
20052	The Oper	Group	Base Resolution bwg98-004 is applied.

jobs **Utilities**

20053 **NAME**

20054 jobs — display status of jobs in the current session

20055 SYNOPSIS

jobs [-1 | -p] [job id...] 20056 UP

20057

20058 DESCRIPTION

The *jobs* utility shall display the status of jobs that were started in the current shell environment; 20059

see Section 2.12 (on page 61). 20060

When jobs reports the termination status of a job, the shell shall remove its process ID from the 20061 list of those "known in the current shell execution environment"; see Section 2.9.3.1 (on page 20062

50). 20063

20064 OPTIONS

The jobs utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 20065 12.2, Utility Syntax Guidelines. 20066

The following options shall be supported: 20067

 $-\mathbf{l}$ (The letter ell.) Provide more information about each job listed. This information 20068 shall include the job number, current job, process group ID, state, and the 20069 command that formed the job. 20070

Display only the process IDs for the process group leaders of the selected jobs. 20071 -p

20072 By default, the jobs utility shall display the status of all stopped jobs, running background jobs and all jobs whose status has changed and have not been reported by the shell. 20073

20074 OPERANDS

20075 The following operand shall be supported:

job_id Specifies the jobs for which the status is to be displayed. If no job id is given, the 20076 status information for all jobs shall be displayed. The format of *job_id* is described 20077 20078 in the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.203, Job Control

Job ID.

20080 STDIN

20079

Not used. 20081

20082 INPUT FILES

None. 20083

20084 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *jobs*: 20085

LANG Provide a default value for the internationalization variables that are unset or null. 20086 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 20087 Internationalization Variables for the precedence of internationalization variables 20088 20089

used to determine the values of locale categories.)

LC ALL If set to a non-empty string value, override the values of all the other 20090 internationalization variables. 20091

Determine the locale for the interpretation of sequences of bytes of text data as 20092 LC_CTYPE 20093

characters (for example, single-byte as opposed to multi-byte characters in

arguments). 20094

LC_MESSAGES 20095

Determine the locale that should be used to affect the format and contents of 20096

jobs Utilities

20097 diagnostic messages written to standard error and informative messages written to 20098 standard output. 20099 XSI NLSPATH Determine the location of message catalogs for the processing of *LC MESSAGES*. 20100 ASYNCHRONOUS EVENTS 20101 Default. 20102 STDOUT 20103 If the $-\mathbf{p}$ option is specified, the output shall consist of one line for each process ID: 20104 "%d\n", cess ID> Otherwise, if the **-l** option is not specified, the output shall be a series of lines of the form: 20105 20106 "[%d] %c %s %sn", <job-number>, <current>, <state>, <command> where the fields shall be as follows: 20107 <current> The character '+' identifies the job that would be used as a default for the fg or bg 20108 utilities; this job can also be specified using the job_id %+ or "%%". The character 20109 '-' identifies the job that would become the default if the current default job were 20110 to exit; this job can also be specified using the job_id %-. For other jobs, this field is 20111 a <space>. At most one job can be identified with '+' and at most one job can be 20112 identified with '-'. If there is any suspended job, then the current job shall be a 20113 suspended job. If there are at least two suspended jobs, then the previous job also 20114 shall be a suspended job. 20115 <job-number> A number that can be used to identify the process group to the wait, fg, bg, and kill 20116 utilities. Using these utilities, the job can be identified by prefixing the job number 20117 with '%'. 20118 20119 <state> One of the following strings (in the POSIX locale): 20120 Running Indicates that the job has not been suspended by a signal and has not exited. 20121 Done Indicates that the job completed and returned exit status zero. 20122 20123 Done(code) Indicates that the job completed normally and that it exited with the specified non-zero exit status, *code*, expressed as a decimal number. 20124 20125 Stopped Indicates that the job was suspended by the SIGTSTP signal. Stopped (SIGTSTP) 20126 Indicates that the job was suspended by the SIGTSTP signal. 20127 Stopped (SIGSTOP) 20128 Indicates that the job was suspended by the SIGSTOP signal. 20129 Stopped (SIGTTIN) 20130 20131 Indicates that the job was suspended by the SIGTTIN signal. Stopped (SIGTTOU) 20132 20133 Indicates that the job was suspended by the SIGTTOU signal. The implementation may substitute the string **Suspended** in place of **Stopped**. If 20134 the job was terminated by a signal, the format of <state> is unspecified, but it shall 20135 be visibly distinct from all of the other *<state>* formats shown here and shall 20136 indicate the name or description of the signal causing the termination. 20137

Utilities jobs

20138 < command> The associated command that was given to the shell.

If the **–l** option is specified, a field containing the process group ID shall be inserted before the <state> field. Also, more processes in a process group may be output on separate lines, using only the process ID and <command> fields.

.

20142 **STDERR**

20143

The standard error shall be used only for diagnostic messages.

20144 OUTPUT FILES

20145 None.

20146 EXTENDED DESCRIPTION

20147 None.

20148 EXIT STATUS

20149 The following exit values shall be returned:

20150 0 Successful completion.

20151 >0 An error occurred.

20152 CONSEQUENCES OF ERRORS

20153 Default.

20154 APPLICATION USAGE

The $-\mathbf{p}$ option is the only portable way to find out the process group of a job because different implementations have different strategies for defining the process group of the job. Usage such as $\$(jobs - \mathbf{p})$ provides a way of referring to the process group of the job in an implementation-independent way.

The *jobs* utility does not work as expected when it is operating in its own utility execution environment because that environment has no applicable jobs to manipulate. See the APPLICATION USAGE section for *bg*. For this reason, *jobs* is generally implemented as a shell regular built-in.

20163 EXAMPLES

20164 None.

20165 RATIONALE

Both "%%" and "%+" are used to refer to the current job. Both forms are of equal validity—the
"%%" mirroring "\$\$" and "%+" mirroring the output of *jobs*. Both forms reflect historical
practice of the KornShell and the C shell with job control.

The job control features provided by *bg, fg,* and *jobs* are based on the KornShell. The standard developers examined the characteristics of the C shell versions of these utilities and found that differences exist. Despite widespread use of the C shell, the KornShell versions were selected for this volume of IEEE Std 1003.1-2001 to maintain a degree of uniformity with the rest of the KornShell features selected (such as the very popular command line editing features).

The *jobs* utility is not dependent on the job control option, as are the seemingly related *bg* and *fg* utilities because *jobs* is useful for examining background jobs, regardless of the condition of job control. When the user has invoked a *set* +**m** command and job control has been turned off, *jobs* can still be used to examine the background jobs associated with that current session. Similarly, *kill* can then be used to kill background jobs with *kill*% <*background job number*>.

The output for terminated jobs is left unspecified to accommodate various historical systems.
The following formats have been witnessed:

jobs Utilities

20181	1. Killed(signal name)
20182	2. signal name
20183	3. signal name(coredump)
20184	4. signal description—core dumped
20185 20186	Most users should be able to understand these formats, although it means that applications have trouble parsing them.
20187 20188	The calculation of job IDs was not described since this would suggest an implementation, which may impose unnecessary restrictions.
20189 20190 20191	In an early proposal, a $-\mathbf{n}$ option was included to "Display the status of jobs that have changed, exited, or stopped since the last status report". It was removed because the shell always writes any changed status of jobs before each prompt.
20192 FUTUF	RE DIRECTIONS
20193	None.
20194 SEE AI	
20195	Section 2.12 (on page 61), bg, fg, kill, wait
20196 CHAN	GE HISTORY
20197	First released in Issue 4.
20198 Issue 6	
20199	This utility is marked as part of the User Portability Utilities option.
20200	The JC shading is removed as job control is mandatory in this issue.

Utilities **join**

20201 **NAME**

20202 join — relational database operator

20203 SYNOPSIS

20204 join [-a file_number | -v file_number] [-e string] [-o list] [-t char]
20205 [-1 field] [-2 field] file1 file2

20206 DESCRIPTION

The *join* utility shall perform an equality join on the files *file1* and *file2*. The joined files shall be written to the standard output.

The join field is a field in each file on which the files are compared. The *join* utility shall write one line in the output for each pair of lines in *file1* and *file2* that have identical join fields. The output line by default shall consist of the join field, then the remaining fields from *file1*, then the remaining fields from *file2*. This format can be changed by using the **–o** option (see below). The **–a** option can be used to add unmatched lines to the output. The **–v** option can be used to output only unmatched lines.

The files *file1* and *file2* shall be ordered in the collating sequence of *sort* –**b** on the fields on which they shall be joined, by default the first in each line. All selected output shall be written in the same collating sequence.

The default input field separators shall be <black>s. In this case, multiple separators shall count as one field separator, and leading separators shall be ignored. The default output field separator shall be a <space>.

The field separator and collating sequence can be changed by using the -t option (see below).

If the same key appears more than once in either file, all combinations of the set of remaining fields in *file1* and the set of remaining fields in *file2* are output in the order of the lines encountered.

If the input files are not in the appropriate collating sequence, the results are unspecified.

20226 **OPTIONS**

20225

20227 20228

20229

20231 20232

20233

20235

2023620237

20238

20239

20241

20242 20243

20244

The *join* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported:

20230 —a file_number

−o list

Produce a line for each unpairable line in file *file_number*, where *file_number* is 1 or 2, in addition to the default output. If both **–a1** and **–a2** are specified, all unpairable lines shall be output.

20234 — e string Replace empty output fields in the list selected by –o with the string string.

Construct the output line to comprise the fields specified in *list*, each element of which shall have one of the following two forms:

- 1. *file_number.field*, where *file_number* is a file number and *field* is a decimal integer field number
- 2. 0 (zero), representing the join field

The elements of *list* shall be either comma-separated or <blank>-separated, as specified in Guideline 8 of the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines. The fields specified by *list* shall be written for all selected output lines. Fields selected by *list* that do not appear in the input shall be treated as empty output fields. (See the –e option.) Only specifically

join Utilities

20245 20246		requested fields shall be written. The application shall ensure that <i>list</i> is a single command line argument.
20247 20248 20249	–t char	Use character <i>char</i> as a separator, for both input and output. Every appearance of <i>char</i> in a line shall be significant. When this option is specified, the collating sequence shall be the same as <i>sort</i> without the $-\mathbf{b}$ option.
20250	−v file_numb	per
20251		Instead of the default output, produce a line only for each unpairable line in
20252 20253		$file_number$, where $file_number$ is 1 or 2. If both $-v1$ and $-v2$ are specified, all unpairable lines shall be output.
20254	−1 field	Join on the <i>field</i> th field of file 1. Fields are decimal integers starting with 1.
20255	−2 field	Join on the <i>field</i> th field of file 2. Fields are decimal integers starting with 1.
20256 OPER	ANDS	
20257	The following	ng operands shall be supported:
20258 20259	file1, file2	A pathname of a file to be joined. If either of the <i>file1</i> or <i>file2</i> operands is $'-'$, the standard input shall be used in its place.
20260 STDIN	J	
20261		d input shall be used only if the <i>file1</i> or <i>file2</i> operand is '-'. See the INPUT FILES
20262	section.	
20263 INPU		
20264	The input fil	les shall be text files.
	RONMENT VA	
20265 ENVII 20266	The following	ng environment variables shall affect the execution of <i>join</i> :
20266 20267		ng environment variables shall affect the execution of <i>join</i> : Provide a default value for the internationalization variables that are unset or null.
20266 20267 20268	The following	ng environment variables shall affect the execution of <i>join</i> : Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
20266 20267	The following	ng environment variables shall affect the execution of <i>join</i> : Provide a default value for the internationalization variables that are unset or null.
20266 20267 20268 20269 20270	The followin	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
20266 20267 20268 20269	The following	ng environment variables shall affect the execution of <i>join</i> : Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables
20266 20267 20268 20269 20270 20271	The followin	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables.
20266 20267 20268 20269 20270 20271 20272 20273 20274	The followin	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. TE Determine the locale of the collating sequence <i>join</i> expects to have been used when
20266 20267 20268 20269 20270 20271 20272 20273	The followin	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables.
20266 20267 20268 20269 20270 20271 20272 20273 20274 20275 20276	The followin	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. TE Determine the locale of the collating sequence <i>join</i> expects to have been used when the input files were sorted. Determine the locale for the interpretation of sequences of bytes of text data as
20266 20267 20268 20269 20270 20271 20272 20273 20274 20275 20276 20277	The followin	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. TE Determine the locale of the collating sequence <i>join</i> expects to have been used when the input files were sorted. 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
20266 20267 20268 20269 20270 20271 20272 20273 20274 20275 20276 20277 20278	The followin LANG LC_ALL LC_COLLAT LC_CTYPE	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. TE Determine the locale of the collating sequence <i>join</i> expects to have been used when the input files were sorted. 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).
20266 20267 20268 20269 20270 20271 20272 20273 20274 20275 20276 20277 20278	The followin	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. TE Determine the locale of the collating sequence <i>join</i> expects to have been used when the input files were sorted. 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). GES
20266 20267 20268 20269 20270 20271 20272 20273 20274 20275 20276 20277 20278	The followin LANG LC_ALL LC_COLLAT LC_CTYPE	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. TE Determine the locale of the collating sequence <i>join</i> expects to have been used when the input files were sorted. 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). GES Determine the locale that should be used to affect the format and contents of
20266 20267 20268 20269 20270 20271 20272 20273 20274 20275 20276 20277 20278 20279 20280	The followin LANG LC_ALL LC_COLLAT LC_CTYPE LC_MESSAG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. TE Determine the locale of the collating sequence join expects to have been used when the input files were sorted. 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). GES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
20266 20267 20268 20269 20270 20271 20272 20273 20274 20275 20276 20277 20278 20279 20280 20281 20282 XSI	The followin LANG LC_ALL LC_COLLAT LC_CTYPE LC_MESSAG NLSPATH	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. TE Determine the locale of the collating sequence join expects to have been used when the input files were sorted. 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). GES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. Determine the location of message catalogs for the processing of LC_MESSAGES.
20266 20267 20268 20269 20270 20271 20272 20273 20274 20275 20276 20277 20278 20279 20280 20281 20282 XSI	The followin LANG LC_ALL LC_COLLAT LC_CTYPE LC_MESSAG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. TE Determine the locale of the collating sequence join expects to have been used when the input files were sorted. 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). GES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. Determine the location of message catalogs for the processing of LC_MESSAGES.

Utilities join

20285 **STDOUT**

The *join* utility output shall be a concatenation of selected character fields. When the $-\mathbf{o}$ option is not specified, the output shall be:

20288 "%s%s%s\n", <join field>, <other file1 fields>,

20289 <other file2 fields>

20290 If the join field is not the first field in a file, the *<other file fields>* for that file shall be:

20291 <fields preceding join field>, <fields following join field>

When the $-\mathbf{o}$ option is specified, the output format shall be:

20293 "%s\n", <concatenation of fields>

where the concatenation of fields is described by the $-\mathbf{o}$ option, above.

For either format, each field (except the last) shall be written with its trailing separator character.

If the separator is the default (<blank>s), a single <space> shall be written after each field

20297 (except the last).

20298 STDERR

20299 The standard error shall be used only for diagnostic messages.

20300 OUTPUT FILES

20301 None.

20302 EXTENDED DESCRIPTION

20303 None.

20304 EXIT STATUS

20305 The following exit values shall be returned:

20306 0 All input files were output successfully.

20307 >0 An error occurred.

20308 CONSEQUENCES OF ERRORS

20309 Default.

20310 APPLICATION USAGE

Pathnames consisting of numeric digits or of the form *string.string* should not be specified directly following the **–o** list.

20313 EXAMPLES

The $-\mathbf{o}$ 0 field essentially selects the union of the join fields. For example, given file **phone**:

 20315
 !Name
 Phone Number

 20316
 Don
 +1 123-456-7890

 20317
 Hal
 +1 234-567-8901

 20318
 Yasushi
 +2 345-678-9012

20319 and file **fax**:

 20320
 ! Name
 Fax Number

 20321
 Don
 +1 123-456-7899

 20322
 Keith
 +1 456-789-0122

 20323
 Yasushi
 +2 345-678-9011

20324 (where the large expanses of white space are meant to each represent a single <tab>), the

20325 command:

join Utilities

```
20326
            join -t "<tab>" -a 1 -a 2 -e '(unknown)' -o 0,1.2,2.2 phone fax
20327
            would produce:
20328
             ! Name
                                Phone Number
                                                            Fax Number
20329
            Don
                                +1 123-456-7890
                                                             +1 123-456-7899
20330
            Hal
                                +1 234-567-8901
                                                             (unknown)
            Keith
                                                             +1 456-789-0122
20331
                                (unknown)
                                                             +2 345-678-9011
            Yasushi
                                +2 345-678-9012
20332
            Multiple instances of the same key will produce combinatorial results. The following:
20333
20334
            fa:
20335
                 a x
20336
                 ау
20337
                 a z
            fb:
20338
20339
                 a p
20340
            will produce:
20341
            ахр
20342
            аур
20343
            azp
            And the following:
20344
            fa:
20345
20346
                 a b c
20347
                 a d e
            fb:
20348
20349
                 a w x
20350
                 a y z
20351
                 аор
            will produce:
20352
20353
            abcwx
20354
            abcyz
            abcop
20355
20356
            a d e w x
            adeyz
20357
20358
            adeop
```

20359 RATIONALE

20360

20361

20362 20363

20364 20365

20366

20367

20368

2036920370

20371

The $-\mathbf{e}$ option is only effective when used with $-\mathbf{o}$ because, unless specific fields are identified using $-\mathbf{o}$, *join* is not aware of what fields might be empty. The exception to this is the join field, but identifying an empty join field with the $-\mathbf{e}$ string is not historical practice and some scripts might break if this were changed.

The 0 field in the $-\mathbf{o}$ list was adopted from the Tenth Edition version of *join* to satisfy international objections that the *join* in the base documents does not support the "full join" or "outer join" described in relational database literature. Although it has been possible to include a join field in the output (by default, or by field number using $-\mathbf{o}$), the join field could not be included for an unpaired line selected by $-\mathbf{a}$. The $-\mathbf{o}$ 0 field essentially selects the union of the join fields.

This sort of outer join was not possible with the *join* commands in the base documents. The $-\mathbf{o}$ 0 field was chosen because it is an upwards-compatible change for applications. An alternative

Utilities join

20372 20373 20374	was considered: have the join field represent the union of the fields in the files (where they are identical for matched lines, and one or both are null for unmatched lines). This was not adopted because it would break some historical applications.
20375	The ability to specify $\mathit{file2}$ as – is not historical practice; it was added for completeness.
20376 20377 20378 20379	The $-\mathbf{v}$ option is not historical practice, but was considered necessary because it permitted the writing of <i>only</i> those lines that do not match on the join field, as opposed to the $-\mathbf{a}$ option, which prints both lines that do and do not match. This additional facility is parallel with the $-\mathbf{v}$ option of <i>grep</i> .
20380 20381 20382	Some historical implementations have been encountered where a blank line in one of the input files was considered to be the end of the file; the description in this volume of IEEE Std 1003.1-2001 does not cite this as an allowable case.
20383 FUTUR 20384	E DIRECTIONS None.
20385 SEE AL 20386	SO awk, comm, sort, uniq
20387 CHAN 0 20388	GE HISTORY First released in Issue 2.
20389 Issue 6 20390	The obsolescent $-\mathbf{j}$ options and the multi-argument $-\mathbf{o}$ option are withdrawn in this issue.
20391	The normative text is reworded to avoid use of the term "must" for application requirements.

kill Utilities

```
      20392 NAME

      20393 kill — terminate or signal processes

      20394 SYNOPSIS

      20395 kill —s signal_name pid ...

      20396 kill —l [exit_status]

      20397 XSI kill [-signal_name] pid ...

      20398 kill [-signal_number] pid ...

      20399
```

20400 DESCRIPTION

20401

20402

20403

20404

20405

20412

20422

20423

20424

20425 20426 The *kill* utility shall send a signal to the process or processes specified by each *pid* operand.

For each *pid* operand, the *kill* utility shall perform actions equivalent to the *kill*() function defined in the System Interfaces volume of IEEE Std 1003.1-2001 called with the following arguments:

- The value of the *pid* operand shall be used as the *pid* argument.
- The *sig* argument is the value specified by the **-s** option, *-signal_number* option, or the *-signal_name* option, or by SIGTERM, if none of these options is specified.

20408 **OPTIONS**

The *kill* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines, except that in the last two SYNOPSIS forms, the *-signal_number* and *-signal_name* options are usually more than a single character.

The following options shall be supported:

 $-\mathbf{l}$ (The letter ell.) Write all values of *signal_name* supported by the implementation, if 20413 20414 no operand is given. If an exit_status operand is given and it is a value of the '?' 20415 shell special parameter (see Section 2.5.2 (on page 34) and wait) corresponding to a 20416 process that was terminated by a signal, the signal_name corresponding to the signal that terminated the process shall be written. If an exit_status operand is 20417 20418 given and it is the unsigned decimal integer value of a signal number, the 20419 signal_name (the symbolic constant name without the SIG prefix defined in the Base Definitions volume of IEEE Std 1003.1-2001) corresponding to that signal 20420 shall be written. Otherwise, the results are unspecified. 20421

-s signal name

Specify the signal to send, using one of the symbolic names defined in the <signal.h> header. Values of signal_name shall be recognized in a case-independent fashion, without the SIG prefix. In addition, the symbolic name 0 shall be recognized, representing the signal value zero. The corresponding signal shall be sent instead of SIGTERM.

20427 sent instead of SIGTERN

–signal name 20428 XSI Equivalent to **-s** *signal_name*. 20429 20430 XSI -signal_number Specify a non-negative decimal integer, signal_number, representing the signal to 20431 be used instead of SIGTERM, as the *sig* argument in the effective call to *kill()*. The 20432 correspondence between integer values and the sig value used is shown in the 20433 following list. 20434 The effects of specifying any signal_number other than those listed below are 20435 undefined. 20436

kill **Utilities**

20437		0 0			
20438		1 SIGHUP			
20439		2 SIGINT			
20440		3 SIGQUIT			
20441		6 SIGABRT			
20442		9 SIGKILL			
20443		14 SIGALRM			
20444		15 SIGTERM			
20445		If the first argument is a negative integer, it shall be interpreted as a <i>-signal_number</i>			
20445		option, not as a negative <i>pid</i> operand specifying a process group.			
20447 OPERA	NDS				
20448	The following operands shall be supported:				
20449	pid	One of the following:			
20450 20451 20452 20453		1. A decimal integer specifying a process or process group to be signaled. The process or processes selected by positive, negative, and zero values of the <i>pid</i>			
20454 20455 20456		operand shall be as described for the <i>kill()</i> function. If process number 0 is specified, all processes in the current process group shall be signaled. For the effects of negative <i>pid</i> numbers, see the <i>kill()</i> function defined in the System Interfaces volume of IEEE Std 1003.1-2001. If the first <i>pid</i> operand is negative, it should be preceded by "" to keep it from being interpreted as an option.			
20455		specified, all processes in the current process group shall be signaled. For the effects of negative <i>pid</i> numbers, see the <i>kill()</i> function defined in the System Interfaces volume of IEEE Std 1003.1-2001. If the first <i>pid</i> operand is negative,			

20464 STDIN

20463

Not used. 20465

20466 INPUT FILES

None. 20467

20468 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *kill*: 20469

terminated by a signal.

20470	LANG	Provide a default value for the internationalization variables that are unset or null.		
20471		(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,		
20472		Internationalization Variables for the precedence of internationalization variables		
20473		used to determine the values of locale categories.)		
20474 20475	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.		
20476 20477	LC_CTYPE	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		
20478		arguments).		

kill Utilities

20479 LC_MESSAGES 20480 Determine the locale that should be used to affect the format and contents of 20481 diagnostic messages written to standard error. NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 20482 XSI 20483 ASYNCHRONOUS EVENTS Default. 20484 20485 **STDOUT** When the **-l** option is not specified, the standard output shall not be used. 20486 20487 When the -l option is specified, the symbolic name of each signal shall be written in the following format: 20488 "%s%c", <signal name>, <separator> 20489 where the *<signal_name>* is in uppercase, without the **SIG** prefix, and the *<separator>* shall be 20490 20491 either a <newline> or a <space>. For the last signal written, <*separator>* shall be a <newline>. When both the -l option and exit_status operand are specified, the symbolic name of the 20492 20493 corresponding signal shall be written in the following format: 20494 "%s\n", < signal name> 20495 STDERR The standard error shall be used only for diagnostic messages. 20496 20497 OUTPUT FILES 20498 None. 20499 EXTENDED DESCRIPTION None. 20500 20501 EXIT STATUS The following exit values shall be returned: 20502 At least one matching process was found for each pid operand, and the specified signal was 20503 20504 successfully processed for at least one matching process. 20505 >0 An error occurred. 20506 CONSEQUENCES OF ERRORS Default. 20507 20508 APPLICATION USAGE 20509 Process numbers can be found by using ps. The job control job ID notation is not required to work as expected when kill is operating in its 20510 20511 own utility execution environment. In either of the following examples: nohup kill %1 & 20512 20513 system("kill %1"); the kill operates in a different environment and does not share the shell's understanding of job 20514 numbers. 20515 20516 EXAMPLES Any of the commands: 20517 kill -9 100 -165 20518 kill -s kill 100 -165 20519 20520 kill -s KILL 100 -165

kill **Utilities**

sends the SIGKILL signal to the process whose process ID is 100 and to all processes whose process group ID is 165, assuming the sending process has permission to send that signal to the specified processes, and that they exist.

The System Interfaces volume of IEEE Std 1003.1-2001 and this volume of IEEE Std 1003.1-2001 do not require specific signal numbers for any signal_names. Even the -signal_number option provides symbolic (although numeric) names for signals. If a process is terminated by a signal, its exit status indicates the signal that killed it, but the exact values are not specified. The kill -1 option, however, can be used to map decimal signal numbers and exit status values into the name of a signal. The following example reports the status of a terminated job:

```
20530
            job
20531
            stat=$?
            if [ $stat -eq 0 ]
20532
20533
20534
                 echo job completed successfully.
            elif [ $stat -qt 128 ]
20535
            then
20536
                 echo job terminated by signal SIG$(kill -1 $stat).
20537
20538
            else
20539
                 echo job terminated with error code $stat.
20540
            fi
```

To send the default signal to a process group (say 123), an application should use a command similar to one of the following:

```
20543
            kill -TERM -123
20544
            kill -- -123
```

20545 RATIONALE

20521

20522

20523

20524 20525

20526

20527 20528

20529

20541 20542

20546

20547 20548

20549 20550

20551

20552 20553

20554

20555

20556

20557

20558

20559

20560

20561 20562

20563

20564

20565

20567

The –l option originated from the C shell, and is also implemented in the KornShell. The C shell output can consist of multiple output lines because the signal names do not always fit on a single line on some terminal screens. The KornShell output also included the implementationdefined signal numbers and was considered by the standard developers to be too difficult for scripts to parse conveniently. The specified output format is intended not only to accommodate the historical C shell output, but also to permit an entirely vertical or entirely horizontal listing on systems for which this is appropriate.

An early proposal invented the name SIGNULL as a signal_name for signal 0 (used by the System Interfaces volume of IEEE Std 1003.1-2001 to test for the existence of a process without sending it a signal). Since the signal_name 0 can be used in this case unambiguously, SIGNULL has been removed.

An early proposal also required symbolic signal_names to be recognized with or without the SIG prefix. Historical versions of kill have not written the SIG prefix for the –l option and have not recognized the SIG prefix on signal_names. Since neither applications portability nor ease-of-use would be improved by requiring this extension, it is no longer required.

To avoid an ambiguity of an initial negative number argument specifying either a signal number or a process group, IEEE Std 1003.1-2001 mandates that it is always considered the former by implementations that support the XSI option. It also requires that conforming applications always use the "--" options terminator argument when specifying a process group, unless an option is also specified.

The -s option was added in response to international interest in providing some form of kill that 20566 meets the Utility Syntax Guidelines.

kill Utilities

20568 The job control job ID notation is not required to work as expected when kill is operating in its 20569 own utility execution environment. In either of the following examples: 20570 nohup kill %1 & system("kill %1"); 20571 20572 the kill operates in a different environment and does not understand how the shell has managed 20573 its job numbers. 20574 **FUTURE DIRECTIONS** 20575 None. **20576 SEE ALSO** 20577 Chapter 2 (on page 29), ps, wait, the System Interfaces volume of IEEE Std 1003.1-2001, kill(), the Base Definitions volume of IEEE Std 1003.1-2001, < signal.h> 20578 20579 CHANGE HISTORY First released in Issue 2. 20580 20581 Issue 6 The obsolescent versions of the SYNOPSIS are turned into non-obsolescent features of the XSI 20582 option, corresponding to a similar change in the *trap* special built-in. 20583

Utilities lex

20584 NAME

20585 lex — generate programs for lexical tasks (**DEVELOPMENT**)

20586 SYNOPSIS

20587 CD lex [-t][-n]-v][file ...]

20588

20589 **DESCRIPTION**

The *lex* utility shall generate C programs to be used in lexical processing of character input, and that can be used as an interface to *yacc*. The C programs shall be generated from *lex* source code and conform to the ISO C standard. Usually, the *lex* utility shall write the program it generates to the file **lex.yy.c**; the state of this file is unspecified if *lex* exits with a non-zero exit status. See the EXTENDED DESCRIPTION section for a complete description of the *lex* input language.

20595 OPTIONS

The *lex* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

20598 The following options shall be supported:

20599 -**n** Suppress the summary of statistics usually written with the -**v** option. If no table sizes are specified in the *lex* source code and the -**v** option is not specified, then -**n** is implied.

20602 —t Write the resulting program to standard output instead of lex.yy.c.

Write a summary of *lex* statistics to the standard output. (See the discussion of *lex* table sizes in **Definitions in lex** (on page 537).) If the -t option is specified and -n is not specified, this report shall be written to standard error. If table sizes are specified in the *lex* source code, and if the -n option is not specified, the -v option may be enabled.

20608 OPERANDS

20609 The following operand shall be supported:

20610 file A pathname of an input file. If more than one such file is specified, all files shall be concatenated to produce a single lex program. If no file operands are specified, or if a file operand is '-', the standard input shall be used.

20613 **STDIN**

The standard input shall be used if no *file* operands are specified, or if a *file* operand is '-'. See INPUT FILES.

20616 INPUT FILES

The input files shall be text files containing *lex* source code, as described in the EXTENDED DESCRIPTION section.

20619 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *lex*:

20621 LANG Provide a default value for the internationalization variables that are unset or null.
20622 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
20623 Internationalization Variables for the precedence of internationalization variables
20624 used to determine the values of locale categories.)

20625 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

20627 LC_COLLATE

Determine the locale for the behavior of ranges, equivalence classes, and multi-

lex **Utilities**

20629 character collating elements within regular expressions. If this variable is not set to 20630 the POSIX locale, the results are unspecified. 20631 LC_CTYPE 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 20632 arguments and input files), and the behavior of character classes within regular 20633 expressions. If this variable is not set to the POSIX locale, the results are 20634 unspecified. 20635

LC_MESSAGES 20636

Determine the locale that should be used to affect the format and contents of 20637 20638 diagnostic messages written to standard error.

NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 20639 XSI

20640 ASYNCHRONOUS EVENTS

20641 Default.

20642 **STDOUT**

20645

20646

20647

20648

20649 20650

20651

20653

20654

20655 20656

20657

20658

20659

20660

20661

20662

20664 20665

If the -t option is specified, the text file of C source code output of lex shall be written to 20643 20644 standard output.

If the -t option is not specified:

- Implementation-defined informational, error, and warning messages concerning the contents of *lex* source code input shall be written to either the standard output or standard error.
- If the $-\mathbf{v}$ option is specified and the $-\mathbf{n}$ option is not specified, lex statistics shall also be written to either the standard output or standard error, in an implementation-defined format. These statistics may also be generated if table sizes are specified with a '%' operator in the Definitions section, as long as the -**n** option is not specified.

20652 STDERR

If the -t option is specified, implementation-defined informational, error, and warning messages concerning the contents of *lex* source code input shall be written to the standard error.

If the **–t** option is not specified:

- 1. Implementation-defined informational, error, and warning messages concerning the contents of lex source code input shall be written to either the standard output or standard error.
- 2. If the $-\mathbf{v}$ option is specified and the $-\mathbf{n}$ option is not specified, lex statistics shall also be written to either the standard output or standard error, in an implementation-defined format. These statistics may also be generated if table sizes are specified with a '%' operator in the *Definitions* section, as long as the **-n** option is not specified.

20663 OUTPUT FILES

A text file containing C source code shall be written to lex.yy.c, or to the standard output if the **-t** option is present.

20666 EXTENDED DESCRIPTION

Each input file shall contain lex source code, which is a table of regular expressions with 20667 corresponding actions in the form of C program fragments. 20668

When **lex.yy.c** is compiled and linked with the *lex* library (using the -1 l operand with c99), the 20669 resulting program shall read character input from the standard input and shall partition it into 20670 strings that match the given expressions. 20671

Utilities lex

20672 When an expression is matched, these actions shall occur:

• The input string that was matched shall be left in *yytext* as a null-terminated string; *yytext* shall either be an external character array or a pointer to a character string. As explained in **Definitions in lex**, the type can be explicitly selected using the **%array** or **%pointer** declarations, but the default is implementation-defined.

- The external **int** *yyleng* shall be set to the length of the matching string.
- The expression's corresponding program fragment, or action, shall be executed.

During pattern matching, *lex* shall search the set of patterns for the single longest possible match. Among rules that match the same number of characters, the rule given first shall be chosen.

The general format of *lex* source shall be:

 20683
 Definitions

 20684
 %%

 20685
 Rules

 20686
 %%

*User*Subroutines

The first "%%" is required to mark the beginning of the rules (regular expressions and actions); the second "%%" is required only if user subroutines follow.

Any line in the *Definitions* section beginning with a <black> shall be assumed to be a C program fragment and shall be copied to the external definition area of the **lex.yy.c** file. Similarly, anything in the *Definitions* section included between delimiter lines containing only "%{ " and "%} " shall also be copied unchanged to the external definition area of the **lex.yy.c** file.

Any such input (beginning with a <black> or within "%{" and "%}" delimiter lines) appearing at the beginning of the *Rules* section before any rules are specified shall be written to **lex.yy.c** after the declarations of variables for the *yylex*() function and before the first line of code in *yylex*(). Thus, user variables local to *yylex*() can be declared here, as well as application code to execute upon entry to *yylex*().

The action taken by *lex* when encountering any input beginning with a
blank> or within "% { " and "% } " delimiter lines appearing in the *Rules* section but coming after one or more rules is undefined. The presence of such input may result in an erroneous definition of the yylex() function.

Definitions in lex

Definitions appear before the first "%%" delimiter. Any line in this section not contained between "%{" and "%}" lines and not beginning with a

blank> shall be assumed to define a lex substitution string. The format of these lines shall be:

20707 name substitute

If a *name* does not meet the requirements for identifiers in the ISO C standard, the result is undefined. The string *substitute* shall replace the string *{name}* when it is used in a rule. The *name* string shall be recognized in this context only when the braces are provided and when it does not appear within a bracket expression or within double-quotes.

In the *Definitions* section, any line beginning with a '%' (percent sign) character and followed by an alphanumeric word beginning with either 's' or 'S' shall define a set of start conditions. Any line beginning with a '%' followed by a word beginning with either 'x' or 'X' shall define a set of exclusive start conditions. When the generated scanner is in a %s state, patterns with no

lex **Utilities**

state specified shall be also active; in a %x state, such patterns shall not be active. The rest of the line, after the first word, shall be considered to be one or more
 slank>-separated names of start conditions. Start condition names shall be constructed in the same way as definition names. Start conditions can be used to restrict the matching of regular expressions to one or more states as described in **Regular Expressions in lex** (on page 539).

Implementations shall accept either of the following two mutually-exclusive declarations in the **Definitions** section:

%array Declare the type of *yytext* to be a null-terminated character array.

%pointer Declare the type of *yytext* to be a pointer to a null-terminated character string.

The default type of yytext is implementation-defined. If an application refers to yytext outside of the scanner source file (that is, via an extern), the application shall include the appropriate **%array** or **%pointer** declaration in the scanner source file.

Implementations shall accept declarations in the *Definitions* section for setting certain internal table sizes. The declarations are shown in the following table.

Table 4-9 Table Size Declarations in *lex*

Declaration	Description	Minimum Value
% p <i>n</i>	Number of positions	2 500
% n <i>n</i>	Number of states	500
% a n	Number of transitions	2 000
% e n	Number of parse tree nodes	1 000
% k <i>n</i>	Number of packed character classes	1 000
% o n	Size of the output array	3 000

In the table, n represents a positive decimal integer, preceded by one or more

 lank>s. The exact meaning of these table size numbers is implementation-defined. The implementation shall document how these numbers affect the lex utility and how they are related to any output that may be generated by the implementation should limitations be encountered during the execution of lex. It shall be possible to determine from this output which of the table size values needs to be modified to permit lex to successfully generate tables for the input language. The values in the column Minimum Value represent the lowest values conforming implementations shall provide.

Rules in lex

The rules in *lex* source files are a table in which the left column contains regular expressions and the right column contains actions (C program fragments) to be executed when the expressions are recognized.

```
ERE action
20750
20751
            ERE action
```

20752

20716

20717

20718 20719

20720

20721

20722 20723

20724

20725

20726

20727

20728

20729

20730

20738

20739

20740

20741

20742 20743

20744

20745

20746 20747

20748

20749

20753

20754 20755

20756

20757

The extended regular expression (ERE) portion of a row shall be separated from action by one or following conditions:

- The entire expression appears within double-quotes.
- The <blank>s appear within double-quotes or square brackets.

Utilities lex

• Each
blank> is preceded by a backslash character.

User Subroutines in lex

r/x

{name}

Anything in the user subroutines section shall be copied to **lex.yy.c** following *yylex*().

Regular Expressions in lex

The *lex* utility shall support the set of extended regular expressions (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.4, Extended Regular Expressions), with the following additions and exceptions to the syntax:

"..." Any string enclosed in double-quotes shall represent the characters within the double-quotes as themselves, except that backslash escapes (which appear in the following table) shall be recognized. Any backslash-escape sequence shall be terminated by the closing quote. For example, "\01""1" represents a single string: the octal value 1 followed by the character '1'.

<state>r, <state1,state2,...>r

The regular expression *r* shall be matched only when the program is in one of the start conditions indicated by *state*, *state1*, and so on; see **Actions in lex** (on page 541). (As an exception to the typographical conventions of the rest of this volume of IEEE Std 1003.1-2001, in this case *<state>* does not represent a metavariable, but the literal angle-bracket characters surrounding a symbol.) The start condition shall be recognized as such only at the beginning of a regular expression.

The regular expression r shall be matched only if it is followed by an occurrence of regular expression x (x is the instance of trailing context, further defined below). The token returned in yytext shall only match r. If the trailing portion of r matches the beginning of x, the result is unspecified. The r expression cannot include further trailing context or the '\$' (match-end-of-line) operator; x cannot include the '\$' (match-beginning-of-line) operator, nor trailing context, nor the '\$' operator. That is, only one occurrence of trailing context is allowed in a lex regular expression, and the '\$' operator only can be used at the beginning of such an expression.

When *name* is one of the substitution symbols from the *Definitions* section, the string, including the enclosing braces, shall be replaced by the *substitute* value. The *substitute* value shall be treated in the extended regular expression as if it were enclosed in parentheses. No substitution shall occur if {name} occurs within a bracket expression or within double-quotes.

Within an ERE, a backslash character shall be considered to begin an escape sequence as specified in the table in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5, File Format Notation ('\\', '\a', '\b', '\f', '\n', '\r', '\t', '\v'). In addition, the escape sequences in the following table shall be recognized.

A literal <newline> cannot occur within an ERE; the escape sequence $' \n'$ can be used to represent a <newline>. A <newline> shall not be matched by a period operator.

lex Utilities

Table 4-10 Escape Sequences in *lex*

20798	Escape		
20799	Sequence	Description	Meaning
20800	\digits	A backslash character followed	The character whose encoding is
20801		by the longest sequence of one,	represented by the one, two, or
20802		two, or three octal-digit	three-digit octal integer. If the
20803		characters (01234567). If all of	size of a byte on the system is
20804		the digits are 0 (that is,	greater than nine bits, the valid
20805		representation of the NUL	escape sequence used to
20806		character), the behavior is	represent a byte is
20807		undefined.	implementation-defined. Multi-
20808			byte characters require multiple,
20809			concatenated escape sequences
20810			of this type, including the
20811			leading $' \setminus '$ for each byte.
20812	\xdigits	A backslash character followed	The character whose encoding is
20813		by the longest sequence of	represented by the hexadecimal
20814		hexadecimal-digit characters	integer.
20815		(01234567abcdefABCDEF). If all	
20816		of the digits are 0 (that is,	
20817		representation of the NUL	
20818		character), the behavior is	
20819		undefined.	
20820	\c	A backslash character followed	The character 'c', unchanged.
20821		by any character not described	
20822		in this table or in the table in the	
20823		Base Definitions volume of	
20824		IEEE Std 1003.1-2001, Chapter 5,	
20825		File Format Notation $(' \setminus \ \)'$,	
20826		'\a','\b','\f','\n','\r',	
20827		'\t','\v').	
	H	 	

Note:

If a '\x' sequence needs to be immediately followed by a hexadecimal digit character, a sequence such as "\x1""1" can be used, which represents a character containing the value 1, followed by the character '1'.

The order of precedence given to extended regular expressions for *lex* differs from that specified in the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.4, Extended Regular Expressions. The order of precedence for *lex* shall be as shown in the following table, from high to low.

Note:

The escaped characters entry is not meant to imply that these are operators, but they are included in the table to show their relationships to the true operators. The start condition, trailing context, and anchoring notations have been omitted from the table because of the placement restrictions described in this section; they can only appear at the beginning or ending of an ERE.

Utilities lex

20840	Table 4-11	ERE Precedence in <i>lex</i>
-------	-------------------	------------------------------

 $\begin{array}{c} 20853 \\ 20854 \end{array}$

Extended Regular Expression	Precedence
collation-related bracket symbols	[= =] [: :] []
escaped characters	\ <special character=""></special>
bracket expression	[]
quoting	""
grouping	()
definition	{name}
single-character RE duplication	* + ?
concatenation	
interval expression	{m,n}
alternation	

The ERE anchoring operators '^' and '\$' do not appear in the table. With *lex* regular expressions, these operators are restricted in their use: the '^' operator can only be used at the beginning of an entire regular expression, and the '\$' operator only at the end. The operators apply to the entire regular expression. Thus, for example, the pattern "(^abc) | (def\$)" is undefined; it can instead be written as two separate rules, one with the regular expression "^abc" and one with "def\$", which share a common action via the special '|' action (see below). If the pattern were written "^abc|def\$", it would match either "abc" or "def" on a line by itself.

Unlike the general ERE rules, embedded anchoring is not allowed by most historical *lex* implementations. An example of embedded anchoring would be for patterns such as " ($^{\circ}$ |)foo(|\$)" to match "foo" when it exists as a complete word. This functionality can be obtained using existing *lex* features:

Note also that '\$' is a form of trailing context (it is equivalent to " \n ") and as such cannot be used with regular expressions containing another instance of the operator (see the preceding discussion of trailing context).

The additional regular expressions trailing-context operator '/' can be used as an ordinary character if presented within double-quotes, "/"; preceded by a backslash, "\/"; or within a bracket expression, "[/]". The start-condition '<' and '>' operators shall be special only in a start condition at the beginning of a regular expression; elsewhere in the regular expression they shall be treated as ordinary characters.

Actions in lex

The action to be taken when an ERE is matched can be a C program fragment or the special actions described below; the program fragment can contain one or more C statements, and can also include special actions. The empty C statement ';' shall be a valid action; any string in the **lex.yy.c** input that matches the pattern portion of such a rule is effectively ignored or skipped. However, the absence of an action shall not be valid, and the action *lex* takes in such a condition is undefined.

The specification for an action, including C statements and special actions, can extend across several lines if enclosed in braces:

```
20883 ERE <one or more blanks> { program statement program statement }
```

lex Utilities

The default action when a string in the input to a **lex.yy.c** program is not matched by any expression shall be to copy the string to the output. Because the default behavior of a program generated by *lex* is to read the input and copy it to the output, a minimal *lex* source program that has just "%%" shall generate a C program that simply copies the input to the output unchanged.

Four special actions shall be available:

ECHO; REJECT; BEGIN

The action $' \mid '$ means that the action for the next rule is the action for this rule. Unlike the other three actions, $' \mid '$ cannot be enclosed in braces or be semicolon-terminated; the application shall ensure that it is specified alone, with no other actions.

ECHO; Write the contents of the string *yytext* on the output.

Usually only a single expression is matched by a given string in the input. **REJECT** means "continue to the next expression that matches the current input", and shall cause whatever rule was the second choice after the current rule to be executed for the same input. Thus, multiple rules can be matched and executed for one input string or overlapping input strings. For example, given the regular expressions "xyz" and "xy" and the input "xyz", usually only the regular expression "xyz" would match. The next attempted match would start after **z**. If the last action in the "xyz" rule is **REJECT**, both this rule and the "xy" rule would be executed. The **REJECT** action may be implemented in such a fashion that flow of control does not continue after it, as if it were equivalent to a **goto** to another part of *yylex*(). The use of **REJECT** may result in somewhat larger and slower scanners.

BEGIN The action:

REJECT;

BEGIN newstate;

switches the state (start condition) to *newstate*. If the string *newstate* has not been declared previously as a start condition in the *Definitions* section, the results are unspecified. The initial state is indicated by the digit '0' or the token **INITIAL**.

The functions or macros described below are accessible to user code included in the *lex* input. It is unspecified whether they appear in the C code output of *lex*, or are accessible only through the –**l l** operand to *c99* (the *lex* library).

int yylex(void)

Performs lexical analysis on the input; this is the primary function generated by the *lex* utility. The function shall return zero when the end of input is reached; otherwise, it shall return non-zero values (tokens) determined by the actions that are selected.

int yymore(void)

When called, indicates that when the next input string is recognized, it is to be appended to the current value of *yytext* rather than replacing it; the value in *yyleng* shall be adjusted accordingly.

int yyless(int n)

Retains *n* initial characters in *yytext*, NUL-terminated, and treats the remaining characters as if they had not been read; the value in *yyleng* shall be adjusted accordingly.

int input(void)

Returns the next character from the input, or zero on end-of-file. It shall obtain input from the stream pointer *yyin*, although possibly via an intermediate buffer. Thus, once scanning has begun, the effect of altering the value of *yyin* is undefined. The character read shall be removed from the input stream of the scanner without any processing by the scanner.

Utilities lex

```
20931 int unput(int c)
```

20935

20936

20938 20939

20940

20941

20943

20944

20954

20955

20956

20957 20958

20959

20960

20961 20962

20963

20964

20965

20966

20967

20969 20970

Returns the character 'c' to the input; *yytext* and *yyleng* are undefined until the next expression is matched. The result of using *unput*() for more characters than have been input is unspecified.

The following functions shall appear only in the *lex* library accessible through the **–l l** operand; they can therefore be redefined by a conforming application:

20937 int yywrap(void)

Called by *yylex()* at end-of-file; the default *yywrap()* shall always return 1. If the application requires *yylex()* to continue processing with another source of input, then the application can include a function *yywrap()*, which associates another file with the external variable **FILE*** *yyin* and shall return a value of zero.

int main(int argc, char *argv[])

Calls *yylex*() to perform lexical analysis, then exits. The user code can contain *main*() to perform application-specific operations, calling *yylex*() as applicable.

Except for *input*(), *unput*(), and *main*(), all external and static names generated by *lex* shall begin with the prefix **yy** or **YY**.

20947 EXIT STATUS

The following exit values shall be returned:

20949 0 Successful completion.

20950 >0 An error occurred.

20951 CONSEQUENCES OF ERRORS

20952 Default.

20953 APPLICATION USAGE

Conforming applications are warned that in the *Rules* section, an ERE without an action is not acceptable, but need not be detected as erroneous by *lex*. This may result in compilation or runtime errors.

The purpose of *input()* is to take characters off the input stream and discard them as far as the lexical analysis is concerned. A common use is to discard the body of a comment once the beginning of a comment is recognized.

The *lex* utility is not fully internationalized in its treatment of regular expressions in the *lex* source code or generated lexical analyzer. It would seem desirable to have the lexical analyzer interpret the regular expressions given in the *lex* source according to the environment specified when the lexical analyzer is executed, but this is not possible with the current *lex* technology. Furthermore, the very nature of the lexical analyzers produced by *lex* must be closely tied to the lexical requirements of the input language being described, which is frequently locale-specific anyway. (For example, writing an analyzer that is used for French text is not automatically useful for processing other languages.)

20968 EXAMPLES

The following is an example of a *lex* program that implements a rudimentary scanner for a Pascal-like syntax:

lex Utilities

```
20977
            DIGIT
                       [0-9]
20978
            ID
                       [a-z][a-z0-9]*
20979
            응응
20980
            {DIGIT}+ {
20981
                 printf("An integer: %s (%d)\n", yytext,
                      atoi(yytext));
20982
20983
            {DIGIT}+"."{DIGIT}*
20984
                 printf("A float: %s (%g)\n", yytext,
20985
                     atof(yytext));
20986
20987
20988
            if | then | begin | end | procedure | function
                 printf("A keyword: %s\n", yytext);
20989
20990
                     printf("An identifier: %s\n", yytext);
20991
            {ID}
            "+"|"-"|"*"|"/"
                                       printf("An operator: %s\n", yytext);
20992
            "{"[^}\n]*"}"
                              /* Eat up one-line comments. */
20993
            \lceil \t \n \rceil +
                               /* Eat up white space. */
20994
20995
                printf("Unrecognized character: %s\n", yytext);
            99
20996
            int main(int argc, char *argv[])
20997
20998
20999
                 ++argv, --argc; /* Skip over program name. */
                 if (argc > 0)
21000
21001
                     yyin = fopen(argv[0], "r");
                 else
21002
21003
                     yyin = stdin;
                 yylex();
21004
21005
```

21006 RATIONALE

21007 21008

21009 21010

21011

21012

21013

21014 21015

21016

21017 21018

21019

21020

21021

Even though the –c option and references to the C language are retained in this description, *lex* may be generalized to other languages, as was done at one time for EFL, the Extended FORTRAN Language. Since the *lex* input specification is essentially language-independent, versions of this utility could be written to produce Ada, Modula-2, or Pascal code, and there are known historical implementations that do so.

The current description of *lex* bypasses the issue of dealing with internationalized EREs in the *lex* source code or generated lexical analyzer. If it follows the model used by *awk* (the source code is assumed to be presented in the POSIX locale, but input and output are in the locale specified by the environment variables), then the tables in the lexical analyzer produced by *lex* would interpret EREs specified in the *lex* source in terms of the environment variables specified when *lex* was executed. The desired effect would be to have the lexical analyzer interpret the EREs given in the *lex* source according to the environment specified when the lexical analyzer is executed, but this is not possible with the current *lex* technology.

The description of octal and hexadecimal-digit escape sequences agrees with the ISO C standard use of escape sequences. See the RATIONALE for *ed* for a discussion of bytes larger than 9 bits

Utilities lex

being represented by octal values. Hexadecimal values can represent larger bytes and multi-byte characters directly, using as many digits as required.

 There is no detailed output format specification. The observed behavior of *lex* under four different historical implementations was that none of these implementations consistently reported the line numbers for error and warning messages. Furthermore, there was a desire that *lex* be allowed to output additional diagnostic messages. Leaving message formats unspecified avoids these formatting questions and problems with internationalization.

Although the x specifier for *exclusive* start conditions is not historical practice, it is believed to be a minor change to historical implementations and greatly enhances the usability of *lex* programs since it permits an application to obtain the expected functionality with fewer statements.

The %array and %pointer declarations were added as a compromise between historical systems. The System V-based *lex* copies the matched text to a *yytext* array. The *flex* program, supported in BSD and GNU systems, uses a pointer. In the latter case, significant performance improvements are available for some scanners. Most historical programs should require no change in porting from one system to another because the string being referenced is null-terminated in both cases. (The method used by *flex* in its case is to null-terminate the token in place by remembering the character that used to come right after the token and replacing it before continuing on to the next scan.) Multi-file programs with external references to *yytext* outside the scanner source file should continue to operate on their historical systems, but would require one of the new declarations to be considered strictly portable.

The description of EREs avoids unnecessary duplication of ERE details because their meanings within a *lex* ERE are the same as that for the ERE in this volume of IEEE Std 1003.1-2001.

The intention in breaking the list of functions into those that may appear in **lex.yy.c** versus those that only appear in **libl.a** is that only those functions in **libl.a** can be reliably redefined by a conforming application.

The descriptions of standard output and standard error are somewhat complicated because historical *lex* implementations chose to issue diagnostic messages to standard output (unless –t was given). IEEE Std 1003.1-2001 allows this behavior, but leaves an opening for the more expected behavior of using standard error for diagnostics. Also, the System V behavior of writing the statistics when any table sizes are given is allowed, while BSD-derived systems can avoid it. The programmer can always precisely obtain the desired results by using either the –t or –n options.

The OPERANDS section does not mention the use of - as a synonym for standard input; not all historical implementations support such usage for any of the *file* operands.

A description of the *translation table* was deleted from early proposals because of its relatively low usage in historical applications.

The change to the definition of the *input*() function that allows buffering of input presents the opportunity for major performance gains in some applications.

The following examples clarify the differences between *lex* regular expressions and regular expressions appearing elsewhere in this volume of IEEE Std 1003.1-2001. For regular expressions

lex Utilities

of the form "r/x", the string matching r is always returned; confusion may arise when the beginning of x matches the trailing portion of r. For example, given the regular expression "a*b/cc" and the input "aaabcc", yytext would contain the string "aaab" on this match. But given the regular expression "x*/xy" and the input "xxxy", the token xxx, not xx, is returned by some implementations because xxx matches "x*".

In the rule "ab*/bc", the "b*" at the end of r extends r's match into the beginning of the trailing context, so the result is unspecified. If this rule were "ab/bc", however, the rule matches the text "ab" when it is followed by the text "bc". In this latter case, the matching of r cannot extend into the beginning of x, so the result is specified.

21078 FUTURE DIRECTIONS

21079 None.

21080 SEE ALSO

21081 *c99*, *ed*, *yacc*

21082 CHANGE HISTORY

First released in Issue 2.

21084 Issue 6

21085 This utility is marked as part of the C-Language Development Utilities option.

21086 The obsolescent -c option is withdrawn in this issue.

The normative text is reworded to avoid use of the term "must" for application requirements.

Utilities link

```
21088 NAME
21089
              link — call link() function
21090 SYNOPSIS
              link file1 file2
21091 XSI
21092
21093 DESCRIPTION
              The link utility shall perform the function call:
21094
              link(file1, file2);
21095
21096
              A user may need appropriate privilege to invoke the link utility.
21097 OPTIONS
              None.
21098
21099 OPERANDS
21100
              The following operands shall be supported:
              file1
                           The pathname of an existing file.
21101
              file2
21102
                           The pathname of the new directory entry to be created.
21103 STDIN
21104
              Not used.
21105 INPUT FILES
              Not used.
21106
21107 ENVIRONMENT VARIABLES
              The following environment variables shall affect the execution of link:
21108
21109
              LANG
                           Provide a default value for the internationalization variables that are unset or null.
                           (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
21110
                           Internationalization Variables for the precedence of internationalization variables
91111
21112
                           used to determine the values of locale categories.)
              LC ALL
                           If set to a non-empty string value, override the values of all the other
21113
21114
                           internationalization variables.
                           Determine the locale for the interpretation of sequences of bytes of text data as
21115
              LC_CTYPE
21116
                           characters (for example, single-byte as opposed to multi-byte characters in
21117
                           arguments).
21118
              LC_MESSAGES
                           Determine the locale that should be used to affect the format and contents of
21119
21120
                           diagnostic messages written to standard error.
              NLSPATH
                           Determine the location of message catalogs for the processing of LC_MESSAGES.
21121
21122 ASYNCHRONOUS EVENTS
21123
              Default.
21124 STDOUT
              None.
21125
21126 STDERR
```

547

The standard error shall be used only for diagnostic messages.

21127

link Utilities

21128 OUTPUT FILES 21129 None. 21130 EXTENDED DESCRIPTION 21131 None. 21132 EXIT STATUS 21133 The following exit values shall be returned: 0 Successful completion. 21134 >0 An error occurred. 21135 21136 CONSEQUENCES OF ERRORS Default. 21138 APPLICATION USAGE 21139 None. 21140 EXAMPLES None. 21141 21142 RATIONALE 21143 None. 21144 FUTURE DIRECTIONS None. 21145 21146 SEE ALSO In, unlink, the System Interfaces volume of IEEE Std 1003.1-2001, link() 21147

First released in Issue 5.

Utilities ln

DESCRIPTION

In the first synopsis form, the *In* utility shall create a new directory entry (link) at the destination path specified by the *target_file* operand. If the –s option is specified, a symbolic link shall be created for the file specified by the *source_file* operand. This first synopsis form shall be assumed when the final operand does not name an existing directory; if more than two operands are specified and the final is not an existing directory, an error shall result.

In the second synopsis form, the *ln* utility shall create a new directory entry (link), or if the **-s** option is specified a symbolic link, for each file specified by a *source_file* operand, at a destination path in the existing directory named by *target_dir*.

If the last operand specifies an existing file of a type not specified by the System Interfaces volume of IEEE Std 1003.1-2001, the behavior is implementation-defined.

The corresponding destination path for each *source_file* shall be the concatenation of the target directory pathname, a slash character, and the last pathname component of the *source_file*. The second synopsis form shall be assumed when the final operand names an existing directory.

For each *source_file*:

- 1. If the destination path exists:
 - a. If the **–f** option is not specified, *ln* shall write a diagnostic message to standard error, do nothing more with the current *source_file*, and go on to any remaining *source_files*.
 - b. Actions shall be performed equivalent to the *unlink()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001, called using *destination* as the *path* argument. If this fails for any reason, *ln* shall write a diagnostic message to standard error, do nothing more with the current *source_file*, and go on to any remaining *source_files*.
- 2. If the **–s** option is specified, *In* shall create a symbolic link named by the destination path and containing as its pathname *source_file*. The *In* utility shall do nothing more with *source_file* and shall go on to any remaining files.
- 3. If <code>source_file</code> is a symbolic link, actions shall be performed equivalent to the <code>link()</code> function using the object that <code>source_file</code> references as the <code>path1</code> argument and the destination path as the <code>path2</code> argument. The <code>ln</code> utility shall do nothing more with <code>source_file</code> and shall go on to any remaining files.
- 4. Actions shall be performed equivalent to the *link()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001 using *source_file* as the *path1* argument, and the destination path as the *path2* argument.

21188 OPTIONS

The *In* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

21191 The following option shall be supported:

21192 — **f** Force existing destination pathnames to be removed to allow the link.

ln Utilities

21193	- s	Create symbolic links instead of hard links.
21194 OPERA		
21195	The followin	g operands shall be supported:
21196 21197 21198	source_file	A pathname of a file to be linked. If the $-s$ option is specified, no restrictions on the type of file or on its existence shall be made. If the $-s$ option is not specified, whether a directory can be linked is implementation-defined.
21199	target_file	The pathname of the new directory entry to be created.
21200	target_dir	A pathname of an existing directory in which the new directory entries are created.
21201 STDIN 21202	Not used.	
21203 INPUT 21204	FILES None.	
21205 ENVIR	ONMENT VA	ARIABLES
21206		g environment variables shall affect the execution of <i>ln</i> :
21207 21208 21209 21210	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
21211 21212	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
21213 21214 21215	LC_CTYPE	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).
21216	LC_MESSAC	GES
21217 21218		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
21219 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .
21220 ASYNC 21221	HRONOUS I Default.	EVENTS
21222 STDOU	${f T}$	
21223	Not used.	
21224 STDER 21225		d error shall be used only for diagnostic messages.
21226 OUTPU 21227	T FILES None.	
21228 EXTEN	DED DESCR	IPTION
21229	None.	
21230 EXIT ST		or arity ralivas aball be notiumed.
21231		ag exit values shall be returned:
21232		specified files were linked successfully.
01000	> 0 Am anno	n o o o uma d

21233

>0 An error occurred.

Utilities ln

21234 CONSEQUENCES OF ERRORS

21235 Default.

21236 APPLICATION USAGE

21237 None.

21238 EXAMPLES

21239 None.

21240 RATIONALE

21241

21242

21243

21244 21245

21246

21247

21248

21249

21250

21251

2125221253

21254

21255

21256 21257

21258 21259

21260

21261

2126221263

21264 21265

21266

21267

Some historic versions of *In* (including the one specified by the SVID) unlink the destination file, if it exists, by default. If the mode does not permit writing, these versions prompt for confirmation before attempting the unlink. In these versions the **-f** option causes *In* not to attempt to prompt for confirmation.

This allows *In* to succeed in creating links when the target file already exists, even if the file itself is not writable (although the directory must be). Early proposals specified this functionality.

This volume of IEEE Std 1003.1-2001 does not allow the *ln* utility to unlink existing destination paths by default for the following reasons:

- The *In* utility has historically been used to provide locking for shell applications, a usage that is incompatible with *In* unlinking the destination path by default. There was no corresponding technical advantage to adding this functionality.
- This functionality gave *In* the ability to destroy the link structure of files, which changes the historical behavior of *In*.
- This functionality is easily replicated with a combination of *rm* and *ln*.
- It is not historical practice in many systems; BSD and BSD-derived systems do not support this behavior. Unfortunately, whichever behavior is selected can cause scripts written expecting the other behavior to fail.
- It is preferable that *In* perform in the same manner as the *link()* function, which does not permit the target to exist already.

This volume of IEEE Std 1003.1-2001 retains the **–f** option to provide support for shell scripts depending on the SVID semantics. It seems likely that shell scripts would not be written to handle prompting by *In* and would therefore have specified the **–f** option.

The **–f** option is an undocumented feature of many historical versions of the *ln* utility, allowing linking to directories. These versions require modification.

Early proposals of this volume of IEEE Std 1003.1-2001 also required a -i option, which behaved like the -i options in *cp* and *mv*, prompting for confirmation before unlinking existing files. This was not historical practice for the *ln* utility and has been omitted.

21268 FUTURE DIRECTIONS

21269 None.

21270 SEE ALSO

chmod, find, pax, rm, the System Interfaces volume of IEEE Std 1003.1-2001, link(), unlink()

21272 CHANGE HISTORY

First released in Issue 2.

ln Utilities

21274 Issue 6

The *ln* utility is updated to include symbolic link processing as defined in the IEEE P1003.2b draft standard.

Utilities locale

21277 NAME	1 1	
21278		t locale-specific information
21279 SYNOF 21280	'SIS locale [–	.a _m1
21281		ck] name
21282 DESCR 21283		atility shall write information about the current locale environment, or all public
21284 21285	locales, to th	ne standard output. For the purposes of this section, a <i>public locale</i> is one provided by entation that is accessible to the application.
21286 21287 21288	environmen	e is invoked without any arguments, it shall summarize the current locale at for each locale category as determined by the settings of the environment variables the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 7, Locale.
21289 21290		ked with operands, it shall write values that have been assigned to the keywords in itegories, as follows:
21291 21292	 Specifying keyword 	ng a keyword name shall select the named keyword and the category containing that l.
21293 21294	 Specifyir category 	ng a category name shall select the named category and all keywords in that
21295 OPTIO	NS	
21296 21297		tility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section Syntax Guidelines.
21298	The following	ng options shall be supported:
21299 21300 21301 21302	- a	Write information about all available public locales. The available locales shall include POSIX , representing the POSIX locale. The manner in which the implementation determines what other locales are available is implementation-defined.
21303 21304 21305 21306	-с	Write the names of selected locale categories; see the STDOUT section. The $-c$ option increases readability when more than one category is selected (for example, via more than one keyword name or via a category name). It is valid both with and without the $-k$ option.
21307 21308	-k	Write the names and values of selected keywords. The implementation may omit values for some keywords; see the OPERANDS section.
21309 21310	- m	Write names of available charmaps; see the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.1, Portable Character Set.
21311 OPERA		
21312	The following	ng operand shall be supported:
21313	name	The name of a locale category as defined in the Base Definitions volume of IEEE Std 1003 1 2001. Chapter 7 I locale the name of a keyword in a locale
21314 21315		IEEE Std 1003.1-2001, Chapter 7, Locale, the name of a keyword in a locale category, or the reserved name charmap . The named category or keyword shall be
21316		selected for output. If a single <i>name</i> represents both a locale category name and a
21317		keyword name in the current locale, the results are unspecified. Otherwise, both

categories *LC_CTYPE* and *LC_COLLATE*.

category and keyword names can be specified as name operands, in any sequence.

It is implementation-defined whether any keyword values are written for the

21317 21318

21319

21320

locale

21321 STDIN	I	
21322	Not used.	
21323 INPUT 21324	FILES None.	
21325 ENVIR 21326	RONMENT VA The followin	ARIABLES ag environment variables shall affect the execution of <i>locale</i> :
21327 21328 21329 21330	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
21331 21332	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
21333 21334 21335	LC_CTYPE	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).
21336 21337 21338	LC_MESSA	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
21339 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
21340 XSI 21341 21342		ion shall ensure that the $LANG$, LC_{-}^* , and $NLSPATH$ environment variables specify locale environment to be written out; they shall be used if the $-\mathbf{a}$ option is not
21343 ASYN 0 21344	CHRONOUS I Default.	EVENTS
	Default.	EVENTS
21344	Default. UT If <i>locale</i> is in LC_* environ the standard format. Only	events avoked without any options or operands, the names and values of the <i>LANG</i> and nament variables described in this volume of IEEE Std 1003.1-2001 shall be written to a output, one variable per line, with <i>LANG</i> first, and each line using the following by those variables set in the environment and not overridden by <i>LC_ALL</i> shall be go this format:
21344 21345 STDO 21346 21347 21348 21349	Default. UT If <i>locale</i> is in <i>LC_*</i> envirouthe standard format. Only written usin	avoked without any options or operands, the names and values of the <i>LANG</i> and nament variables described in this volume of IEEE Std 1003.1-2001 shall be written to a loutput, one variable per line, with <i>LANG</i> first, and each line using the following by those variables set in the environment and not overridden by <i>LC_ALL</i> shall be
21344 21345 STDO 21346 21347 21348 21349 21350	Default. UT If locale is in LC_* environ the standard format. Only written usin "%s=%s\n" The names of IEEE Std 100	avoked without any options or operands, the names and values of the <i>LANG</i> and nament variables described in this volume of IEEE Std 1003.1-2001 shall be written to doutput, one variable per line, with <i>LANG</i> first, and each line using the following y those variables set in the environment and not overridden by <i>LC_ALL</i> shall be g this format:
21344 21345 STDO 21346 21347 21348 21349 21350 21351 21352 21353	Default. UT If locale is in LC_* environ the standard format. Only written usin "%s=%s\n" The names of IEEE Std 100 written in the	avoked without any options or operands, the names and values of the <i>LANG</i> and nament variables described in this volume of IEEE Std 1003.1-2001 shall be written to doutput, one variable per line, with <i>LANG</i> first, and each line using the following by those variables set in the environment and not overridden by <i>LC_ALL</i> shall be githis format: Variable_name Value
21344 21345 STDO 21346 21347 21348 21349 21350 21351 21352 21353 21354	Default. UT If locale is in LC_* environ the standard format. Only written usin, "%s=%s\n" The names of IEEE Std 100 written in the "%s=\"%s\ The <implies implementations.<="" td=""><td>twoked without any options or operands, the names and values of the $LANG$ and nament variables described in this volume of IEEE Std 1003.1-2001 shall be written to doutput, one variable per line, with $LANG$ first, and each line using the following y those variables set in the environment and not overridden by LC_ALL shall be githis format: variable_name>, <value> of those LC_* variables associated with locale categories defined in this volume of 13.1-2001 that are not set in the environment or are overridden by LC_ALL shall be the following format:</value></td></implies>	twoked without any options or operands, the names and values of the $LANG$ and nament variables described in this volume of IEEE Std 1003.1-2001 shall be written to doutput, one variable per line, with $LANG$ first, and each line using the following y those variables set in the environment and not overridden by LC_ALL shall be githis format: variable_name>, <value> of those LC_* variables associated with locale categories defined in this volume of 13.1-2001 that are not set in the environment or are overridden by LC_ALL shall be the following format:</value>
21344 21345 STDO 21346 21347 21348 21349 21350 21351 21352 21353 21354 21355 21356 21357	Default. UT If locale is in LC_* environ the standard format. Only written usin "%s=%s\n" The names of IEEE Std 100 written in the "%s=\"%s\ The <implies implementation="" shell."<="" td="" the="" to=""><td>avoked without any options or operands, the names and values of the <i>LANG</i> and ament variables described in this volume of IEEE Std 1003.1-2001 shall be written to doutput, one variable per line, with <i>LANG</i> first, and each line using the following by those variables set in the environment and not overridden by <i>LC_ALL</i> shall be githis format: (**variable_name**), **value** of those *LC_** variables associated with locale categories defined in this volume of 13.1-2001 that are not set in the environment or are overridden by <i>LC_ALL</i> shall be given format: (**variable_name**), **implied value** (**value**) shall be the name of the locale that has been selected for that category by the tion, based on the values in <i>LANG</i> and <i>LC_ALL</i>, as described in the Base Definitions</td></implies>	avoked without any options or operands, the names and values of the <i>LANG</i> and ament variables described in this volume of IEEE Std 1003.1-2001 shall be written to doutput, one variable per line, with <i>LANG</i> first, and each line using the following by those variables set in the environment and not overridden by <i>LC_ALL</i> shall be githis format: (**variable_name**), **value** of those *LC_** variables associated with locale categories defined in this volume of 13.1-2001 that are not set in the environment or are overridden by <i>LC_ALL</i> shall be given format: (**variable_name**), **implied value** (**value**) shall be the name of the locale that has been selected for that category by the tion, based on the values in <i>LANG</i> and <i>LC_ALL</i> , as described in the Base Definitions

Utilities locale

```
21364
              "LC ALL=\n"
21365
              If any arguments are specified:
                1. If the -a option is specified, the names of all the public locales shall be written, each in the
21366
21367
                   following format:
                    "%s\n", <locale name>
21368
                2. If the –c option is specified, the names of all selected categories shall be written, each in the
21369
                   following format:
21370
21371
                    "%s\n", <category name>
                   If keywords are also selected for writing (see following items), the category name output
21372
21373
                   shall precede the keyword output for that category.
                   If the -c option is not specified, the names of the categories shall not be written; only the
21374
                   keywords, as selected by the <name> operand, shall be written.
21375
                3. If the -k option is specified, the names and values of selected keywords shall be written. If
21376
                    a value is non-numeric, it shall be written in the following format:
21377
21378
                    "%s=\"%s\"\n", <keyword name>, <keyword value>
                   If the keyword was charmap, the name of the charmap (if any) that was specified via the
21379
                    localedef – f option when the locale was created shall be written, with the word charmap as
21380
                    <keyword name>.
21381
                   If a value is numeric, it shall be written in one of the following formats:
21382
                    "%s=%d\n", <keyword name>, <keyword value>
21383
                    "%s=%c%o\n", <keyword name>, <escape character>, <keyword value>
21384
                    "%s=%cx%x\n", <keyword name>, <escape character>, <keyword value>
21385
                    where the <escape character> is that identified by the escape_char keyword in the current
21386
                   locale; see the Base Definitions volume of IEEE Std 1003.1-2001, Section 7.3, Locale
21387
                   Definition.
21388
                    Compound keyword values (list entries) shall be separated in the output by semicolons.
21389
21390
                    When included in keyword values, the semicolon, the double-quote, the backslash, and
21391
                    any control character shall be preceded (escaped) with the escape character.
21392
                4. If the -k option is not specified, selected keyword values shall be written, each in the
21393
                   following format:
                    "%s\n", <keyword value>
21394
                   If the keyword was charmap, the name of the charmap (if any) that was specified via the
21395
                   localedef – f option when the locale was created shall be written.
21396
21397
                5. If the -\mathbf{m} option is specified, then a list of all available charmaps shall be written, each in
                    the format:
21398
                    "%s\n", <charmap>
21399
```

where *<charmap>* is in a format suitable for use as the option-argument to the *localedef* –f

21400

21401

option.

locale Utilities

```
21402 STDERR
21403
             The standard error shall be used only for diagnostic messages.
21404 OUTPUT FILES
             None.
21405
21406 EXTENDED DESCRIPTION
             None.
21407
21408 EXIT STATUS
             The following exit values shall be returned:
21409
21410
                 All the requested information was found and output successfully.
                An error occurred.
21412 CONSEQUENCES OF ERRORS
21413
             Default.
21414 APPLICATION USAGE
21415
             If the LANG environment variable is not set or set to an empty value, or one of the LC_{-}^{*}
21416
             environment variables is set to an unrecognized value, the actual locales assumed (if any) are
21417
             implementation-defined as described in the Base Definitions volume of IEEE Std 1003.1-2001,
             Chapter 8, Environment Variables.
21418
21419
             Implementations are not required to write out the actual values for keywords in the categories
             LC_CTYPE and LC_COLLATE; however, they must write out the categories (allowing an
21420
21421
             application to determine, for example, which character classes are available).
21422 EXAMPLES
21423
             In the following examples, the assumption is that locale environment variables are set as
21424
             follows:
21425
             LANG=locale x
21426
             LC COLLATE=locale y
21427
             The command locale would result in the following output:
21428
             LANG=locale x
             LC CTYPE="locale x"
21429
21430
             LC_COLLATE=locale_y
21431
             LC TIME="locale x"
             LC NUMERIC="locale x"
21432
             LC MONETARY="locale x"
21433
             LC MESSAGES="locale x"
21434
21435
             LC ALL=
21436
             The order of presentation of the categories is not specified by this volume of
             IEEE Std 1003.1-2001.
21437
             The command:
21438
             LC_ALL=POSIX locale -ck decimal_point
21439
21440
             would produce:
             LC NUMERIC
21441
21442
             decimal point="."
21443
             The following command shows an application of locale to determine whether a user-supplied
```

21444

response is affirmative:

locale **Utilities**

```
21445
              if printf "%s\n" "$response" | grep -Eq "$(locale yesexpr)"
21446
             then
21447
                   affirmative processing goes here
21448
             else
21449
                   non-affirmative processing goes here
21450
             fi
21451 RATIONALE
21452
             The output for categories LC_CTYPE and LC_COLLATE has been made implementation-defined
21453
             because there is a questionable value in having a shell script receive an entire array of characters.
21454
             It is also difficult to return a logical collation description, short of returning a complete localedef
21455
             source.
21456
             The -m option was included to allow applications to query for the existence of charmaps. The
             output is a list of the charmaps (implementation-supplied and user-supplied, if any) on the
21457
             system.
21458
             The -c option was included for readability when more than one category is selected (for
21459
             example, via more than one keyword name or via a category name). It is valid both with and
21460
21461
             without the -\mathbf{k} option.
             The charmap keyword, which returns the name of the charmap (if any) that was used when the
21462
21463
             current locale was created, was included to allow applications needing the information to
             retrieve it.
21464
21465 FUTURE DIRECTIONS
21466
             None.
21467 SEE ALSO
             localedef, the Base Definitions volume of IEEE Std 1003.1-2001, Section 7.3, Locale Definition
21468
21469 CHANGE HISTORY
21470
             First released in Issue 4.
21471 Issue 5
21472
             The FUTURE DIRECTIONS section is added.
21473 Issue 6
21474
             The normative text is reworded to avoid use of the term "must" for application requirements.
21475
             IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/30 is applied, correcting an editorial error
             in the STDOUT section.
```

21476

localedef Utilities

or and NIANGE		
21477 NAME 21478		define locale environment
		deline locale environment
21479 SYNOI 21480		[-c][-f charmap][-i sourcefile][-u code set name] name
		[c][I charmap][I boarcerire][u code_bee_name] name
21481 DESCE		utility shall convert source definitions for locale categories into a format usable by
21482 21483		s and utilities whose operational behavior is determined by the setting of the locale
21484		t variables defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter
21485	7, Locale. It	is implementation-defined whether users have the capability to create new locales,
21486		to those supplied by the implementation. If the symbolic constant
21487 XSI		CALEDEF is defined, the system supports the creation of new locales. On XSI-
21488		systems, the symbolic constant POSIX2_LOCALEDEF shall be defined.
21489 21490		hall read source definitions for one or more locale categories belonging to the same he file named in the – i option (if specified) or from standard input.
21491	The <i>name</i> op	perand identifies the target locale. The utility shall support the creation of <i>public</i> , or
21492		cessible locales, as well as <i>private</i> , or restricted-access locales. Implementations may
21493	restrict the ca	apability to create or modify public locales to users with the appropriate privileges.
21494	Each categor	ry source definition shall be identified by the corresponding environment variable
21495		erminated by an END category-name statement. The following categories shall be
21496	supported. I	n addition, the input may contain source for implementation-defined categories.
21497	LC_CTYPE	Defines character classification and case conversion.
21498	LC_COLLAT	
21499		Defines collation rules.
21500	LC_MONETA	
21501		Defines the format and symbols used in formatting of monetary information.
21502	LC_NUMER	IC
21503		Defines the decimal delimiter, grouping, and grouping symbol for non-monetary
21504		numeric editing.
21505	LC_TIME	Defines the format and content of date and time information.
21506	LC_MESSAC	
21507		Defines the format and values of affirmative and negative responses.
21508 OPTIO	NS	
21509		utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
21510	12.2, Utility S	Syntax Guidelines.
21511	The followin	g options shall be supported:
21512	-с	Create permanent output even if warning messages have been issued.
21513	− f charmap	Specify the pathname of a file containing a mapping of character symbols and
21514		collating element symbols to actual character encodings. The format of the
21515		charmap is described in the Base Definitions volume of IEEE Std 1003.1-2001,
21516		Section 6.4, Character Set Description File. The application shall ensure that this
21517		option is specified if symbolic names (other than collating symbols defined in a
21518		collating-symbol keyword) are used. If the -f option is not present, an implementation defined character mapping shall be used

implementation-defined character mapping shall be used.

21519

Utilities localedef

21520	–i inputfile	The pathname of a file containing the source definitions. If this option is not
21521		present, source definitions shall be read from standard input. The format of the
21522		inputfile is described in the Base Definitions volume of IEEE Std 1003.1-2001,
21523		Section 7.3, Locale Definition.
21524	-u code_set_1	name
21525		Specify the name of a codeset used as the target mapping of character symbols and
21526		collating element symbols whose encoding values are defined in terms of the
21527		ISO/IEC 10646-1: 2000 standard position constant values.
21528 OPERA	NDS	

21529 The following operand shall be supported:

Identifies the locale; see the Base Definitions volume of IEEE Std 1003.1-2001, 21530 name Chapter 7, Locale for a description of the use of this name. If the name contains one 21531 or more slash characters, name shall be interpreted as a pathname where the 21532 created locale definitions shall be stored. If name does not contain any slash 21533 characters, the interpretation of the name is implementation-defined and the locale 21534 shall be public. This capability may be restricted to users with appropriate 21535 21536 privileges. (As a consequence of specifying one name, although several categories 21537 can be processed in one execution, only categories belonging to the same locale can 21538 be processed.)

21539 **STDIN**

21540 21541

21542 21543

21545

21546

21547

21548 21549

21550

21552

Unless the **–i** option is specified, the standard input shall be a text file containing one or more locale category source definitions, as described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 7.3, Locale Definition. When lines are continued using the escape character mechanism, there is no limit to the length of the accumulated continued line.

21544 INPUT FILES

The character set mapping file specified as the *charmap* option-argument is described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.4, Character Set Description File. If a locale category source definition contains a **copy** statement, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 7, Locale, and the **copy** statement names a valid, existing locale, then *localedef* shall behave as if the source definition had contained a valid category source definition for the named locale.

21551 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *localedef*:

21553 LANG
Provide a default value for the internationalization variables that are unset or null.
21554 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
21555 Internationalization Variables for the precedence of internationalization variables
21556 used to determine the values of locale categories.)

21557 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

21559 LC_COLLATE

21560 (This variable has no affect on *localedef*; the POSIX locale is used for this category.)

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). This variable has no affect on the processing of *localedef* input data; the POSIX locale is used for this purpose, regardless of the value of this variable.

localedefUtilities

21566 LC_MESSAGES

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

21569 XSI NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

21570 ASYNCHRONOUS EVENTS

21571 Default.

21572 STDOUT

21573 The utility shall report all categories successfully processed, in an unspecified format.

21574 STDERR

21580

21581

21582

21583

21584

21585 21586

21590

2159121592

21596

21597

21604

21607

21608

21575 The standard error shall be used only for diagnostic messages.

21576 OUTPUT FILES

The format of the created output is unspecified. If the *name* operand does not contain a slash, the existence of an output file for the locale is unspecified.

21579 EXTENDED DESCRIPTION

When the **–u** option is used, the *code_set_name* option-argument shall be interpreted as an implementation-defined name of a codeset to which the ISO/IEC 10646-1:2000 standard position constant values shall be converted via an implementation-defined method. Both the ISO/IEC 10646-1:2000 standard position constant values and other formats (decimal, hexadecimal, or octal) shall be valid as encoding values within the *charmap* file. The codeset represented by the implementation-defined name can be any codeset that is supported by the implementation.

When conflicts occur between the *charmap* specification of *<code_set_name>*, *<mb_cur_max>*, or *<mb_cur_min>* and the implementation-defined interpretation of these respective items for the codeset represented by the **–u** option-argument *code_set_name*, the result is unspecified.

When conflicts occur between the *charmap* encoding values specified for symbolic names of characters of the portable character set and the implementation-defined assignment of character encoding values, the result is unspecified.

If a non-printable character in the *charmap* has a width specified that is not **-1**, *localedef* shall generate a warning.

21595 EXIT STATUS

The following exit values shall be returned:

- 0 No errors occurred and the locales were successfully created.
- 21598 1 Warnings occurred and the locales were successfully created.
- 2 The locale specification exceeded implementation limits or the coded character set or sets used were not supported by the implementation, and no locale was created.
- The capability to create new locales is not supported by the implementation.
- 21602 >3 Warnings or errors occurred and no output was created.

21603 CONSEQUENCES OF ERRORS

If an error is detected, no permanent output shall be created.

If warnings occur, permanent output shall be created if the **-c** option was specified. The following conditions shall cause warning messages to be issued:

If a symbolic name not found in the *charmap* file is used for the descriptions of the *LC_CTYPE* or *LC_COLLATE* categories (for other categories, this shall be an error condition).

Utilities localedef

- If the number of operands to the **order** keyword exceeds the {COLL_WEIGHTS_MAX} limit.
- If optional keywords not supported by the implementation are present in the source.
- If a non-printable character has a width specified other than −1.
- 21612 Other implementation-defined conditions may also cause warnings.

21613 APPLICATION USAGE

The *charmap* definition is optional, and is contained outside the locale definition. This allows both completely self-defined source files, and generic sources (applicable to more than one codeset). To aid portability, all *charmap* definitions must use the same symbolic names for the portable character set. As explained in the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.4, Character Set Description File, it is implementation-defined whether or not users or applications can provide additional character set description files. Therefore, the –f option might be operable only when an implementation-defined *charmap* is named.

21621 EXAMPLES

21622 None.

21623 RATIONALE

The output produced by the *localedef* utility is implementation-defined. The *name* operand is used to identify the specific locale. (As a consequence, although several categories can be processed in one execution, only categories belonging to the same locale can be processed.)

21627 FUTURE DIRECTIONS

21628 None.

21629 **SEE ALSO**

locale, the Base Definitions volume of IEEE Std 1003.1-2001, Section 7.3, Locale Definition

21631 CHANGE HISTORY

First released in Issue 4.

21633 Issue 6

21634 The -**u** option is added, as specified in the IEEE P1003.2b draft standard.

The normative text is reworded to avoid use of the term "must" for application requirements.

logger Utilities

21636 NAME 21637 logger — log messages 21638 SYNOPSIS 21639 logger string ... 21640 DESCRIPTION 21641 The *logger* utility saves a message, in an unspecified manner and format, containing the *string* 21642 operands provided by the user. The messages are expected to be evaluated later by personnel 21643 performing system administration tasks. 21644 It is implementation-defined whether messages written in locales other than the POSIX locale 21645 are effective. 21646 OPTIONS None. 21647 21648 OPERANDS The following operand shall be supported: 21649 string One of the string arguments whose contents are concatenated together, in the 21650 order specified, separated by single <space>s. 21651 21652 **STDIN** 21653 Not used. 21654 INPUT FILES None. 21655 21656 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *logger*: 21657 21658 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 21659 Internationalization Variables for the precedence of internationalization variables 21660 21661 used to determine the values of locale categories.) LC ALL 21662 If set to a non-empty string value, override the values of all the other 21663 internationalization variables. 21664 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 21665 characters (for example, single-byte as opposed to multi-byte characters in arguments). 21666 21667 LC_MESSAGES Determine the locale that should be used to affect the format and contents of 21668 diagnostic messages written to standard error. (This means diagnostics from logger 21669 to the user or application, not diagnostic messages that the user is sending to the 21670 21671 system administrator.) NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 21672 XSI 21673 ASYNCHRONOUS EVENTS Default. 21674 21675 STDOUT

21676

Not used.

Utilities logger

21677 STDERR

The standard error shall be used only for diagnostic messages.

21679 OUTPUT FILES

21680 Unspecified.

21681 EXTENDED DESCRIPTION

21682 None.

21683 EXIT STATUS

The following exit values shall be returned:

21685 0 Successful completion.

21686 >0 An error occurred.

21687 CONSEQUENCES OF ERRORS

21688 Default.

21689 APPLICATION USAGE

This utility allows logging of information for later use by a system administrator or programmer in determining why non-interactive utilities have failed. The locations of the saved messages, their format, and retention period are all unspecified. There is no method for a conforming application to read messages, once written.

21694 EXAMPLES

A batch application, running non-interactively, tries to read a configuration file and fails; it may attempt to notify the system administrator with:

logger myname: unable to read file foo. [timestamp]

21698 RATIONALE

21699

21700

2170121702

21703

21704

The standard developers believed strongly that some method of alerting administrators to errors was necessary. The obvious example is a batch utility, running non-interactively, that is unable to read its configuration files or that is unable to create or write its results file. However, the standard developers did not wish to define the format or delivery mechanisms as they have historically been (and will probably continue to be) very system-specific, as well as involving functionality clearly outside the scope of this volume of IEEE Std 1003.1-2001.

The text with $LC_MESSAGES$ about diagnostic messages means diagnostics from logger to the user or application, not diagnostic messages that the user is sending to the system administrator.

21707 Multiple *string* arguments are allowed, similar to *echo*, for ease-of-use.

Like the utilities *mailx* and *lp, logger* is admittedly difficult to test. This was not deemed sufficient justification to exclude these utilities from this volume of IEEE Std 1003.1-2001. It is also arguable that they are, in fact, testable, but that the tests themselves are not portable.

21711 FUTURE DIRECTIONS

21712 None.

21713 SEE ALSO

21714 lp, mailx, write

21715 CHANGE HISTORY

First released in Issue 4.

logname Utilities

```
21717 NAME
21718
              logname — return the user's login name
21719 SYNOPSIS
21720
              logname
21721 DESCRIPTION
21722
              The logname utility shall write the user's login name to standard output. The login name shall be
              the string that would be returned by the getlogin() function defined in the System Interfaces
21723
              volume of IEEE Std 1003.1-2001. Under the conditions where the getlogin() function would fail,
21724
              the logname utility shall write a diagnostic message to standard error and exit with a non-zero
21725
21726
              exit status.
21727 OPTIONS
21728
              None.
21729 OPERANDS
21730
              None.
21731 STDIN
              Not used.
21732
21733 INPUT FILES
21734
              None.
21735 ENVIRONMENT VARIABLES
21736
              The following environment variables shall affect the execution of logname:
              LANG
                           Provide a default value for the internationalization variables that are unset or null.
21737
21738
                           (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
21739
                           Internationalization Variables for the precedence of internationalization variables
21740
                           used to determine the values of locale categories.)
21741
              LC_ALL
                           If set to a non-empty string value, override the values of all the other
                           internationalization variables.
21742
                           Determine the locale for the interpretation of sequences of bytes of text data as
              LC_CTYPE
21743
                           characters (for example, single-byte as opposed to multi-byte characters in
21744
                           arguments).
21745
              LC_MESSAGES
21746
                           Determine the locale that should be used to affect the format and contents of
21747
                           diagnostic messages written to standard error.
21748
              NLSPATH
                           Determine the location of message catalogs for the processing of LC_MESSAGES.
21749 XSI
21750 ASYNCHRONOUS EVENTS
              Default.
21751
21752 STDOUT
              The logname utility output shall be a single line consisting of the user's login name:
21753
21754
              "%s\n", <login name>
21755 STDERR
```

The standard error shall be used only for diagnostic messages.

21756

Utilities logname

21757 OUTPUT FILES

21758 None.

21759 EXTENDED DESCRIPTION

21760 None.

21761 EXIT STATUS

21762 The following exit values shall be returned:

21763 0 Successful completion.

>0 An error occurred.

21765 CONSEQUENCES OF ERRORS

21766 Default.

21767 APPLICATION USAGE

21768 The *logname* utility explicitly ignores the *LOGNAME* environment variable because environment

21769 changes could produce erroneous results.

21770 EXAMPLES

21771 None.

21772 RATIONALE

21773 The **passwd** file is not listed as required because the implementation may have other means of

21774 mapping login names.

21775 FUTURE DIRECTIONS

21776 None.

21777 SEE ALSO

id, who, the System Interfaces volume of IEEE Std 1003.1-2001, getlogin()

21779 CHANGE HISTORY

First released in Issue 2.

lp Utilities

```
21781 NAME
21782
              lp — send files to a printer
21784
              lp [-c] [-d dest] [-n copies] [-msw] [-o option] ... [-t title] [file...]
21785 DESCRIPTION
              The lp utility shall copy the input files to an output destination in an unspecified manner. The
21786
              default output destination should be to a hardcopy device, such as a printer or microfilm
21787
              recorder, that produces non-volatile, human-readable documents. If such a device is not
21788
21789
              available to the application, or if the system provides no such device, the lp utility shall exit with
21790
              a non-zero exit status.
              The actual writing to the output device may occur some time after the lp utility successfully
21791
21792
              exits. During the portion of the writing that corresponds to each input file, the implementation
              shall guarantee exclusive access to the device.
21793
21794
              The lp utility shall associate a unique request ID with each request.
              Normally, a banner page is produced to separate and identify each print job. This page may be
21795
21796
              suppressed by implementation-defined conditions, such as an operator command or one of the
21797
              −o option values.
21798 OPTIONS
              The lp utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2,
21799
              Utility Syntax Guidelines.
21800
              The following options shall be supported:
21801
21802
                            Exit only after further access to any of the input files is no longer required. The
              -C
                            application can then safely delete or modify the files without affecting the output
21803
                            operation. Normally, files are not copied, but are linked whenever possible. If the
21804
                            -c option is not given, then the user should be careful not to remove any of the
21805
                            files before the request has been printed in its entirety. It should also be noted that
21806
21807
                            in the absence of the -c option, any changes made to the named files after the
                            request is made but before it is printed may be reflected in the printed output. On
21808
                            some implementations, -c may be on by default.
21809
              -d dest
                            Specify a string that names the destination (dest). If dest is a printer, the request
21810
                            shall be printed only on that specific printer. If dest is a class of printers, the request
21811
                            shall be printed on the first available printer that is a member of the class. Under
21812
21813
                            certain conditions (printer unavailability, file space limitation, and so on), requests
21814
                            for specific destinations need not be accepted. Destination names vary between
                            systems.
21815
                            If -\mathbf{d} is not specified, and neither the LPDEST nor PRINTER environment variable
21816
                            is set, an unspecified destination is used. The -d dest option shall take precedence
21817
                            over LPDEST, which in turn shall take precedence over PRINTER. Results are
21818
                            undefined when dest contains a value that is not a valid destination name.
21819
                            Send mail (see mailx) after the files have been printed. By default, no mail is sent
21820
              -m
                            upon normal completion of the print request.
21821
                            Write copies number of copies of the files, where copies is a positive decimal integer.
21822
              -n copies
                            The methods for producing multiple copies and for arranging the multiple copies
21823
                            when multiple file operands are used are unspecified, except that each file shall be
21824
```

21825

output as an integral whole, not interleaved with portions of other files.

Utilities lp

21826 21827	−o option	Specify printer-dependent or class-dependent <i>options</i> . Several such <i>options</i> may be collected by specifying the $-\mathbf{o}$ option more than once.
21828	-s	Suppress messages from lp.
21829	−t title	Write title on the banner page of the output.
21830 21831	- w	Write a message on the user's terminal after the files have been printed. If the user is not logged in, then mail shall be sent instead.
21832 OPER A 21833		ng operand shall be supported:
21834 21835 21836 21837 21838	file	A pathname of a file to be output. If no <i>file</i> operands are specified, or if a <i>file</i> operand is '-', the standard input shall be used. If a <i>file</i> operand is used, but the -c option is not specified, the process performing the writing to the output device may have user and group permissions that differ from that of the process invoking <i>lp</i> .
21839 STDIN 21840 21841	The standar	d input shall be used only if no <i>file</i> operands are specified, or if a <i>file</i> operand is $'-'$. UT FILES section.
21842 INPUT 21843		les shall be text files.
21844 ENVIR 21845	ONMENT VA The followin	ARIABLES ng environment variables shall affect the execution of <i>lp</i> :
21846 21847 21848 21849	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
21850 21851	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
21852 21853 21854	LC_CTYPE	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).
21855	LC_MESSA	
21856 21857 21858		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.
21859 21860	LC_TIME	Determine the format and contents of date and time strings displayed in the lp banner page, if any.
21861 21862 21863 21864	LPDEST	Determine the destination. If the <i>LPDEST</i> environment variable is not set, the <i>PRINTER</i> environment variable shall be used. The –d <i>dest</i> option takes precedence over <i>LPDEST</i> . Results are undefined when –d is not specified and <i>LPDEST</i> contains a value that is not a valid destination name.
21865 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
21866 21867 21868 21869	PRINTER	Determine the output device or destination. If the <i>LPDEST</i> and <i>PRINTER</i> environment variables are not set, an unspecified output device is used. The $-\mathbf{d}$ dest option and the <i>LPDEST</i> environment variable shall take precedence over <i>PRINTER</i> . Results are undefined when $-\mathbf{d}$ is not specified, <i>LPDEST</i> is unset, and

lp Utilities

21870 PRINTER contains a value that is not a valid device or destination name.

Determine the timezone used to calculate date and time strings displayed in the *lp*

banner page, if any. If TZ is unset or null, an unspecified default timezone shall be

21873 used.

21874 ASYNCHRONOUS EVENTS

21875 Default.

21876 **STDOUT**

The *lp* utility shall write a *request ID* to the standard output, unless –**s** is specified. The format of the message is unspecified. The request ID can be used on systems supporting the historical cancel and *lpstat* utilities.

21880 STDERR

The standard error shall be used only for diagnostic messages.

21882 OUTPUT FILES

21883 None.

21884 EXTENDED DESCRIPTION

21885 None.

21886 EXIT STATUS

The following exit values shall be returned:

21888 0 All input files were processed successfully.

21889 >0 No output device was available, or an error occurred.

21890 CONSEQUENCES OF ERRORS

21891 Default.

21892 APPLICATION USAGE

The *pr* and *fold* utilities can be used to achieve reasonable formatting for the implementation's default page size.

A conforming application can use one of the *file* operands only with the –c option or if the file is publicly readable and guaranteed to be available at the time of printing. This is because IEEE Std 1003.1-2001 gives the implementation the freedom to queue up the request for printing at some later time by a different process that might not be able to access the file.

21899 EXAMPLES

21895

21896

21897 21898

21901

21902

21900 1. To print file *file*:

2. To print multiple files with headers:

21904 RATIONALE

The *lp* utility was designed to be a basic version of a utility that is already available in many historical implementations. The standard developers considered that it should be implementable simply as:

after appropriate processing of options, if that is how the implementation chose to do it and if exclusive access could be granted (so that two users did not write to the device simultaneously). Although in the future the standard developers may add other options to this utility, it should Utilities lp

always be able to execute with no options or operands and send the standard input to an unspecified output device.

This volume of IEEE Std 1003.1-2001 makes no representations concerning the format of the printed output, except that it must be "human-readable" and "non-volatile". Thus, writing by default to a disk or tape drive or a display terminal would not qualify. (Such destinations are not prohibited when **–d** *dest*, *LPDEST*, or *PRINTER* are used, however.)

This volume of IEEE Std 1003.1-2001 is worded such that a "print job" consisting of multiple input files, possibly in multiple copies, is guaranteed to print so that any one file is not intermixed with another, but there is no statement that all the files or copies have to print out together.

The -c option may imply a spooling operation, but this is not required. The utility can be implemented to wait until the printer is ready and then wait until it is finished. Because of that, there is no attempt to define a queuing mechanism (priorities, classes of output, and so on).

On some historical systems, the request ID reported on the STDOUT can be used to later cancel or find the status of a request using utilities not defined in this volume of IEEE Std 1003.1-2001.

Although the historical System V lp and BSD lpr utilities have provided similar functionality, they used different names for the environment variable specifying the destination printer. Since the name of the utility here is lp, LPDEST (used by the System V lp utility) was given precedence over PRINTER (used by the BSD lpr utility). Since environments of users frequently contain one or the other environment variable, the lp utility is required to recognize both. If this was not done, many applications would send output to unexpected output devices when users moved from system to system.

Some have commented that lp has far too little functionality to make it worthwhile. Requests have proposed additional options or operands or both that added functionality. The requests included:

- Wording requiring the output to be "hardcopy"
- A requirement for multiple printers

Options for supporting various page-description languages

Given that a compliant system is not required to even have a printer, placing further restrictions upon the behavior of the printer is not useful. Since hardcopy format is so application-dependent, it is difficult, if not impossible, to select a reasonable subset of functionality that should be required on all compliant systems.

The term *unspecified* is used in this section in lieu of *implementation-defined* as most known implementations would not be able to make definitive statements in their conformance documents; the existence and usage of printers is very dependent on how the system administrator configures each individual system.

Since the default destination, device type, queuing mechanisms, and acceptable forms of input are all unspecified, usage guidelines for what a conforming application can do are as follows:

- Use the command in a pipeline, or with -c, so that there are no permission problems and the files can be safely deleted or modified.
- Limit output to text files of reasonable line lengths and printable characters and include no device-specific formatting information, such as a page description language. The meaning of "reasonable" in this context can only be answered as a quality-of-implementation issue, but it should be apparent from historical usage patterns in the industry and the locale. The *pr* and *fold* utilities can be used to achieve reasonable formatting for the default page size of the

lp Utilities

21957 implementation. 21958 Alternatively, the application can arrange its installation in such a way that it requires the 21959 system administrator or operator to provide the appropriate information on lp options and environment variable values. 21960 21961 At a minimum, having this utility in this volume of IEEE Std 1003.1-2001 tells the industry that conforming applications require a means to print output and provides at least a command name 21962 and LPDEST routing mechanism that can be used for discussions between vendors, application 21963 writers, and users. The use of "should" in the DESCRIPTION of lp clearly shows the intent of 21964 the standard developers, even if they cannot mandate that all systems (such as laptops) have 21965 21966 printers. This volume of IEEE Std 1003.1-2001 does not specify what the ownership of the process 21967 performing the writing to the output device may be. If -c is not used, it is unspecified whether 21968 the process performing the writing to the output device has permission to read file if there are 21969 any restrictions in place on who may read file until after it is printed. Also, if -c is not used, the 21970 results of deleting *file* before it is printed are unspecified. 21971 21972 FUTURE DIRECTIONS None. 21973 21974 SEE ALSO 21975 mailx 21976 CHANGE HISTORY First released in Issue 2. 21977 21978 Issue 6 The following new requirements on POSIX implementations derive from alignment with the 21979 Single UNIX Specification: 21980 21981 • In the DESCRIPTION, the requirement to associate a unique request ID, and the normal generation of a banner page is added. 21982 • In the OPTIONS section: 21983 — The **-d** *dest* description is expanded, but references to *lpstat* are removed. 21984 — The $-\mathbf{m}$, $-\mathbf{o}$, $-\mathbf{s}$, $-\mathbf{t}$, and $-\mathbf{w}$ options are added. 21985 21986 • In the ENVIRONMENT VARIABLES section, *LC_TIME* may now affect the execution. • The STDOUT section is added. 21987 21988 The normative text is reworded to avoid use of the term "must" for application requirements.

The *TZ* entry is added to the ENVIRONMENT VARIABLES section.

21989

Utilities ls

```
21990 NAME
```

21995

21996

21997

21998 21999

22000

22001

22002

22003 22004

22005 22006

22007

22008

22009

22010

22012 22013

22014

22015

22016 22017

22018

21991 ls — list directory contents

21992 SYNOPSIS

21993 XSI ls [-CFRacdilqrtu1] [-H | -L] [-fgmnopsx] [file...]

21994 DESCRIPTION

For each operand that names a file of a type other than directory or symbolic link to a directory, *ls* shall write the name of the file as well as any requested, associated information. For each operand that names a file of type directory, *ls* shall write the names of files contained within the directory as well as any requested, associated information. If one of the $-\mathbf{d}$, $-\mathbf{F}$, or $-\mathbf{l}$ options are specified, and one of the $-\mathbf{H}$ or $-\mathbf{L}$ options are not specified, for each operand that names a file of type symbolic link to a directory, *ls* shall write the name of the file as well as any requested, associated information. If none of the $-\mathbf{d}$, $-\mathbf{F}$, or $-\mathbf{l}$ options are specified, or the $-\mathbf{H}$ or $-\mathbf{L}$ options are specified, for each operand that names a file of type symbolic link to a directory, *ls* shall write the names of files contained within the directory as well as any requested, associated information.

If no operands are specified, *Is* shall write the contents of the current directory. If more than one operand is specified, *Is* shall write non-directory operands first; it shall sort directory and non-directory operands separately according to the collating sequence in the current locale.

The *ls* utility shall detect infinite loops; that is, entering a previously visited directory that is an ancestor of the last file encountered. When it detects an infinite loop, *ls* shall write a diagnostic message to standard error and shall either recover its position in the hierarchy or terminate.

22011 OPTIONS

The *ls* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported:

- -C Write multi-text-column output with entries sorted down the columns, according to the collating sequence. The number of text columns and the column separator characters are unspecified, but should be adapted to the nature of the output device.
- 22019 -F Do not follow symbolic links named as operands unless the -H or -L options are specified. Write a slash ('/') immediately after each pathname that is a directory, an asterisk ('*') after each that is executable, a vertical bar ('|') after each that is a FIFO, and an at sign ('@') after each that is a symbolic link. For other file types, other symbols may be written.
- 22024 —H If a symbolic link referencing a file of type directory is specified on the command line, *Is* shall evaluate the file information and file type to be those of the file referenced by the link, and not the link itself; however, *Is* shall write the name of the link itself and not the file referenced by the link.
- Evaluate the file information and file type for all symbolic links (whether named on the command line or encountered in a file hierarchy) to be those of the file referenced by the link, and not the link itself; however, *Is* shall write the name of the link itself and not the file referenced by the link. When –L is used with –l, write the contents of symbolic links in the long format (see the STDOUT section).
- 22033 —R Recursively list subdirectories encountered.
- Write out all directory entries, including those whose names begin with a period ('.'). Entries beginning with a period shall not be written out unless explicitly

ls **Utilities**

22036 22037		referenced, the $-\mathbf{a}$ option is supplied, or an implementation-defined condition shall cause them to be written.
22038 22039 22040	-с	Use time of last modification of the file status information (see $<$ sys/stat.h $>$ in the System Interfaces volume of IEEE Std 1003.1-2001) instead of last modification of the file itself for sorting ($-$ t) or writing ($-$ l).
22041 22042 22043	−d	Do not follow symbolic links named as operands unless the $-\mathbf{H}$ or $-\mathbf{L}$ options are specified. Do not treat directories differently than other types of files. The use of $-\mathbf{d}$ with $-\mathbf{R}$ produces unspecified results.
22044 XSI 22045 22046	−f	Force each argument to be interpreted as a directory and list the name found in each slot. This option shall turn off $-\mathbf{l}$, $-\mathbf{t}$, $-\mathbf{s}$, and $-\mathbf{r}$, and shall turn on $-\mathbf{a}$; the order is the order in which entries appear in the directory.
22047 XSI	–g	The same as $-\mathbf{l}$, except that the owner shall not be written.
22048 22049	−i	For each file, write the file's file serial number (see $stat()$ in the System Interfaces volume of IEEE Std 1003.1-2001).
22050 22051 22052	-l	(The letter ell.) Do not follow symbolic links named as operands unless the $-\mathbf{H}$ or $-\mathbf{L}$ options are specified. Write out in long format (see the STDOUT section). When $-\mathbf{l}$ (ell) is specified, -1 (one) shall be assumed.
22053 XSI	-m	Stream output format; list files across the page, separated by commas.
22054 XSI 22055	–n	The same as –l, except that the owner's UID and GID numbers shall be written, rather than the associated character strings.
22056 XSI	-0	The same as $-\mathbf{l}$, except that the group shall not be written.
22057 XSI	- p	Write a slash $('/')$ after each filename if that file is a directory.
22058 22059 22060	- q	Force each instance of non-printable filename characters and $<$ tab $>$ s to be written as the question-mark ($'$? $'$) character. Implementations may provide this option by default if the output is to a terminal device.
22061	-r	Reverse the order of the sort to get reverse collating sequence or oldest first.
22062 XSI 22063	-s	Indicate the total number of file system blocks consumed by each file displayed. The block size is implementation-defined.
22064 22065	-t	Sort with the primary key being time modified (most recently modified first) and the secondary key being filename in the collating sequence.
22066 22067	−u	Use time of last access (see $<$ sys/stat.h $>$) instead of last modification of the file for sorting (-t) or writing (-l).
22068 XSI 22069	- x	The same as -C , except that the multi-text-column output is produced with entries sorted across, rather than down, the columns.
22070	-1	(The numeric digit one.) Force output to be one entry per line.
22071 22072 XSI 22073	considered a	more than one of the options in the following mutually-exclusive pairs shall not be an error: $-\mathbf{C}$ and $-\mathbf{l}$ (ell), $-\mathbf{m}$ and $-\mathbf{l}$ (ell), $-\mathbf{x}$ and $-\mathbf{l}$ (ell), $-\mathbf{C}$ and $-\mathbf{l}$ (one), $-\mathbf{H}$ and $-\mathbf{L}$, The last option specified in each pair shall determine the output format.
22074 OPERA 22075		ng operand shall be supported:
		• • • • • • • • • • • • • • • • • • • •
22076 22077	file	A pathname of a file to be written. If the file specified is not found, a diagnostic message shall be output on standard error.

22077 message shall be output on standard error. Utilities ls

22078 STDIN 22079	Not used.	
22080 INPUT 22081	None.	
	ONMENT VA	ADIARI FS
22082 EIN VIIC 22083		ng environment variables shall affect the execution of <i>ls</i> :
22084 22085 22086 22087 22088 22089 22090 22091		Determine the user's preferred column position width for writing multiple text-column output. If this variable contains a string representing a decimal integer, the <i>ls</i> utility shall calculate how many pathname text columns to write (see –C) based on the width provided. If <i>COLUMNS</i> is not set or invalid, an implementation-defined number of column positions shall be assumed, based on the implementation's knowledge of the output device. The column width chosen to write the names of files in any given directory shall be constant. Filenames shall not be truncated to fit into the multiple text-column output.
22092 22093 22094 22095	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
22096 22097	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
22098	LC_COLLAT	TE .
22099 22100		Determine the locale for character collation information in determining the pathname collation sequence.
22101 22102 22103	LC_CTYPE	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 which characters are defined as printable (character class print).
22104	LC_MESSA	GES
22105 22106		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
22107	LC_TIME	Determine the format and contents for date and time strings written by <i>ls</i> .
22108 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .
22109 22110	TZ	Determine the timezone for date and time strings written by <i>ls.</i> If <i>TZ</i> is unset or null, an unspecified default timezone shall be used.
22111 ASYN (CHRONOUS	EVENTS
22112	Default.	
22113 STDOU 22114 22115 XSI 22116	The default terminals or	format shall be to list one entry per line to standard output; the exceptions are to when one of the $-\mathbf{C}$, $-\mathbf{m}$, or $-\mathbf{x}$ options is specified. If the output is to a terminal, the plementation-defined.
22117 XSI	When - m is	specified, the format used shall be:
22118	"%s, %s,	\n", <filename1>, <filename2></filename2></filename1>
22119	where the la	rgest number of filenames shall be written without exceeding the length of the line.
22120		ion is specified, the file's file serial number (see < sys/stat.h >) shall be written in the

following format before any other output for the corresponding entry:

22121

ls Utilities

```
22122
              %u ", <file serial number>
              If the –l option is specified without –L, the following information shall be written:
22123
              "%s %u %s %s %u %s %s\n", <file mode>, <number of links>,
22124
22125
                   <owner name>, <group name>, <number of bytes in the file>,
22126
                   <date and time>, <pathname>
              If the file is a symbolic link, this information shall be about the link itself and the <pathname>
22127
              field shall be of the form:
22128
              "%s -> %s", <pathname of link>, <contents of link>
22129
22130
              If both –I and –L are specified, the following information shall be written:
22131
              "%s %u %s %s %u %s %s\n", <file mode>, <number of links>,
22132
                   <owner name>, <group name>, <number of bytes in the file>,
                   <date and time>, <pathname of link>
22133
              where all fields except <pathname of link> shall be for the file resolved from the symbolic link.
22134
              The -g, -n, and -o options use the same format as -l, but with omitted items and their
22135 XSI
              associated <blank>s. See the OPTIONS section.
22136
              In both the preceding –l forms, if <owner name> or <group name> cannot be determined, or if –n
22137 XSI
22138
              is given, they shall be replaced with their associated numeric values using the format %u.
              The <date and time> field shall contain the appropriate date and timestamp of when the file was
22139
22140
              last modified. In the POSIX locale, the field shall be the equivalent of the output of the following
22141
              date command:
              date "+%b %e %H:%M"
22142
              if the file has been modified in the last six months, or:
22143
22144
              date "+%b %e %Y"
22145
              (where two <space>s are used between %e and %Y) if the file has not been modified in the last six
22146
              months or if the modification date is in the future, except that, in both cases, the final <newline>
              produced by date shall not be included and the output shall be as if the date command were
22147
              executed at the time of the last modification date of the file rather than the current time. When
22148
22149
              the LC_TIME locale category is not set to the POSIX locale, a different format and order of
22150
              presentation of this field may be used.
22151
              If the file is a character special or block special file, the size of the file may be replaced with
22152
              implementation-defined information associated with the device in question.
              If the pathname was specified as a file operand, it shall be written as specified.
22153
22154 XSI
              The file mode written under the -\mathbf{l}, -\mathbf{g}, -\mathbf{n}, and -\mathbf{o} options shall consist of the following format:
              "%c%s%s%s%c", <entry type>, <owner permissions>,
22155
22156
                   <group permissions>, <other permissions>,
22157
                   <optional alternate access method flag>
              The <optional alternate access method flag> shall be a single <space> if there is no alternate or
22158
              additional access control method associated with the file; otherwise, a printable character shall
22159
              be used.
22160
              The <entry type> character shall describe the type of file, as follows:
22161
22162
              d
                       Directory.
```

Utilities ls

22163	b Block special file.
22164	C Character special file.
22165	1 (ell) Symbolic link.
22166	p FIFO.
22167	- Regular file.
22168	Implementations may add other characters to this list to represent other implementation-defined
22169	file types.
22170	The next three fields shall be three characters each:
22171	<pre><owner permissions=""> Permissions for the file owner class (see the Base Definitions volume of</owner></pre>
22172 22173	Permissions for the file owner class (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 4.4, File Access Permissions).
22174	<pre><group permissions=""></group></pre>
22175	Permissions for the file group class.
22176 22177	<pre><other permissions=""> Permissions for the file other class.</other></pre>
22178	Each field shall have three character positions:
22179	1. If $'r'$, the file is readable; if $'-'$, the file is not readable.
22180	2. If 'w', the file is writable; if ' $-$ ', the file is not writable.
22181	3. The first of the following that applies:
22182 22183	If in <i><owner permissions=""></owner></i> , the file is not executable and set-user-ID mode is set. If in <i><group permissions=""></group></i> , the file is not executable and set-group-ID mode is set.
22184 22185	s If in <i><owner permissions=""></owner></i> , the file is executable and set-user-ID mode is set. If in <i><group permissions=""></group></i> , the file is executable and set-group-ID mode is set.
22186 XSI 22187	T If in <i><other permissions=""></other></i> and the file is a directory, search permission is not granted to others, and the restricted deletion flag is set.
22188 XSI 22189	t If in <i><other permissions=""></other></i> and the file is a directory, search permission is granted to others, and the restricted deletion flag is set.
22190	x The file is executable or the directory is searchable.
22191	- None of the attributes of 'S', 's', 'T', 't', or 'x' applies.
22192	Implementations may add other characters to this list for the third character position. Such
22193 22194	additions shall, however, be written in lowercase if the file is executable or searchable, and in uppercase if it is not.
	If any of the $-\mathbf{l}$, $-\mathbf{g}$, $-\mathbf{n}$, $-\mathbf{o}$, or $-\mathbf{s}$ options is specified, each list of files within the directory shall be
22195 XSI 22196	preceded by a status line indicating the number of file system blocks occupied by files in the
22197	directory in 512-byte units, rounded up to the next integral number of units, if necessary. In the
22198	POSIX locale, the format shall be:
22199	"total %u\n", <number directory="" in="" of="" the="" units=""></number>
22200 22201 22202	If more than one directory, or a combination of non-directory files and directories are written, either as a result of specifying multiple operands, or the $-\mathbf{R}$ option, each list of files within a directory shall be preceded by:

ls Utilities

22203 "\n%s:\n", <directory name>

22204 If this string is the first thing to be written, the first <newline> shall not be written. This output

shall precede the number of units in the directory.

22206 XSI If the $-\mathbf{s}$ option is given, each file shall be written with the number of blocks used by the file.

Along with $-\mathbf{C} = -\mathbf{1}$ -m or $-\mathbf{x}$ the number and a < space > shall precede the filename: with $-\mathbf{g} = -\mathbf{1}$

Along with -C, -1, -m, or -x, the number and a <space> shall precede the filename; with -g, -1,

 $-\mathbf{n}$, or $-\mathbf{o}$, they shall precede each line describing a file.

22209 STDERR

22208

22225

22226

22227

22228 22229

22230 22231

22232

22233

22234

22235

22236 22237

22238 22239

22240

22241

22242

22243 22244

22245

22246

22247

The standard error shall be used only for diagnostic messages.

22211 OUTPUT FILES

22212 None.

22213 EXTENDED DESCRIPTION

22214 None.

22215 EXIT STATUS

The following exit values shall be returned:

22217 0 Successful completion.

>0 An error occurred.

22219 CONSEQUENCES OF ERRORS

22220 Default.

22221 APPLICATION USAGE

Many implementations use the equal sign ('=') to denote sockets bound to the file system for the $-\mathbf{F}$ option. Similarly, many historical implementations use the 's' character to denote sockets as the entry type characters for the $-\mathbf{I}$ option.

It is difficult for an application to use every part of the file modes field of *ls* –**l** in a portable manner. Certain file types and executable bits are not guaranteed to be exactly as shown, as implementations may have extensions. Applications can use this field to pass directly to a user printout or prompt, but actions based on its contents should generally be deferred, instead, to the *test* utility.

The output of *ls* (with the –l and related options) contains information that logically could be used by utilities such as *chmod* and *touch* to restore files to a known state. However, this information is presented in a format that cannot be used directly by those utilities or be easily translated into a format that can be used. A character has been added to the end of the permissions string so that applications at least have an indication that they may be working in an area they do not understand instead of assuming that they can translate the permissions string into something that can be used. Future issues or related documents may define one or more specific characters to be used based on different standard additional or alternative access control mechanisms.

As with many of the utilities that deal with filenames, the output of *ls* for multiple files or in one of the long listing formats must be used carefully on systems where filenames can contain embedded white space. Systems and system administrators should institute policies and user training to limit the use of such filenames.

The number of disk blocks occupied by the file that it reports varies depending on underlying file system type, block size units reported, and the method of calculating the number of blocks. On some file system types, the number is the actual number of blocks occupied by the file (counting indirect blocks and ignoring holes in the file); on others it is calculated based on the file size (usually making an allowance for indirect blocks, but ignoring holes).

Utilities ls

22248 EXAMPLES

22249 An example of a small directory tree being fully listed with <i>Is</i> – laRF a in the POSIX

2225	60	total 11								
2225	1	drwxr-xr-x	3	hlj	prog	64	Jul	4	12:07	./
2225	2	drwxrwxrwx	4	hlj	prog	3264	Jul	4	12:09	/
2225	3	drwxr-xr-x	2	hlj	prog	48	Jul	4	12:07	b/
2225	4	-rwxrr	1	hlj	prog	572	Jul	4	12:07	foo*
2225	5	a/b:								
2225	6	total 4								
2225	7	drwxr-xr-x	2	hlj	prog	48	Jul	4	12:07	./
2225	8	drwxr-xr-x	3	hlj	prog	64	Jul	4	12:07	/
2225	9	-rw-rr	1	hlj	prog	700	Jul	4	12:07	bar

22260 RATIONALE

Some historical implementations of the ls utility show all entries in a directory except dot and dot-dot when a superuser invokes ls without specifying the -a option. When "normal" users invoke ls without specifying -a, they should not see information about any files with names beginning with a period unless they were named as file operands.

Implementations are expected to traverse arbitrary depths when processing the $-\mathbf{R}$ option. The only limitation on depth should be based on running out of physical storage for keeping track of untraversed directories.

The -1 (one) option was historically found in BSD and BSD-derived implementations only. It is required in this volume of IEEE Std 1003.1-2001 so that conforming applications might ensure that output is one entry per line, even if the output is to a terminal.

Generally, this volume of IEEE Std 1003.1-2001 is silent about what happens when options are given multiple times. In the cases of -C, -I, and -I, however, it does specify the results of these overlapping options. Since Is is one of the most aliased commands, it is important that the implementation perform intuitively. For example, if the alias were:

```
alias ls="ls -C"
```

22276 and the user typed ls-1, single-text-column output should result, not an error.

The BSD ls provides a $-\mathbf{A}$ option (like $-\mathbf{a}$, but dot and dot-dot are not written out). The small difference from $-\mathbf{a}$ did not seem important enough to require both.

Implementations may make $-\mathbf{q}$ the default for terminals to prevent trojan horse attacks on terminals with special escape sequences. This is not required because:

- Some control characters may be useful on some terminals; for example, a system might write them as "\001" or "^A".
- Special behavior for terminals is not relevant to applications portability.

An early proposal specified that the optional alternate access method flag had to be '+' if there was an alternate access method used on the file or <space> if there was not. This was changed to be <space> if there is not and a single printable character if there is. This was done for three reasons:

- 1. There are historical implementations using characters other than '+'.
- 2. There are implementations that vary this character used in that position to distinguish between various alternate access methods in use.

ls Utilities

3. The standard developers did not want to preclude future specifications that might need a way to specify more than one alternate access method.

Nonetheless, implementations providing a single alternate access method are encouraged to use '+'.

In an early proposal, the units used to specify the number of blocks occupied by files in a directory in an ls –I listing were implementation-defined. This was because BSD systems have historically used 1024-byte units and System V systems have historically used 512-byte units. It was pointed out by BSD developers that their system has used 512-byte units in some places and 1024-byte units in other places. (System V has consistently used 512.) Therefore, this volume of IEEE Std 1003.1-2001 usually specifies 512. Future releases of BSD are expected to consistently provide 512 bytes as a default with a way of specifying 1024-byte units where appropriate.

The *<date and time>* field in the *-*I format is specified only for the POSIX locale. As noted, the format can be different in other locales. No mechanism for defining this is present in this volume of IEEE Std 1003.1-2001, as the appropriate vehicle is a messaging system; that is, the format should be specified as a "message".

22306 FUTURE DIRECTIONS

The $-\mathbf{s}$ uses implementation-defined units and cannot be used portably; it may be withdrawn in a future version.

22309 SEE ALSO

22295

22296

22297

22298

22299 22300

22301

22302

22303

22304 22305

chmod, *find*, the System Interfaces volume of IEEE Std 1003.1-2001, *stat*(), the Base Definitions volume of IEEE Std 1003.1-2001, *<sys/stat.h>*

22312 CHANGE HISTORY

First released in Issue 2.

22314 Issue 5

22315 A second FUTURE DIRECTION is added.

22316 Issue 6

22319

22320

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

In the –F option, other symbols are allowed for other file types.

Treatment of symbolic links is added, as defined in the IEEE P1003.2b draft standard.

The Open Group Base Resolution bwg2001-010 is applied, adding the ${\tt T}$ and ${\tt t}$ fields as an XSI extension.

Utilities m4

22323 **NAME** 22324 m4 — macro processor (**DEVELOPMENT**) 22325 SYNOPSIS [-s][-D name[=val]]...[-U name]... file...22326 XSI 22327 22328 DESCRIPTION The *m4* utility is a macro processor that shall read one or more text files, process them according 22329 to their included macro statements, and write the results to standard output. 22330 22331 OPTIONS 22332 The *m4* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines, except that the order of the -**D** and -**U** options shall be significant. 22333 The following options shall be supported: 22334 Enable line synchronization output for the c99 preprocessor phase (that is, #line 22335 -sdirectives). 22336 $-\mathbf{D}$ name[=val]22337 Define *name* to *val* or to null if = *val* is omitted. 22338 -U name Undefine name. 22339 22340 OPERANDS The following operand shall be supported: 22341 22342 file A pathname of a text file to be processed. If no file is given, or if it is '-', the standard input shall be read. 22343 22344 **STDIN** 22345 The standard input shall be a text file that is used if no *file* operand is given, or if it is '-'. 22346 INPUT FILES The input file named by the *file* operand shall be a text file. 22347 22348 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *m4*: 22349 LANG Provide a default value for the internationalization variables that are unset or null. 22350 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 22351 Internationalization Variables for the precedence of internationalization variables 22352 used to determine the values of locale categories.) 22353 LC_ALL If set to a non-empty string value, override the values of all the other 22354 internationalization variables. 22355 Determine the locale for the interpretation of sequences of bytes of text data as 22356 LC_CTYPE 22357 characters (for example, single-byte as opposed to multi-byte characters in 22358 arguments and input files). LC_MESSAGES 22359 Determine the locale that should be used to affect the format and contents of 22360 22361 diagnostic messages written to standard error. **NLSPATH**

Determine the location of message catalogs for the processing of *LC_MESSAGES*.

22362

m4 Utilities

22363 ASYNCHRONOUS EVENTS

22364 Default.

22365 STDOUT

The standard output shall be the same as the input files, after being processed for macro expansion.

22368 STDERR

The standard error shall be used to display strings with the **errprint** macro, macro tracing enabled by the **traceon** macro, the defined text for macros written by the **dumpdef** macro, or for diagnostic messages.

22372 OUTPUT FILES

22373 None.

22375

22376

22377

22378

2237922380

22383

22384

22385

22386

22387

22388

22389 22390

22391

2239222393

22394

22395

22396

22397

22398 22399

22400

22401 22402 22403

22404

22405

22406

22407 22408

22374 EXTENDED DESCRIPTION

The *m4* utility shall compare each token from the input against the set of built-in and user-defined macros. If the token matches the name of a macro, then the token shall be replaced by the macro's defining text, if any, and rescanned for matching macro names. Once no portion of the token matches the name of a macro, it shall be written to standard output. Macros may have arguments, in which case the arguments shall be substituted into the defining text before it is rescanned.

22381 Macro calls have the form:

22382 name(arg1, arg2, ..., argn)

Macro names shall consist of letters, digits, and underscores, where the first character is not a digit. Tokens not of this form shall not be treated as macros.

The application shall ensure that the left parenthesis immediately follows the name of the macro. If a token matching the name of a macro is not followed by a left parenthesis, it is handled as a use of that macro without arguments.

If a macro name is followed by a left parenthesis, its arguments are the comma-separated tokens between the left parenthesis and the matching right parenthesis. Unquoted

blank>s and <newline>s preceding each argument shall be ignored. All other characters, including trailing

 characters and <newline>s, are retained. Commas enclosed between left and right parenthesis characters do not delimit arguments.

Arguments are positionally defined and referenced. The string "\$1" in the defining text shall be replaced by the first argument. Systems shall support at least nine arguments; only the first nine can be referenced, using the strings "\$1" to "\$9", inclusive. The string "\$0" is replaced with the name of the macro. The string "\$#" is replaced by the number of arguments as a string. The string "\$*" is replaced by a list of all of the arguments, separated by commas. The string "\$@" is replaced by a list of all of the arguments separated by commas, and each argument is quoted using the current left and right quoting strings.

If fewer arguments are supplied than are in the macro definition, the omitted arguments are taken to be null. It is not an error if more arguments are supplied than are in the macro definition.

No special meaning is given to any characters enclosed between matching left and right quoting strings, but the quoting strings are themselves discarded. By default, the left quoting string consists of a grave accent (' ' ') and the right quoting string consists of an acute accent (' ' '); see also the **changequote** macro.

Comments are written but not scanned for matching macro names; by default, the begincomment string consists of the number sign character and the end-comment string consists of a Utilities m4

22409	<newline>.</newline>	See also the changecom and dnl macros.			
22410	The <i>m4</i> utili	ity shall make available the following built-in macros. They can be redefined, but			
22411		done the original meaning is lost. Their values shall be null unless otherwise stated.			
22412	In the descriptions below, the term <i>defining text</i> refers to the value of the macro: the second				
22413		the define macro, among other things. Except for the first argument to the eval			
22414		umeric arguments to built-in macros shall be interpreted as decimal values. The			
22415	string value	s produced as the defining text of the decr, divnum, incr, index, len, and sysval			
22416	built-in mac	ros shall be in the form of a decimal-constant as defined in the C language.			
22417	changecom	0			
22418		With no arguments, the comment mechanism shall be disabled. With a single			
22419		argument, that argument shall become the begin-comment string and the			
22420		<newline> shall become the end-comment string. With two arguments, the first</newline>			
22421		argument shall become the begin-comment string and the second argument shall			
22422		become the end-comment string. Systems shall support comment strings of at least			
22423		five characters.			
22424	changequot	e The changequote macro shall set the begin-quote and end-quote strings. With no			
22425		arguments, the quote strings shall be set to the default values (that is, ''). With a			
22426		single argument, that argument shall become the begin-quote string and the			
22427		<newline> shall become the end-quote string. With two arguments, the first</newline>			
22428		argument shall become the begin-quote string and the second argument shall			
22429		become the end-quote string. Systems shall support quote strings of at least five			
22430		characters.			
22431	decr	The defining text of the decr macro shall be its first argument decremented by 1. It			
22432		shall be an error to specify an argument containing any non-numeric characters.			
22433	define	The second argument shall become the defining text of the macro whose name is			
22434	define	the first argument.			
	1.0				
22435	defn	The defining text of the defn macro shall be the quoted definition (using the			
22436		current quoting strings) of its arguments.			
22437	divert	The <i>m4</i> utility maintains nine temporary buffers, numbered 1 to 9, inclusive. When			
22438		the last of the input has been processed, any output that has been placed in these			
22439		buffers shall be written to standard output in buffer-numerical order. The divert			
22440		macro shall divert future output to the buffer specified by its argument. Specifying			
22441		no argument or an argument of 0 shall resume the normal output process. Output			
22442		diverted to a stream other than 0 to 9 shall be discarded. It shall be an error to			
22443		specify an argument containing any non-numeric characters.			
22444	divnum	The defining text of the divnum macro shall be the number of the current output			
22445		stream as a string.			
22446	dnl	The dnl macro shall cause <i>m4</i> to discard all input characters up to and including			
22446	um	the next <newline>.</newline>			
22447	_				
22448	dumpdef	The dumpdef macro shall write the defined text to standard error for each of the			
22449		macros specified as arguments, or, if no arguments are specified, for all macros.			
22450	errprint	The errprint macro shall write its arguments to standard error.			
22451	eval	The eval macro shall evaluate its first argument as an arithmetic expression, using			
22452		32-bit signed integer arithmetic. All of the C-language operators shall be			
22453		supported, except for:			
		**			

m4 Utilities

22454		
22455		->
22456		++
22457		
22458		(type)
22459		unary *
22460		sizeof
22461		
22462		
22463		?:
22464		unary &
22465		and all assignment operators. It shall be an error to specify any of these operators.
22466		Precedence and associativity shall be as in the ISO C standard. Systems shall
22467		support octal and hexadecimal numbers as in the ISO C standard. The second
22468		argument, if specified, shall set the radix for the result; the default is 10. The third
22469		argument, if specified, sets the minimum number of digits in the result. It shall be
22470		an error to specify the second or third argument containing any non-numeric
22471		characters.
22472	ifdef	If the first argument to the ifdef macro is defined, the defining text shall be the
22473	11401	second argument. Otherwise, the defining text shall be the third argument, if
22474		specified, or the null string, if not.
22475	ifelse	The ifelse macro takes three or more arguments. If the first two arguments
22476	neise	compare as equal strings (after macro expansion of both arguments), the defining
22477		text shall be the third argument. If the first two arguments do not compare as
		equal strings and there are three arguments, the defining text shall be null. If the
22478		first two arguments do not compare as equal strings and there are four or five
22479		arguments, the defining text shall be the fourth argument. If the first two
22480		
22481		arguments do not compare as equal strings and there are six or more arguments,
22482		the first three arguments shall be discarded and processing shall restart with the
22483		remaining arguments.
22484 22485	include	The defining text for the include macro shall be the contents of the file named by the first argument. It shall be an error if the file cannot be read.
22486 22487	incr	The defining text of the incr macro shall be its first argument incremented by 1. It shall be an error to specify an argument containing any non-numeric characters.
22488	index	The defining text of the index macro shall be the first character position (as a
22489	Huck	string) in the first argument where a string matching the second argument begins
22490		(zero origin), or –1 if the second argument does not occur.
22430		(2010 origin), or a rather second argument does not occur.
22491 22492	len	The defining text of the len macro shall be the length (as a string) of the first argument.
22493	m4exit	Exit from the <i>m4</i> utility. If the first argument is specified, it is the exit code. The
22494		default is zero. It shall be an error to specify an argument containing any non-
22495		numeric characters.
22496	m4wrap	The first argument shall be processed when EOF is reached. If the m4wrap macro
22497		is used multiple times, the arguments specified shall be processed in the order in
22498		which the m4wrap macros were processed.
22499	maketemp	The defining text shall be the first argument, with any trailing 'X' characters
22500	1	replaced with the current process ID as a string.
		,

Utilities m4

22501 22502 22503	popdef	The popdef macro shall delete the current definition of its arguments, replacing that definition with the previous one. If there is no previous definition, the macro is undefined.		
22504 22505	pushdef	The pushdef macro shall be equivalent to the define macro with the exception that it shall preserve any current definition for future retrieval using the popdef macro.		
22506 22507	shift	The defining text for the shift macro shall be all of its arguments except for the first one.		
22508 22509	sinclude	The sinclude macro shall be equivalent to the include macro, except that it shall not be an error if the file is inaccessible.		
22510 22511 22512 22513 22514 22515 22516	substr	The defining text for the substr macro shall be the substring of the first argument beginning at the zero-offset character position specified by the second argument. The third argument, if specified, shall be the number of characters to select; if not specified, the characters from the starting point to the end of the first argument shall become the defining text. It shall not be an error to specify a starting point beyond the end of the first argument and the defining text shall be null. It shall be an error to specify an argument containing any non-numeric characters.		
22517 22518 22519 22520	syscmd	The syscmd macro shall interpret its first argument as a shell command line. The defining text shall be the string result of that command. No output redirection shall be performed by the <i>m4</i> utility. The exit status value from the command can be retrieved using the sysval macro.		
22521 22522	sysval	The defining text of the sysval macro shall be the exit value of the utility last invoked by the syscmd macro (as a string).		
22523 22524 22525	traceon	The traceon macro shall enable tracing for the macros specified as arguments, or, if no arguments are specified, for all macros. The trace output shall be written to standard error in an unspecified format.		
22526 22527	traceoff	The traceoff macro shall disable tracing for the macros specified as arguments, or, if no arguments are specified, for all macros.		
22528 22529 22530	translit	The defining text of the translit macro shall be the first argument with every character that occurs in the second argument replaced with the corresponding character from the third argument.		
22531 22532	undefine	The undefine macro shall delete all definitions (including those preserved using the pushdef macro) of the macros named by its arguments.		
22533 22534 22535 22536 22537	undivert	The undivert macro shall cause immediate output of any text in temporary buffers named as arguments, or all temporary buffers if no arguments are specified. Buffers can be undiverted into other temporary buffers. Undiverting shall discard the contents of the temporary buffer. It shall be an error to specify an argument containing any non-numeric characters.		
22538 EXIT		and another and a like the materials.		
22539		ng exit values shall be returned:		
22540 22541		0 Successful completion.>0 An error occurred		
22542		If the m4exit macro is used, the exit value can be specified by the input file.		
www.ib	II die IIITCAI	i macro to abou, the one raide can be specified by the hiput me.		

m4 Utilities

22543 CONSEQUENCES OF ERRORS

22544 Default.

```
22545 APPLICATION USAGE
```

```
22546
            The defn macro is useful for renaming macros, especially built-ins.
22547 EXAMPLES
22548
            If the file m4src contains the lines:
22549
            The value of 'VER' is "VER".
            ifdef('VER', ''VER'' is defined to be VER., VER is not defined.)
22550
            ifelse(VER, 1, ''VER'' is 'VER'.)
22551
            ifelse(VER, 2, ''VER'' is 'VER'., ''VER'' is not 2.)
22552
22553
            then the command
22554
            m4 m4src
22555
22556
            or the command:
            m4 -U VER m4src
22557
            produces the output:
22558
            The value of VER is "VER".
22559
22560
            VER is not defined.
22561
            VER is not 2.
22562
            end
            The command:
22563
22564
            m4 -D VER m4src
22565
            produces the output:
22566
            The value of VER is "".
            VER is defined to be .
22567
22568
            VER is not 2.
            end
22569
            The command:
22570
            m4 -D VER=1 m4src
22571
22572
            produces the output:
            The value of VER is "1".
22573
            VER is defined to be 1.
22574
            VER is 1.
22575
22576
            VER is not 2.
            end
22577
            The command:
22578
22579
            m4 -D VER=2 m4src
22580
            produces the output:
            The value of VER is "2".
22581
            VER is defined to be 2.
22582
```

Utilities m4

22583 VER is 2. 22584 end 22585 RATIONALE None. 22586 22587 FUTURE DIRECTIONS 22588 None. **22589 SEE ALSO** 22590 c99 22591 CHANGE HISTORY 22592 First released in Issue 2. 22593 Issue 5 The phrase "the defined text for macros written by the dumpdef macro" is added to the 22594 description of STDERR, and the description of dumpdef is updated to indicate that output is 22595 written to standard error. The description of eval is updated to indicate that the list of excluded 22596 C operators excludes unary '&' and '.'. In the description of **ifdef**, the phrase "and it is not 22597 defined to be zero" is deleted. 22598 22599 Issue 6 22600 In the EXTENDED DESCRIPTION, the eval text is updated to include a '&' character in the excepted list. 22601 The EXTENDED DESCRIPTION of divert is updated to clarify that there are only nine diversion 22602 22603 buffers. The normative text is reworded to avoid use of the term "must" for application requirements. 22604 The Open Group Base Resolution bwg2000-006 is applied. 22605 IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/31 is applied, replacing the EXAMPLES 22606

section.

22607

22608 NAME

22609 mailx — process messages

22610 SYNOPSIS

22611 Send Mode

22612 mailx [-s subject] address...

22613 Receive Mode

22614 mailx -e

22615 mailx [-HiNn] [-F] [-u user]
22616 mailx -f [-HiNn] [-F] [file]

22617 **DESCRIPTION**

The *mailx* utility provides a message sending and receiving facility. It has two major modes, selected by the options used: Send Mode and Receive Mode.

On systems that do not support the User Portability Utilities option, an application using *mailx* shall have the ability to send messages in an unspecified manner (Send Mode). Unless the first character of one or more lines is tilde ($^{\prime}$ $^{\sim}$ $^{\prime}$), all characters in the input message shall appear in the delivered message, but additional characters may be inserted in the message before it is retrieved.

On systems supporting the User Portability Utilities option, mail-receiving capabilities and other interactive features, Receive Mode, described below, also shall be enabled.

22627 Send Mode

Send Mode can be used by applications or users to send messages from the text in standard input.

22630 Receive Mode

Receive Mode is more oriented towards interactive users. Mail can be read and sent in this interactive mode.

When reading mail, *mailx* provides commands to facilitate saving, deleting, and responding to messages. When sending mail, *mailx* allows editing, reviewing, and other modification of the message as it is entered.

Incoming mail shall be stored in one or more unspecified locations for each user, collectively called the system *mailbox* for that user. When *mailx* is invoked in Receive Mode, the system mailbox shall be the default place to find new mail. As messages are read, they shall be marked to be moved to a secondary file for storage, unless specific action is taken. This secondary file is called the **mbox** and is normally located in the directory referred to by the *HOME* environment variable (see *MBOX* in the ENVIRONMENT VARIABLES section for a description of this file). Messages shall remain in this file until explicitly removed. When the **–f** option is used to read mail messages from secondary files, messages shall be retained in those files unless specifically removed. All three of these locations—system mailbox, **mbox**, and secondary file—are referred to in this section as simply "mailboxes", unless more specific identification is required.

22633

22634 22635

22636 22637

22638

22639

22640

22641

22642

22643

22644

22645

22646 OPTIO 22647		tility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section				
22648		12.2, Utility Syntax Guidelines.				
22649 22650 22651		ng options shall be supported. (Only the -s <i>subject</i> option shall be required on all e other options are required only on systems supporting the User Portability Utilities				
22652 22653	−e	Test for the presence of mail in the system mailbox. The <i>mailx</i> utility shall write nothing and exit with a successful return code if there is mail to read.				
22654 22655 22656	−f	Read messages from the file named by the <i>file</i> operand instead of the system mailbox. (See also folder .) If no <i>file</i> operand is specified, read messages from mbox instead of the system mailbox.				
22657 22658 22659 22660	-F	Record the message in a file named after the first recipient. The name is the login-name portion of the address found first on the To: line in the mail header. Overrides the record variable, if set (see Internal Variables in mailx (on page 593).)				
22661	- H	Write a header summary only.				
22662	- i	Ignore interrupts. (See also ignore .)				
22663 22664	-n	Do not initialize from the system default start-up file. See the EXTENDED DESCRIPTION section.				
22665	-N	Do not write an initial header summary.				
22666 22667 22668	− s subject	Set the Subject header field to <i>subject</i> . All characters in the <i>subject</i> string shall appear in the delivered message. The results are unspecified if <i>subject</i> is longer than {LINE_MAX} – 10 bytes or contains a <newline>.</newline>				
22669 22670 22671	–u user	Read the system mailbox of the login name <i>user</i> . This shall only be successful if the invoking user has the appropriate privileges to read the system mailbox of that user.				
22672 OPERANDS						
22673		ng operands shall be supported:				
22674 22675 22676 22677 22678	address	Addressee of message. When -n is specified and no user start-up files are accessed (see the EXTENDED DESCRIPTION section), the user or application shall ensure this is an address to pass to the mail delivery system. Any system or user start-up files may enable aliases (see alias under Commands in mailx (on page 596)) that may modify the form of <i>address</i> before it is passed to the mail delivery system.				
22679 22680 22681	file	A pathname of a file to be read instead of the system mailbox when –f is specified. The meaning of the <i>file</i> option-argument shall be affected by the contents of the folder internal variable; see Internal Variables in mailx (on page 593).				
22682 STDIN 22683 22684 22685 22686	message to be accepted	is invoked in Send Mode (the first synopsis line), standard input shall be the be delivered to the specified addresses. When in Receive Mode, user commands shall from <i>stdin</i> . If the User Portability Utilities option is not supported, standard input ing with a tilde ('~') character produce unspecified results.				
22687 22688 22689	If the User Portability Utilities option is supported, then in both Send and Receive Modes, standard input lines beginning with the escape character (usually tilde ('~')) shall affect processing as described in Command Escapes in mailx (on page 604).					

22690 INPUT FILES

22691 When mailx is used as described by this volume of IEEE Std 1003.1-2001, the file optionargument (see the -f option) and the mbox shall be text files containing mail messages, 22692 22693 formatted as described in the OUTPUT FILES section. The nature of the system mailbox is 22694 unspecified; it need not be a file.

22695 ENVIRONMENT VARIABLES

22696	The following	ng environment variables shall affect the execution of <i>mailx</i> :
22697 22698 22699 22700 22701	DEAD	Determine the pathname of the file in which to save partial messages in case of interrupts or delivery errors. The default shall be dead.letter in the directory named by the <i>HOME</i> variable. The behavior of <i>mailx</i> in saving partial messages is unspecified if the User Portability Utilities option is not supported and <i>DEAD</i> is not defined with the value / dev/null .
22702 22703 22704 XSI 22705 22706	EDITOR	Determine the name of a utility to invoke when the edit (see Commands in mailx (on page 596)) or ~e (see Command Escapes in mailx (on page 604)) command is used. The default editor is unspecified. On XSI-conformant systems it is <i>ed</i> . The effects of this variable are unspecified if the User Portability Utilities option is not supported.
22707	HOME	Determine the pathname of the user's home directory.
22708 22709 22710 22711	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
22712 22713	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
22714 22715 22716 22717	LC_CTYPE	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) and the handling of case-insensitive address and header-field comparisons.
22718	LC_TIME	Determine the format and contents of the date and time strings written by <i>mailx</i> .
22719 22720 22721 22722	LC_MESSAC	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.
22723 22724	LISTER	Determine a string representing the command for writing the contents of the folder directory to standard output when the folders command is given (see folders in Commands in maily (on page 596)). Any string acceptable as a

ıе ee folders in Commands in mails (on page 596)). Any string acceptable as a 22725 *command string* operand to the *sh* –c command shall be valid. If this variable is null 22726 or not set, the output command shall be ls. The effects of this variable are 22727 22728

unspecified if the User Portability Utilities option is not supported.

Determine the pathname of the start-up file. The default shall be .mailrc in the directory referred to by the HOME environment variable. The behavior of mailx is unspecified if the User Portability Utilities option is not supported and MAILRC is not defined with the value /dev/null.

Determine a pathname of the file to save messages from the system mailbox that have been read. The exit command shall override this function, as shall saving the message explicitly in another file. The default shall be mbox in the directory

22729

22730

22731

22732

22733 22734

22735

MAILRC

MBOX

22736 22737		named by the <i>HOME</i> variable. The effects of this variable are unspecified if the User Portability Utilities option is not supported.
22738 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
22739 22740 22741 22742 22743 22744 22745 22746 22747	PAGER	Determine a string representing an output filtering or pagination command for writing the output to the terminal. Any string acceptable as a <code>command_string</code> operand to the <code>sh-c</code> command shall be valid. When standard output is a terminal device, the message output shall be piped through the command if the <code>mailx</code> internal variable <code>crt</code> is set to a value less the number of lines in the message; see <code>Internal Variables in mailx</code> (on page 593). If the <code>PAGER</code> variable is null or not set, the paginator shall be either <code>more</code> or another paginator utility documented in the system documentation. The effects of this variable are unspecified if the User Portability Utilities option is not supported.
22748 22749 22750	SHELL	Determine the name of a preferred command interpreter. The default shall be <i>sh</i> . The effects of this variable are unspecified if the User Portability Utilities option is not supported.
22751 22752 22753 22754 22755	TERM	If the internal variable screen is not specified, determine the name of the terminal type to indicate in an unspecified manner the number of lines in a screenful of headers. If <i>TERM</i> is not set or is set to null, an unspecified default terminal type shall be used and the value of a screenful is unspecified. The effects of this variable are unspecified if the User Portability Utilities option is not supported.
22756 22757 22758	TZ	This variable may determine the timezone used to calculate date and time strings written by <i>mailx</i> . If <i>TZ</i> is unset or null, an unspecified default timezone shall be used.
22759 22760 22761 22762 22763	VISUAL	Determine a pathname of a utility to invoke when the visual command (see Commands in mailx (on page 596)) or " v command-escape (see Command Escapes in mailx (on page 604)) is used. If this variable is null or not set, the full-screen editor shall be <i>vi</i> . The effects of this variable are unspecified if the User Portability Utilities option is not supported.
22764 ASYNC	HRONOUS I	EVENTS

22764 ASYNCHRONOUS EVENTS

22765

22766

22767

22768

22769

2277022771

22772

22773

22774 22775

22776 22777

22778

22779

22780

When *mailx* is in Send Mode and standard input is not a terminal, it shall take the standard action for all signals.

In Receive Mode, or in Send Mode when standard input is a terminal, if a SIGINT signal is received:

- 1. If in command mode, the current command, if there is one, shall be aborted, and a command-mode prompt shall be written.
- 2. If in input mode:
 - a. If **ignore** is set, *mailx* shall write "@\n", discard the current input line, and continue processing, bypassing the message-abort mechanism described in item 2b.
 - b. If the interrupt was received while sending mail, either when in Receive Mode or in Send Mode, a message shall be written, and another subsequent interrupt, with no other intervening characters typed, shall be required to abort the mail message. If in Receive Mode and another interrupt is received, a command-mode prompt shall be written. If in Send Mode and another interrupt is received, *mailx* shall terminate with a non-zero status.

In both cases listed in item b, if the message is not empty:

i. If **save** is enabled and the file named by *DEAD* can be created, the message shall be written to the file named by *DEAD*. If the file exists, the message shall be written to replace the contents of the file.

ii. If **save** is not enabled, or the file named by *DEAD* cannot be created, the message shall not be saved.

22786 The *mailx* utility shall take the standard action for all other signals.

22787 STDOUT

In command and input modes, all output, including prompts and messages, shall be written to standard output.

22790 STDERR

22791 The standard error shall be used only for diagnostic messages.

22792 OUTPUT FILES

Various *mailx* commands and command escapes can create or add to files, including the **mbox**, the dead-letter file, and secondary mailboxes. When *mailx* is used as described in this volume of IEEE Std 1003.1-2001, these files shall be text files, formatted as follows:

```
22796 line beginning with From<space>
22797 [one or more header-lines; see Commands in mailx (on page 596)]
22798 empty line
22799 [zero or more body lines
22800 empty line]
22801 [line beginning with From<space>...]
```

where each message begins with the **From <space>** line shown, preceded by the beginning of the file or an empty line. (The **From <space>** line is considered to be part of the message header, but not one of the header-lines referred to in **Commands in mailx** (on page 596); thus, it shall not be affected by the **discard**, **ignore**, or **retain** commands.) The formats of the remainder of the **From <space>** line and any additional header lines are unspecified, except that none shall be empty. The format of a message body line is also unspecified, except that no line following an empty line shall start with **From <space>**; *mailx* shall modify any such user-entered message body lines (following an empty line and beginning with **From <space>**) by adding one or more characters to precede the 'F'; it may add these characters to **From <space>** lines that are not preceded by an empty line.

When a message from the system mailbox or entered by the user is not a text file, it is implementation-defined how such a message is stored in files written by *mailx*.

22814 EXTENDED DESCRIPTION

The entire EXTENDED DESCRIPTION section shall apply only to implementations supporting the User Portability Utilities option.

The *mailx* utility cannot guarantee support for all character encodings in all circumstances. For example, inter-system mail may be restricted to 7-bit data by the underlying network, 8-bit data need not be portable to non-internationalized systems, and so on. Under these circumstances, it is recommended that only characters defined in the ISO/IEC 646: 1991 standard International Reference Version (equivalent to ASCII) 7-bit range of characters be used.

When *mailx* is invoked using one of the Receive Mode synopsis forms, it shall write a page of header-summary lines (if –N was not specified and there are messages, see below), followed by a prompt indicating that *mailx* can accept regular commands (see **Commands in mailx** (on page 596)); this is termed *command mode*. The page of header-summary lines shall contain the first new message if there are new messages, or the first unread message if there are unread messages, or the first message. When *mailx* is invoked using the Send Mode synopsis and

standard input is a terminal, if no subject is specified on the command line and the **asksub** variable is set, a prompt for the subject shall be written. At this point, *mailx* shall be in input mode. This input mode shall also be entered when using one of the Receive Mode synopsis forms and a reply or new message is composed using the **reply**, **Reply**, **followup**, **Followup**, or **mail** commands and standard input is a terminal. When the message is typed and the end of the message is encountered, the message shall be passed to the mail delivery software. Commands can be entered by beginning a line with the escape character (by default, tilde (' ~')) followed by a single command letter and optional arguments. See **Commands in mailx** (on page 596) for a summary of these commands. It is unspecified what effect these commands will have if standard input is not a terminal when a message is entered using either the Send Mode synopsis, or the Read Mode commands **reply**, **Reply**, **followup**, **Followup**, or **mail**.

Note: For notational convenience, this section uses the default escape character, tilde, in all references and examples.

At any time, the behavior of *mailx* shall be governed by a set of environmental and internal variables. These are flags and valued parameters that can be set and cleared via the *mailx* set and unset commands.

Regular commands are of the form:

unread

read

[command] [msqlist] [arqument ...]

If no *command* is specified in command mode, **next** shall be assumed. In input mode, commands shall be recognized by the escape character, and lines not treated as commands shall be taken as input for the message.

In command mode, each message shall be assigned a sequential number, starting with 1.

All messages have a state that shall affect how they are displayed in the header summary and how they are retained or deleted upon termination of *mailx*. There is at any time the notion of a *current* message, which shall be marked by a '>' at the beginning of a line in the header summary. When *mailx* is invoked using one of the Receive Mode synopsis forms, the current message shall be the first new message, if there is a new message, or the first unread message if there is an unread message, or the first message if there are any messages, or unspecified if there are no messages in the mailbox. Each command that takes an optional list of messages (*msglist*) or an optional single message (*message*) on which to operate shall leave the current message set to the highest-numbered message of the messages specified, unless the command deletes messages, in which case the current message shall be set to the first undeleted message (that is, a message not in the deleted state) after the highest-numbered message deleted by the command, if one exists, or to an unspecified value if there are no remaining undeleted messages. All messages shall be in one of the following states:

new The message is present in the system mailbox and has not been viewed by the user or moved to any other state. Messages in state *new* when *mailx* quits shall be retained in the system mailbox.

The message has been present in the system mailbox for more than one invocation of *mailx* and has not been viewed by the user or moved to any other state. Messages in state *unread* when *mailx* quits shall be retained in the system mailbox.

The message has been processed by one of the following commands: "f, "m, "F, "M, copy, mbox, next, pipe, print, Print, top, type, Type, undelete. The delete, dp, and dt commands may also cause the next message to be marked as *read*, depending on the value of the **autoprint** variable. Messages that are in the system mailbox and in state *read* when *mailx* quits shall be saved in the **mbox**, unless the internal variable **hold** was set. Messages that are in the **mbox** or in a secondary mailbox and in state

22876		read when mailx quits shall be retained in their current location.
22877 22878 22879 22880 22881 22882 22883 22884	deleted	The message has been processed by one of the following commands: delete , dp , dt . Messages in state <i>deleted</i> when <i>mailx</i> quits shall be deleted. Deleted messages shall be ignored until <i>mailx</i> quits or changes mailboxes or they are specified to the undelete command; for example, the message specification /string shall only search the subject lines of messages that have not yet been deleted, unless the command operating on the list of messages is undelete . No deleted message or deleted message header shall be displayed by any <i>mailx</i> command other than undelete .
22885 22886	preserved	The message has been processed by a preserve command. When <i>mailx</i> quits, the message shall be retained in its current location.
22887 22888 22889 22890 22891 22892 22893	saved	The message has been processed by one of the following commands: save or write . If the current mailbox is the system mailbox, and the internal variable keepsave is set, messages in the state saved shall be saved to the file designated by the <i>MBOX</i> variable (see the ENVIRONMENT VARIABLES section). If the current mailbox is the system mailbox, messages in the state <i>saved</i> shall be deleted from the current mailbox, when the quit or file command is used to exit the current mailbox.
22894	The head	der-summary line for each message shall indicate the state of the message.
22895 22896 22897		ommands take an optional list of messages (<i>msglist</i>) on which to operate, which defaults arrent message. A <i>msglist</i> is a list of message specifications separated by <blank>s, which ade:</blank>
22898	n	Message number n .
22899	+	The next undeleted message, or the next deleted message for the undelete command.
22900 22901	_	The next previous undeleted message, or the next previous deleted message for the undelete command.
22902	•	The current message.
22903	^	The first undeleted message, or the first deleted message for the undelete command.
22904	\$	The last message.
22905	*	All messages.
22906	n-m	An inclusive range of message numbers.
22907 22908	address	All messages from <i>address</i> ; any address as shown in a header summary shall be matchable in this form.
22909	/string	All messages with <i>string</i> in the subject line (case ignored).
22910	:C	All messages of type c , where c shall be one of:
22911		d Deleted messages.
22912		n New messages.
22913		Old messages (any not in state <i>read</i> or <i>new</i>).
22914		r Read messages.
22915		u Unread messages.

Other commands take an optional message (*message*) on which to operate, which defaults to the current message. All of the forms allowed for *msglist* are also allowed for *message*, but if more than one message is specified, only the first shall be operated on.

Other arguments are usually arbitrary strings whose usage depends on the command involved.

Start-Up in mailx

22919

22920 22921

2292222923

22924 22925

22926

22927

22928

22929

22930

22931 22932

22933

22934

22936

22944

22945 22946

22947

22948

22949

22950

22951

22952

22953

22954

At start-up time, *mailx* shall take the following steps in sequence:

- 1. Establish all variables at their stated default values.
- 2. Process command line options, overriding corresponding default values.
 - 3. Import any of the *DEAD*, *EDITOR*, *MBOX*, *LISTER*, *PAGER*, *SHELL*, or *VISUAL* variables that are present in the environment, overriding the corresponding default values.
 - 4. Read *mailx* commands from an unspecified system start-up file, unless the −**n** option is given, to initialize any internal *mailx* variables and aliases.
 - 5. Process the start-up file of *mailx* commands named in the user *MAILRC* variable.

Most regular *mailx* commands are valid inside start-up files, the most common use being to set up initial display options and alias lists. The following commands shall be invalid in the start-up file: !, edit, hold, mail, preserve, reply, Reply, shell, visual, Copy, followup, and Followup. Any errors in the start-up file shall either cause *mailx* to terminate with a diagnostic message and a non-zero status or to continue after writing a diagnostic message, ignoring the remainder of the lines in the start-up file.

22935 A blank line in a start-up file shall be ignored.

Internal Variables in mailx

The following variables are internal *mailx* variables. Each internal variable can be set via the *mailx* **set** command at any time. The **unset** and **set no** *name* commands can be used to erase variables.

In the following list, variables shown as:

22941 variable

represent Boolean values. Variables shown as:

22943 variable=*value*

shall be assigned string or numeric values. For string values, the rules in **Commands in mailx** (on page 596) concerning filenames and quoting shall also apply.

The defaults specified here may be changed by the implementation-defined system start-up file unless the user specifies the $-\mathbf{n}$ option.

All network names whose login name components match shall be treated as identical. This shall cause the *msglist* message specifications to behave similarly. The default shall be **noallnet**. See also the **alternates** command and the **metoo** variable.

append Append messages to the end of the **mbox** file upon termination instead of placing them at the beginning. The default shall be **noappend**. This variable shall not affect the **save** command when saving to **mbox**.

22955 **ask, asksub** Prompt for a subject line on outgoing mail if one is not specified on the command line with the **-s** option. The **ask** and **asksub** forms are synonyms; the system shall

22957 22958 22959 22960		refer to asksub and noasksub in its messages, but shall accept ask and noask as user input to mean asksub and noasksub . It shall not be possible to set both ask and noasksub , or noask and asksub . The default shall be asksub , but no prompting shall be done if standard input is not a terminal.
22961	askbcc	Prompt for the blind copy list. The default shall be noaskbcc .
22962	askcc	Prompt for the copy list. The default shall be noaskcc .
22963 22964	autoprint	Enable automatic writing of messages after delete and undelete commands. The default shall be noautoprint .
22965 22966 22967 22968	bang	Enable the special-case treatment of exclamation marks ('!') in escape command lines; see the escape command and Command Escapes in mailx (on page 604). The default shall be nobang , disabling the expansion of '!' in the <i>command</i> argument to the "! command and the " command escape.</td
22969 22970 22971	cmd=comman	Set the default command to be invoked by the pipe command. The default shall be nocmd .
22972 22973 22974	crt=number	Pipe messages having more than <i>number</i> lines through the command specified by the value of the <i>PAGER</i> variable. The default shall be nocrt . If it is set to null, the value used is implementation-defined.
22975 XSI 22976	debug	Enable verbose diagnostics for debugging. Messages are not delivered. The default shall be nodebug .
22977 22978 22979 22980	dot	When dot is set, a period on a line by itself during message input from a terminal shall also signify end-of-file (in addition to normal end-of-file). The default shall be nodot . If ignoreeof is set (see below), a setting of nodot shall be ignored and the period is the only method to terminate input mode.
22981 22982 22983	escape=c	Set the command escape character to be the character 'c'. By default, the command escape character shall be tilde. If escape is unset, tilde shall be used; if it is set to null, command escaping shall be disabled.
22984	flipr	Reverse the meanings of the ${\bf R}$ and ${\bf r}$ commands. The default shall be ${\bf noflipr}$.
22985 22986 22987 22988 22989 22990 22991 22992 22993	folder=direct	The default directory for saving mail files. User-specified filenames beginning with a plus sign $('+')$ shall be expanded by preceding the filename with this directory name to obtain the real pathname. If <i>directory</i> does not start with a slash $('/')$, the contents of $HOME$ shall be prefixed to it. The default shall be nofolder . If folder is unset or set to null, user-specified filenames beginning with $'+'$ shall refer to files in the current directory that begin with the literal $'+'$ character. See also outfolder below. The folder value need not affect the processing of the files named in $MBOX$ and $DEAD$.
22994 22995	header	Enable writing of the header summary when entering <i>mailx</i> in Receive Mode. The default shall be header .
22996 22997	hold	Preserve all messages that are read in the system mailbox instead of putting them in the mbox save file. The default shall be nohold .
22998	ignore	Ignore interrupts while entering messages. The default shall be noignore .
22999 23000 23001	ignoreeof	Ignore normal end-of-file during message input. Input can be terminated only by entering a period $('\cdot,')$ on a line by itself or by the $\tilde{\ }$. command escape. The default shall be noignoreeof . See also dot above.

23002	indentprefix	x=string
23003 23004		A string that shall be added as a prefix to each line that is inserted into the message by the "m command escape. This variable shall default to one <tab>.</tab>
23005 23006	keep	When a system mailbox, secondary mailbox, or mbox is empty, truncate it to zero length instead of removing it. The default shall be nokeep .
23007 23008 23009	keepsave	Keep the messages that have been saved from the system mailbox into other files in the file designated by the variable <i>MBOX</i> , instead of deleting them. The default shall be nokeepsave .
23010 23011	metoo	Suppress the deletion of the login name of the user from the recipient list when replying to a message or sending to a group. The default shall be nometoo .
23012 XSI 23013 23014 23015 23016	onehop	When responding to a message that was originally sent to several recipients, the other recipient addresses are normally forced to be relative to the originating author's machine for the response. This flag disables alteration of the recipients' addresses, improving efficiency in a network where all machines can send directly to all other machines (that is, one hop away). The default shall be noonehop .
23017 23018 23019	outfolder	Cause the files used to record outgoing messages to be located in the directory specified by the folder variable unless the pathname is absolute. The default shall be nooutfolder . See the record variable.
23020 23021	page	Insert a <form-feed> after each message sent through the pipe created by the pipe command. The default shall be nopage.</form-feed>
23022 23023 23024	prompt=stri	Set the command-mode prompt to <i>string</i> . If <i>string</i> is null or if noprompt is set, no prompting shall occur. The default shall be to prompt with the string "?".
23025 23026	quiet	Refrain from writing the opening message and version when entering <i>mailx</i> . The default shall be noquiet .
23027 23028	record=file	Record all outgoing mail in the file with the pathname <i>file</i> . The default shall be norecord . See also outfolder above.
23029 23030	save	Enable saving of messages in the dead-letter file on interrupt or delivery error. See the variable $DEAD$ for the location of the dead-letter file. The default shall be save .
23031	screen=num	
23032 23033 23034 23035		Set the number of lines in a screenful of headers for the headers and z commands. If screen is not specified, a value based on the terminal type identified by the <i>TERM</i> environment variable, the window size, the baud rate, or some combination of these shall be used.
23036 23037	sendwait	Wait for the background mailer to finish before returning. The default shall be nosendwait .
23038 23039 23040	showto	When the sender of the message was the user who is invoking <i>mailx</i> , write the information from the To: line instead of the From: line in the header summary. The default shall be noshowto .
23041 23042 23043 23044	sign=string	Set the variable inserted into the text of a message when the $\mathbf{\tilde{a}}$ command escape is given. The default shall be nosign . The character sequences '\t' and '\n' shall be recognized in the variable as <tab>s and <newline>s, respectively. (See also $\mathbf{\tilde{i}}$ in Command Escapes in mailx (on page 604).)</newline></tab>
23045 23046	Sign=string	Set the variable inserted into the text of a message when the ${}^{\sim}\!A$ command escape is given. The default shall be noSign . The character sequences '\t' and '\n' shall

be recognized in the variable as <tab>s and <newline>s, respectively.

toplines=number

Set the number of lines of the message to write with the **top** command. The default shall be 5.

Commands in mailx

The following *mailx* commands shall be provided. In the following list, header refers to lines from the message header, as shown in the OUTPUT FILES section. Header-line refers to lines within the header that begin with one or more non-white-space characters, immediately followed by a colon and white space and continuing until the next line beginning with a non-white-space character or an empty line. Header-field refers to the portion of a header line prior to the first colon in that line.

For each of the commands listed below, the command can be entered as the abbreviation (those characters in the Synopsis command word preceding the '['), the full command (all characters shown for the command word, omitting the '[' and ']'), or any truncation of the full command down to the abbreviation. For example, the **exit** command (shown as **ex[it]** in the Synopsis) can be entered as **ex**, **exi**, or **exit**.

The arguments to commands can be quoted, using the following methods:

- An argument can be enclosed between paired double-quotes ("") or single-quotes (''); any white space, shell word expansion, or backslash characters within the quotes shall be treated literally as part of the argument. A double-quote shall be treated literally within single-quotes and *vice versa*. These special properties of the quote marks shall occur only when they are paired at the beginning and end of the argument.
- A backslash outside of the enclosing quotes shall be discarded and the following character treated literally as part of the argument.
- An unquoted backslash at the end of a command line shall be discarded and the next line shall continue the command.

Filenames, where expected, shall be subjected to the following transformations, in sequence:

- If the filename begins with an unquoted plus sign, and the folder variable is defined (see the
 folder variable), the plus sign shall be replaced by the value of the folder variable followed
 by a slash. If the folder variable is unset or is set to null, the filename shall be unchanged.
- Shell word expansions shall be applied to the filename (see Section 2.6 (on page 36)). If more than a single pathname results from this expansion and the command is expecting one file, the effects are unspecified.

Declare Aliases

```
Synopsis: a[lias] [alias [address...]]
q[roup] [alias [address...]]
```

Add the given addresses to the alias specified by alias. The names shall be substituted when alias is used as a recipient address specified by the user in an outgoing message (that is, other recipients addressed indirectly through the **reply** command shall not be substituted in this manner). Mail address alias substitution shall apply only when the alias string is used as a full address; for example, when **hlj** is an alias, *hlj@posix.com* does not trigger the alias substitution. If no arguments are given, write a listing of the current aliases to standard output. If only an alias argument is given, write a listing of the specified alias to standard output. These listings need not reflect the same order of addresses that were entered.

23091 **Declare Alternatives** 23092 Synopsis: alt[ernates] name... (See also the **metoo** command.) Declare a list of alternative names for the user's login. When 23093 23094 responding to a message, these names shall be removed from the list of recipients for the 23095 response. The comparison of names shall be in a case-insensitive manner. With no arguments, **alternates** shall write the current list of alternative names. 23096 **Change Current Directory** 23097 Synopsis: cd [directory] 23098 ch[dir] [directory] 23099 23100 Change directory. If *directory* is not specified, the contents of *HOME* shall be used. **Copy Messages** 23101 c[opy] [file] 23102 Synopsis: 23103 c[opy] [msglist] file 23104 C[opy] [msglist] Copy messages to the file named by the pathname *file* without marking the messages as saved. 23105 Otherwise, it shall be equivalent to the **save** command. 23106 In the capitalized form, save the specified messages in a file whose name is derived from the 23107 23108 author of the message to be saved, without marking the messages as saved. Otherwise, it shall be equivalent to the **Save** command. 23109 **Delete Messages** 23110 23111 d[elete] [msglist] Synopsis: 23112 Mark messages for deletion from the mailbox. The deletions shall not occur until mailx quits (see 23113 the **quit** command) or changes mailboxes (see the **folder** command). If **autoprint** is set and there 23114 are messages remaining after the delete command, the current message shall be written as 23115 described for the **print** command (see the **print** command); otherwise, the *mailx* prompt shall be 23116 written. **Discard Header Fields** 23117 di[scard] [header-field...] 23118 Synopsis: 23119 iq[nore] [header-field...] Suppress the specified header fields when writing messages. Specified header-fields shall be 23120 23121 added to the list of suppressed header fields. Examples of header fields to ignore are status and 23122

Suppress the specified header fields when writing messages. Specified header-fields shall be added to the list of suppressed header fields. Examples of header fields to ignore are **status** and **cc**. The fields shall be included when the message is saved. The **Print** and **Type** commands shall override this command. The comparison of header fields shall be in a case-insensitive manner. If no arguments are specified, write a list of the currently suppressed header fields to standard output; the listing need not reflect the same order of header fields that were entered.

If both **retain** and **discard** commands are given, **discard** commands shall be ignored.

23123 23124

2312523126

23127	Delete Messages and Display
23128 23129	Synopsis: dp [msglist] dt [msglist]
23130 23131 23132 23133 23134	Delete the specified messages as described for the delete command, except that the autoprint variable shall have no effect, and the current message shall be written only if it was set to a message after the last message deleted by the command. Otherwise, an informational message to the effect that there are no further messages in the mailbox shall be written, followed by the <i>mailx</i> prompt.
23135	Echo a String
23136	Synopsis: ec[ho] string
23137	Echo the given strings, equivalent to the shell <i>echo</i> utility.
23138	Edit Messages
23139	Synopsis: e[dit] [msglist]
23140 23141 23142	Edit the given messages. The messages shall be placed in a temporary file and the utility named by the <i>EDITOR</i> variable is invoked to edit each file in sequence. The default <i>EDITOR</i> is unspecified.
23143	The edit command does not modify the contents of those messages in the mailbox.
23144	Exit
23145 23146	Synopsis: ex[it] x[it]
23147 23148	Exit from <i>mailx</i> without changing the mailbox. No messages shall be saved in the mbox (see also quit).
23149	Change Folder
23150 23151	Synopsis: fi[le] [file] fold[er] [file]
23152 23153 23154	Quit (see the quit command) from the current file of messages and read in the file named by the pathname <i>file</i> . If no argument is given, the name and status of the current mailbox shall be written.
23155 23156	Several unquoted special characters shall be recognized when used as <i>file</i> names, with the following substitutions:
23157	% The system mailbox for the invoking user.
23158	%user The system mailbox for user.
23159	# The previous file.
23160	& The current mbox .
23161	+file The named file in the folder directory. (See the folder variable.)
23162	The default file shall be the current mailbox.

23163	Display List of Folders
23164	Synopsis: folders
23165 23166	Write the names of the files in the directory set by the folder variable. The command specified by the <i>LISTER</i> environment variable shall be used (see the ENVIRONMENT VARIABLES section).
23167	Follow Up Specified Messages
23168 23169	Synopsis: fo[llowup] [message] F[ollowup] [msglist]
23170 23171	In the lowercase form, respond to a message, recording the response in a file whose name is derived from the author of the message. See also the save and copy commands and outfolder .
23172 23173 23174 23175	In the capitalized form, respond to the first message in the <i>msglist</i> , sending the message to the author of each message in the <i>msglist</i> . The subject line shall be taken from the first message and the response shall be recorded in a file whose name is derived from the author of the first message. See also the Save and Copy commands and outfolder .
23176	Both forms shall override the record variable, if set.
23177	Display Header Summary for Specified Messages
23178	Synopsis: f[rom] [msglist]
23179	Write the header summary for the specified messages.
	write the neutral summary for the specified messages.
23180	Display Header Summary
23180 23181	
	Display Header Summary
23181 23182 23183 23184 23185	Display Header Summary Synopsis: h [eaders] [message] Write the page of headers that includes the message specified. If the message argument is not specified, the current message shall not change. However, if the message argument is specified, the current message shall become the message that appears at the top of the page of headers that includes the message specified. The screen variable sets the number of headers per page. See
23181 23182 23183 23184 23185 23186	Display Header Summary Synopsis: h [eaders] [message] Write the page of headers that includes the message specified. If the message argument is not specified, the current message shall not change. However, if the message argument is specified, the current message shall become the message that appears at the top of the page of headers that includes the message specified. The screen variable sets the number of headers per page. See also the z command.
23181 23182 23183 23184 23185 23186 23187	Display Header Summary Synopsis: h [eaders] [message] Write the page of headers that includes the message specified. If the message argument is not specified, the current message shall not change. However, if the message argument is specified, the current message shall become the message that appears at the top of the page of headers that includes the message specified. The screen variable sets the number of headers per page. See also the z command. Help Synopsis: hel[p]
23181 23182 23183 23184 23185 23186 23187 23188 23189	Display Header Summary Synopsis: h[eaders] [message] Write the page of headers that includes the message specified. If the message argument is not specified, the current message shall not change. However, if the message argument is specified, the current message shall become the message that appears at the top of the page of headers that includes the message specified. The screen variable sets the number of headers per page. See also the z command. Help Synopsis: hel[p] ?
23181 23182 23183 23184 23185 23186 23187 23188 23189 23190	Display Header Summary Synopsis: h [eaders] [message] Write the page of headers that includes the message specified. If the message argument is not specified, the current message shall not change. However, if the message argument is specified, the current message shall become the message that appears at the top of the page of headers that includes the message specified. The screen variable sets the number of headers per page. See also the z command. Help Synopsis: hel[p] ? Write a summary of commands.

23198	Execute Commands Conditionally
23199 23200 23201 23202 23203	Synopsis: i[f] s r mail-commands el[se] mail-commands en[dif]
23204 23205 23206	Execute commands conditionally, where if s executes the following <i>mail-commands</i> , up to an else or endif , if the program is in Send Mode, and if r shall cause the <i>mail-commands</i> to be executed only in Receive Mode.
23207	List Available Commands
23208	Synopsis: l[ist]
23209	Write a list of all commands available. No explanation shall be given.
23210	Mail a Message
23211	Synopsis: m[ail] address
23212	Mail a message to the specified addresses or aliases.
23213	Direct Messages to mbox
23214	Synopsis: mb[ox] [msglist]
23215 23216	Arrange for the given messages to end up in the mbox save file when <i>mailx</i> terminates normally. See <i>MBOX</i> . See also the exit and quit commands.
23217	Process Next Specified Message
23218	Synopsis: n[ext] [message]
23219 23220 23221 23222 23223 23224	If the current message has not been written (for example, by the print command) since <i>mailx</i> started or since any other message was the current message, behave as if the print command was entered. Otherwise, if there is an undeleted message after the current message, make it the current message and behave as if the print command was entered. Otherwise, an informational message to the effect that there are no further messages in the mailbox shall be written, followed by the <i>mailx</i> prompt.
23225	Pipe Message
23226 23227	Synopsis: pi[pe] [[msglist] command] [[msglist] command]
23228 23229 23230 23231 23232	Pipe the messages through the given <i>command</i> by invoking the command interpreter specified by <i>SHELL</i> with two arguments: $-\mathbf{c}$ and <i>command</i> . (See also $sh - \mathbf{c}$.) The application shall ensure that the command is given as a single argument. Quoting, described previously, can be used to accomplish this. If no arguments are given, the current message shall be piped through the command specified by the value of the cmd variable. If the page variable is set, a <form-feed></form-feed>

shall be inserted after each message.

23233

23234	Display Message with Headers
23235 23236	Synopsis: P[rint] [msglist] T[ype] [msglist]
23237 23238 23239 23240	Write the specified messages, including all header lines, to standard output. Override suppression of lines by the discard , ignore , and retain commands. If crt is set, the messages longer than the number of lines specified by the crt variable shall be paged through the command specified by the <i>PAGER</i> environment variable.
23241	Display Message
23242 23243	Synopsis: p[rint] [msglist] t[ype] [msglist]
23244 23245 23246	Write the specified messages to standard output. If crt is set, the messages longer than the number of lines specified by the crt variable shall be paged through the command specified by the <i>PAGER</i> environment variable.
23247	Quit
23248 23249	Synopsis: q[uit] end-of-file
23250 23251 23252 23253	Terminate <i>mailx</i> , storing messages that were read in mbox (if the current mailbox is the system mailbox and unless hold is set), deleting messages that have been explicitly saved (unless keepsave is set), discarding messages that have been deleted, and saving all remaining messages in the mailbox.
23254	Reply to a Message List
23255 23256	Synopsis: R[eply] [msglist] R[espond] [msglist]
23257 23258 23259 23260	Mail a reply message to the sender of each message in the <i>msglist</i> . The subject line shall be formed by concatenating Re : <space> (unless it already begins with that string) and the subject from the first message. If record is set to a filename, the response shall be saved at the end of that file.</space>
23261	See also the flipr variable.
23262	Reply to a Message
23263 23264	Synopsis: r[eply] [message] r[espond] [message]
23265 23266 23267 23268	Mail a reply message to all recipients included in the header of the message. The subject line shall be formed by concatenating Re: <space> (unless it already begins with that string) and the subject from the message. If record is set to a filename, the response shall be saved at the end of that file.</space>
23269	See also the flipr variable.

23270	Retain Header Fields
23271	Synopsis: ret[ain] [header-field]
23272 23273 23274 23275	Retain the specified header fields when writing messages. This command shall override all discard and ignore commands. The comparison of header fields shall be in a case-insensitive manner. If no arguments are specified, write a list of the currently retained header fields to standard output; the listing need not reflect the same order of header fields that were entered.
23276	Save Messages
23277 23278 23279	Synopsis: s[ave] [file] s[ave] [msglist] file S[ave] [msglist]
23280 23281 23282 23283 23284	Save the specified messages in the file named by the pathname <i>file</i> , or the mbox if the <i>file</i> argument is omitted. The file shall be created if it does not exist; otherwise, the messages shall be appended to the file. The message shall be put in the state <i>saved</i> , and shall behave as specified in the description of the <i>saved</i> state when the current mailbox is exited by the quit or file command.
23285 23286 23287 23288	In the capitalized form, save the specified messages in a file whose name is derived from the author of the first message. The name of the file shall be taken to be the author's name with all network addressing stripped off. See also the Copy , followup , and Followup commands and outfolder variable.
23289	Set Variables
23290	Synopsis: se[t] [name[=[string]]] [name=number] [noname]
23291 23292 23293 23294 23295 23296	Define one or more variables called <i>name</i> . The variable can be given a null, string, or numeric value. Quoting and backslash escapes can occur anywhere in <i>string</i> , as described previously, as if the <i>string</i> portion of the argument were the entire argument. The forms <i>name</i> and <i>name</i> = shall be equivalent to <i>name</i> ="" for variables that take string values. The set command without arguments shall write a list of all defined variables and their values. The no <i>name</i> form shall be equivalent to unset <i>name</i> .
23297	Invoke a Shell
23298	Synopsis: sh[ell]
23299	Invoke an interactive command interpreter (see also SHELL).
23300	Display Message Size
23301	Synopsis: si[ze] [msglist]
23302	Write the size in bytes of each of the specified messages.
23303	Read mailx Commands From a File
23304	Synopsis: so[urce] file
23305	Read and execute commands from the file named by the pathname file and return to command

mode.

23306

23307	Display Beginning of Messages
23308	Synopsis: to[p] [msglist]
23309 23310	Write the top few lines of each of the specified messages. If the toplines variable is set, it is taken as the number of lines to write. The default shall be 5.
23311	Touch Messages
23312	Synopsis: tou[ch] [msglist]
23313 23314	Touch the specified messages. If any message in <i>msglist</i> is not specifically deleted nor saved in a file, it shall be placed in the mbox upon normal termination. See exit and quit .
23315	Delete Aliases
23316	Synopsis: una[lias] [alias]
23317	Delete the specified alias names. If a specified alias does not exist, the results are unspecified.
23318	Undelete Messages
23319	Synopsis: u[ndelete] [msglist]
23320 23321 23322	Change the state of the specified messages from deleted to read. If autoprint is set, the last message of those restored shall be written. If <i>msglist</i> is not specified, the message shall be selected as follows:
23323 23324	• If there are any deleted messages that follow the current message, the first of these shall be chosen.
23325	• Otherwise, the last deleted message that also precedes the current message shall be chosen.
23326	Unset Variables
23327	Synopsis: uns[et] name
23328	Cause the specified variables to be erased.
23329	Edit Message with Full-Screen Editor
23330	Synopsis: v[isual] [msglist]
23331	Edit the given messages with a screen editor. Each message shall be placed in a temporary file,
23332	and the utility named by the VISUAL variable shall be invoked to edit each file in sequence. The
23333	default editor shall be <i>vi</i> .
23334	The visual command does not modify the contents of those messages in the mailbox.
23335	Write Messages to a File
23336	Synopsis: w[rite] [msglist] file
23337 23338	Write the given messages to the file specified by the pathname <i>file</i> , minus the message header. Otherwise, it shall be equivalent to the save command.

23339	Scroll Header Display
23340	<i>Synopsis</i> : z [+ -]
23341 23342 23343	Scroll the header display forward (if $'+'$ is specified or if no option is specified) or backward (if $'-'$ is specified) one screenful. The number of headers written shall be set by the screen variable.
23344	Invoke Shell Command
23345	Synopsis: ! command
23346 23347 23348	Invoke the command interpreter specified by <i>SHELL</i> with two arguments: $-\mathbf{c}$ and <i>command</i> . (See also sh $-\mathbf{c}$.) If the bang variable is set, each unescaped occurrence of '!' in <i>command</i> shall be replaced with the command executed by the previous! command or "! command escape.
23349	Null Command
23350	Synopsis: # comment
23351	This null command (comment) shall be ignored by <i>mailx</i> .
23352	Display Current Message Number
23353	Synopsis: =
23354	Write the current message number.
23355	Command Escapes in mailx
23356 23357 23358	The following commands can be entered only from input mode, by beginning a line with the escape character (by default, tilde ($'^{\sim}$)). See the escape variable description for changing this special character. The format for the commands shall be:
23359	<pre><escape-character><command-char><separator>[<arguments>]</arguments></separator></command-char></escape-character></pre>
23360	where the < <i>separator</i> > can be zero or more <blank>s.</blank>
23361 23362 23363 23364	In the following descriptions, the application shall ensure that the argument <i>command</i> (but not <i>mailx-command</i>) is a shell command string. Any string acceptable to the command interpreter specified by the <i>SHELL</i> variable when it is invoked as <i>SHELL</i> – <i>c command_string</i> shall be valid. The command can be presented as multiple arguments (that is, quoting is not required).
23365 23366	Command escapes that are listed with <i>msglist</i> or <i>mailx-command</i> arguments are invalid in Send Mode and produce unspecified results.
23367 23368 23369 23370	~! command Invoke the command interpreter specified by SHELL with two arguments: -c and command; and then return to input mode. If the bang variable is set, each unescaped occurrence of '!' in command shall be replaced with the command executed by the previous! command or ~! command escape.
23371	~. Simulate end-of-file (terminate message input).
23372 23373	~: mailx-command, ~_ mailx-command Perform the command-level request.
23374	~? Write a summary of command escapes.
23375	~A This shall be equivalent to ~i Sign.
23376	~a This shall be equivalent to ~i sign.

23377	~b name	Add the names to the blind carbon copy (Bcc) list.
23378	~c name	Add the <i>names</i> to the carbon copy (Cc) list.
23379	~ d	Read in the dead-letter file. See <i>DEAD</i> for a description of this file.
23380 23381	~e	Invoke the editor, as specified by the $\it EDITOR$ environment variable, on the partial message.
23382 23383 23384 23385	~f [msglist]	Forward the specified messages. The specified messages shall be inserted into the current message without alteration. This command escape also shall insert message headers into the message with field selection affected by the discard , ignore , and retain commands.
23386 23387 23388	~F [msglist]	This shall be the equivalent of the "f command escape, except that all headers shall be included in the message, regardless of previous discard , ignore , and retain commands.
23389 23390 23391	~h	If standard input is a terminal, prompt for a Subject line and the To , Cc , and Bcc lists. Other implementation-defined headers may also be presented for editing. If the field is written with an initial value, it can be edited as if it had just been typed.
23392 23393	~i string	Insert the value of the named variable, followed by a <newline>, into the text of the message. If the string is unset or null, the message shall not be changed.</newline>
23394 23395 23396 23397	~m [msglist]	Insert the specified messages into the message, prefixing non-empty lines with the string in the indentprefix variable. This command escape also shall insert message headers into the message, with field selection affected by the discard , ignore , and retain commands.
23398 23399 23400	~M [msglist]	This shall be the equivalent of the \tilde{m} command escape, except that all headers shall be included in the message, regardless of previous discard , ignore , and retain commands.
23401 23402 23403	~ p	Write the message being entered. If the message is longer than crt lines (see Internal Variables in mailx (on page 593)), the output shall be paginated as described for the <i>PAGER</i> variable.
23404 23405 23406	~ q	Quit (see the quit command) from input mode by simulating an interrupt. If the body of the message is not empty, the partial message shall be saved in the deadletter file. See $DEAD$ for a description of this file.
23407 23408 23409 23410 23411 23412	~ r file, ~< file	e, "r !command, "< !command Read in the file specified by the pathname file. If the argument begins with an exclamation mark ('!'), the rest of the string shall be taken as an arbitrary system command; the command interpreter specified by SHELL shall be invoked with two arguments: -c and command. The standard output of command shall be inserted into the message.
23413	~s string	Set the subject line to <i>string</i> .
23414	~t name	Add the given names to the To list.
23415 23416	~ V	Invoke the full-screen editor, as specified by the \emph{VISUAL} environment variable, on the partial message.
23417 23418 23419	~w file	Write the partial message, without the header, onto the file named by the pathname <i>file</i> . The file shall be created or the message shall be appended to it if the file exists.

23420 ~**x** Exit as with ~**q**, except the message shall not be saved in the dead-letter file.

23421 ~| command Pipe the body of the message through the given command by invoking the command interpreter specified by SHELL with two arguments: -**c** and command.

23423 If the command returns a successful exit status, the standard output of the command shall replace the message. Otherwise, the message shall remain unchanged. If the command fails, an error message giving the exit status shall be written.

23427 EXIT STATUS

23436

23437 23438

23439 23440

23441

23445

23446

23447

23448

23449

23450

23451

2345223453

23454

23455

23456

23457

23458

23459

23460

23461 23462

23463

When the $-\mathbf{e}$ option is specified, the following exit values are returned:

- 23429 0 Mail was found.
- 23430 >0 Mail was not found or an error occurred.
- Otherwise, the following exit values are returned:
- 23432 0 Successful completion; note that this status implies that all messages were *sent*, but it gives no assurances that any of them were actually *delivered*.
- 23434 >0 An error occurred.

23435 CONSEQUENCES OF ERRORS

When in input mode (Receive Mode) or Send Mode:

- If an error is encountered processing a command escape (see **Command Escapes in mailx** (on page 604)), a diagnostic message shall be written to standard error, and the message being composed may be modified, but this condition shall not prevent the message from being sent.
- Other errors shall prevent the sending of the message.
- 23442 When in command mode:
- 23443 Default.

23444 APPLICATION USAGE

Delivery of messages to remote systems requires the existence of communication paths to such systems. These need not exist.

Input lines are limited to {LINE_MAX} bytes, but mailers between systems may impose more severe line-length restrictions. This volume of IEEE Std 1003.1-2001 does not place any restrictions on the length of messages handled by *mailx*, and for delivery of local messages the only limitations should be the normal problems of available disk space for the target mail file. When sending messages to external machines, applications are advised to limit messages to less than 100 000 bytes because some mail gateways impose message-length restrictions.

The format of the system mailbox is intentionally unspecified. Not all systems implement system mailboxes as flat files, particularly with the advent of multimedia mail messages. Some system mailboxes may be multiple files, others records in a database. The internal format of the messages themselves is specified with the historical format from Version 7, but only after the messages have been saved in some file other than the system mailbox. This was done so that many historical applications expecting text-file mailboxes are not broken.

Some new formats for messages can be expected in the future, probably including binary data, bit maps, and various multimedia objects. As described here, *mailx* is not prohibited from handling such messages, but it must store them as text files in secondary mailboxes (unless some extension, such as a variable or command line option, is used to change the stored format). Its method of doing so is implementation-defined and might include translating the data into

23464 text file-compatible or readable form or omitting certain portions of the message from the stored output.

The **discard** and **ignore** commands are not inverses of the **retain** command. The **retain** command discards all header-fields except those explicitly retained. The **discard** command keeps all header-fields except those explicitly discarded. If headers exist on the retained header list, **discard** and **ignore** commands are ignored.

23470 EXAMPLES

23471 None.

23472 RATIONALE

 The standard developers felt strongly that a method for applications to send messages to specific users was necessary. The obvious example is a batch utility, running non-interactively, that wishes to communicate errors or results to a user. However, the actual format, delivery mechanism, and method of reading the message are clearly beyond the scope of this volume of IEEE Std 1003.1-2001.

The intent of this command is to provide a simple, portable interface for sending messages non-interactively. It merely defines a "front-end" to the historical mail system. It is suggested that implementations explicitly denote the sender and recipient in the body of the delivered message. Further specification of formats for either the message envelope or the message itself were deliberately not made, as the industry is in the midst of changing from the current standards to a more internationalized standard and it is probably incorrect, at this time, to require either one.

Implementations are encouraged to conform to the various delivery mechanisms described in the CCITT X.400 standards or to the equivalent Internet standards, described in Internet Request for Comment (RFC) documents RFC 819, RFC 822, RFC 920, RFC 921, and RFC 1123.

Many historical systems modified each body line that started with **From** by prefixing the 'F' with '>'. It is unnecessary, but allowed, to do that when the string does not follow a blank line because it cannot be confused with the next header.

The **edit** and **visual** commands merely edit the specified messages in a temporary file. They do not modify the contents of those messages in the mailbox; such a capability could be added as an extension, such as by using different command names.

The restriction on a subject line being {LINE_MAX}-10 bytes is based on the historical format that consumes 10 bytes for **Subject**: and the trailing <newline>. Many historical mailers that a message may encounter on other systems are not able to handle lines that long, however.

Like the utilities *logger* and *lp*, *mailx* admittedly is difficult to test. This was not deemed sufficient justification to exclude this utility from this volume of IEEE Std 1003.1-2001. It is also arguable that it is, in fact, testable, but that the tests themselves are not portable.

When *mailx* is being used by an application that wishes to receive the results as if none of the User Portability Utilities option features were supported, the *DEAD* environment variable must be set to /dev/null. Otherwise, it may be subject to the file creations described in *mailx* ASYNCHRONOUS EVENTS. Similarly, if the *MAILRC* environment variable is not set to /dev/null, historical versions of *mailx* and *Mail* read initialization commands from a file before processing begins. Since the initialization that a user specifies could alter the contents of messages an application is trying to send, such applications must set *MAILRC* to /dev/null.

The description of *LC_TIME* uses "may affect" because many historical implementations do not or cannot manipulate the date and time strings in the incoming mail headers. Some headers found in incoming mail do not have enough information to determine the timezone in which the mail originated, and, therefore, *mailx* cannot convert the date and time strings into the internal form that then is parsed by routines like *strftime*() that can take *LC_TIME* settings into account.

23511	Changing all these times to a user-specified format is allowed, but not required.
23512 23513 23514 23515 23516 23517 23518	The paginator selected when <i>PAGER</i> is null or unset is partially unspecified to allow the System V historical practice of using <i>pg</i> as the default. Bypassing the pagination function, such as by declaring that <i>cat</i> is the paginator, would not meet with the intended meaning of this description. However, any "portable user" would have to set <i>PAGER</i> explicitly to get his or her preferred paginator on all systems. The paginator choice was made partially unspecified, unlike the <i>VISUAL</i> editor choice (mandated to be <i>vi</i>) because most historical pagers follow a common theme of user input, whereas editors differ dramatically.
23519 23520	Options to specify addresses as cc (carbon copy) or bcc (blind carbon copy) were considered to be format details and were omitted.
23521 23522 23523	A zero exit status implies that all messages were <i>sent</i> , but it gives no assurances that any of them were actually <i>delivered</i> . The reliability of the delivery mechanism is unspecified and is an appropriate marketing distinction between systems.
23524 23525 23526	In order to conform to the Utility Syntax Guidelines, a solution was required to the optional <i>file</i> option-argument to $-\mathbf{f}$. By making <i>file</i> an operand, the guidelines are satisfied and users remain portable. However, it does force implementations to support usage such as:
23527	mailx -fin mymail.box
23528 23529 23530	The no <i>name</i> method of unsetting variables is not present in all historical systems, but it is in System V and provides a logical set of commands corresponding to the format of the display of options from the <i>mailx set</i> command without arguments.
23531 23532	The ask and asksub variables are the names selected by BSD and System V, respectively, for the same feature. They are synonyms in this volume of IEEE Std 1003.1-2001.
23533 23534	The <i>mailx echo</i> command was not documented in the BSD version and has been omitted here because it is not obviously useful for interactive users.
23535 23536 23537 23538	The default prompt on the System V <i>mailx</i> is a question mark, on BSD <i>Mail</i> an ampersand. Since this volume of IEEE Std 1003.1-2001 chose the <i>mailx</i> name, it kept the System V default, assuming that BSD users would not have difficulty with this minor incompatibility (that they can override).
23539 23540 23541 23542	The meanings of r and R are reversed between System V <i>mailx</i> and SunOS <i>Mail</i> . Once again, since this volume of IEEE Std 1003.1-2001 chose the <i>mailx</i> name, it kept the System V default, but allows the SunOS user to achieve the desired results using flipr , an internal variable in System V <i>mailx</i> , although it has not been documented in the SVID.
23543 23544	The indentprefix variable, the retain and unalias commands, and the F and M command escapes were adopted from 4.3 BSD <i>Mail</i> .
23545 23546 23547 23548	The version command was not included because no sufficiently general specification of the version information could be devised that would still be useful to a portable user. This command name should be used by suppliers who wish to provide version information about the <i>mailx</i> command.
23549 23550	The "implementation-specific (unspecified) system start-up file" historically has been named /etc/mailx.rc, but this specific name and location are not required.
23551 23552 23553	The intent of the wording for the next command is that if any command has already displayed the current message it should display a following message, but, otherwise, it should display the current message. Consider the command sequence:
23554 23555	next 3 delete 3

23556	next
23557 23558 23559 23560 23561 23562 23563	where the autoprint option was not set. The normative text specifies that the second next command should display a message following the third message, because even though the current message has not been displayed since it was set by the delete command, it has been displayed since the current message was anything other than message number 3. This does not always match historical practice in some implementations, where the command file address followed by next (or the default command) would skip the message for which the user had searched.
	URE DIRECTIONS
23565	None.
23566 SEE 23567	Chapter 2 (on page 29), ed, ls, more, vi
23568 CHA 23569	ANGE HISTORY First released in Issue 2.
23570 Issue 23571 23572 23573	The description of the <i>EDITOR</i> environment variable is changed to indicate that <i>ed</i> is the default editor if this variable is not set. In previous issues, this default was not stated explicitly at this point but was implied further down in the text.
23574	The FUTURE DIRECTIONS section is added.
23575 Issue 23576 23577	The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:
23578	• The – F option is added.
23579	 The allnet, debug, and sendwait internal variables are added.
23580	• The C, ec, fo, F, and S mailx commands are added.
23581 23582	In the DESCRIPTION and ENVIRONMENT VARIABLES sections, text stating "HOME directory" is replaced by "directory referred to by the HOME environment variable".
23583 23584 23585 23586	The <i>mailx</i> utility is aligned with the IEEE P1003.2b draft standard, which includes various clarifications to resolve IEEE PASC Interpretations submitted for the ISO POSIX-2:1993 standard. In particular, the changes here address IEEE PASC Interpretations 1003.2 #10, #11, #103, #106, #108, #114, #115, #122, and #129.
23587	The normative text is reworded to avoid use of the term "must" for application requirements.
23588	The TZ entry is added to the ENVIRONMENT VARIABLES section.
23589 23590	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/32 is applied, applying a change to the EXTENDED DESCRIPTION, raised by IEEE PASC Interpretation 1003.2 #122, which was everlooked in the first version of IEEE Std 1003.1-2001

overlooked in the first version of IEEE Std 1003.1-2001.

23591

make Utilities

```
23592 NAME
```

23599

23600 23601

23602

23603

23604

23605

23606

23607 23608

23609

23610

23613 23614

23615

23616

23617

23618

23620 23621

23622

23593 make — maintain, update, and regenerate groups of programs (**DEVELOPMENT**)

23594 SYNOPSIS

```
23595 SD make [-einpqrst] [-f makefile]...[ -k | -S] [macro=value]...
23596 [target_name...]
23597
```

23598 **DESCRIPTION**

The *make* utility shall update files that are derived from other files. A typical case is one where object files are derived from the corresponding source files. The *make* utility examines time relationships and shall update those derived files (called targets) that have modified times earlier than the modified times of the files (called prerequisites) from which they are derived. A description file (makefile) contains a description of the relationships between files, and the commands that need to be executed to update the targets to reflect changes in their prerequisites. Each specification, or rule, shall consist of a target, optional prerequisites, and optional commands to be executed when a prerequisite is newer than the target. There are two types of rule:

- 1. Inference rules, which have one target name with at least one period (' . ') and no slash (' /')
- 2. Target rules, which can have more than one target name

In addition, *make* shall have a collection of built-in macros and inference rules that infer prerequisite relationships to simplify maintenance of programs.

To receive exactly the behavior described in this section, the user shall ensure that a portable makefile shall:

- Include the special target .POSIX
- Omit any special target reserved for implementations (a leading period followed by uppercase letters) that has not been specified by this section

The behavior of *make* is unspecified if either or both of these conditions are not met.

23619 OPTIONS

The *make* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported:

- 23623 e Cause environment variables, including those with null values, to override macro assignments within makefiles.
- 23625 f makefile Specify a different makefile. The argument makefile is a pathname of a description file, which is also referred to as the makefile. A pathname of '-' shall denote the standard input. There can be multiple instances of this option, and they shall be processed in the order specified. The effect of specifying the same optionargument more than once is unspecified.
- 23630 —i Ignore error codes returned by invoked commands. This mode is the same as if the special target .IGNORE were specified without prerequisites.
- Continue to update other targets that do not depend on the current target if a non-ignored error occurs while executing the commands to bring a target up-to-date.
- Write commands that would be executed on standard output, but do not execute them. However, lines with a plus sign ('+') prefix shall be executed. In this mode,

23636		lines with an at sign ('@') character prefix shall be written to standard output.
23637 23638	-p	Write to standard output the complete set of macro definitions and target descriptions. The output format is unspecified.
23639 23640 23641 23642	-q	Return a zero exit value if the target file is up-to-date; otherwise, return an exit value of 1. Targets shall not be updated if this option is specified. However, a makefile command line (associated with the targets) with a plus sign $('+')$ prefix shall be executed.
23643	-r	Clear the suffix list and do not use the built-in rules.
23644 23645	-S	Terminate <i>make</i> if an error occurs while executing the commands to bring a target up-to-date. This shall be the default and the opposite of $-\mathbf{k}$.
23646 23647 23648	−s	Do not write makefile command lines or touch messages (see $-t$) to standard output before executing. This mode shall be the same as if the special target .SILENT were specified without prerequisites.
23649 23650 23651 23652 23653 23654 23655	-t	Update the modification time of each target as though a <i>touch target</i> had been executed. Targets that have prerequisites but no commands (see Target Rules (on page 614)), or that are already up-to-date, shall not be touched in this manner. Write messages to standard output for each target file indicating the name of the file and that it was touched. Normally, the <i>makefile</i> command lines associated with each target are not executed. However, a command line with a plus sign ('+') prefix shall be executed.
23656 23657 23658 23659 23660	options spec on the <i>make</i> specified sh	is specified in the <i>MAKEFLAGS</i> environment variable shall be evaluated before any cified on the <i>make</i> utility command line. If the $-\mathbf{k}$ and $-\mathbf{S}$ options are both specified utility command line or by the <i>MAKEFLAGS</i> environment variable, the last option all take precedence. If the $-\mathbf{f}$ or $-\mathbf{p}$ options appear in the <i>MAKEFLAGS</i> environment result is undefined.
23661 OPERA 23662		ng operands shall be supported:
23663 23664 23665	target_name	Target names, as defined in the EXTENDED DESCRIPTION section. If no target is specified, while <i>make</i> is processing the makefiles, the first target that <i>make</i> encounters that is not a special target or an inference rule shall be used.
23666	macro=value	Macro definitions, as defined in Macros (on page 616).
23667 23668		_name and macro=value operands are intermixed on the make utility command line, re unspecified.
23669 STDIN 23670 23671	The standar	d input shall be used only if the <i>makefile</i> option-argument is $'-'$. See the INPUT n.
23672 INPUT 23673 23674	The input fil	le, otherwise known as the makefile, is a text file containing rules, macro definitions, nts. See the EXTENDED DESCRIPTION section.
23675 ENVIR 23676	ONMENT VA	ARIABLES ng environment variables shall affect the execution of <i>make</i> :
23677 23678 23679 23680	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)

23681 LC_ALL If set to a non-empty string value, override the values of all the other 23682 internationalization variables. 23683 LC_CTYPE 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 23684 23685 arguments and input files). LC_MESSAGES 23686 Determine the locale that should be used to affect the format and contents of 23687 diagnostic messages written to standard error. 23688 MAKEFLAGS 23689 This variable shall be interpreted as a character string representing a series of 23690 option characters to be used as the default options. The implementation shall 23691 accept both of the following formats (but need not accept them when intermixed): 23692 The characters are option letters without the leading hyphens or <blank> 23693 separation used on a *make* utility command line. 23694 • The characters are formatted in a manner similar to a portion of the *make* utility 23695 23696 command line: options are preceded by hyphens and

blank>-separated as described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, 23697 Utility Syntax Guidelines. The macro=value macro definition operands can also 23698 be included. The difference between the contents of MAKEFLAGS and the make 23699 utility command line is that the contents of the variable shall not be subjected 23700 23701 to the word expansions (see Section 2.6 (on page 36)) associated with parsing the command line values. 23702 23703 XSI NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. **PROJECTDIR** 23704 XSI 23705 Provide a directory to be used to search for SCCS files not found in the current directory. In all of the following cases, the search for SCCS files is made in the 23706 directory SCCS in the identified directory. If the value of PROJECTDIR begins 23707 with a slash, it shall be considered an absolute pathname; otherwise, the value of 23708 23709 **PROJECTIOIR** is treated as a user name and that user's initial working directory shall be examined for a subdirectory **src** or **source**. If such a directory is found, it 23710 shall be used. Otherwise, the value is used as a relative pathname. 23711 23712 If PROJECTDIR is not set or has a null value, the search for SCCS files shall be made in the directory **SCCS** in the current directory. 23713 23714 The setting of *PROJECTDIR* affects all files listed in the remainder of this utility description for files with a component named SCCS. 23715 The value of the SHELL environment variable shall not be used as a macro and shall not be 23716 modified by defining the SHELL macro in a makefile or on the command line. All other 23717 environment variables, including those with null values, shall be used as macros, as defined in 23718 23719 **Macros** (on page 616). 23720 ASYNCHRONOUS EVENTS If not already ignored, make shall trap SIGHUP, SIGTERM, SIGINT, and SIGQUIT and remove 23721 the current target unless the target is a directory or the target is a prerequisite of the special 23722 target .**PRECIOUS** or unless one of the -n, -p, or -q options was specified. Any targets removed 23723 in this manner shall be reported in diagnostic messages of unspecified format, written to 23724

signals.

23725 23726 standard error. After this cleanup process, if any, make shall take the standard action for all other

23727 STDOUT

The *make* utility shall write all commands to be executed to standard output unless the –s option was specified, the command is prefixed with an at sign, or the special target .SILENT has either the current target as a prerequisite or has no prerequisites. If *make* is invoked without any work needing to be done, it shall write a message to standard output indicating that no action was taken. If the –t option is present and a file is touched, *make* shall write to standard output a message of unspecified format indicating that the file was touched, including the filename of the file.

23735 STDERR

23747

23748 23749

2375023751

2375223753

23754

23756

23757

23758

23736 The standard error shall be used only for diagnostic messages.

23737 OUTPUT FILES

Files can be created when the -t option is present. Additional files can also be created by the utilities invoked by *make*.

23740 EXTENDED DESCRIPTION

The *make* utility attempts to perform the actions required to ensure that the specified targets are up-to-date. A target is considered out-of-date if it is older than any of its prerequisites or if it does not exist. The *make* utility shall treat all prerequisites as targets themselves and recursively ensure that they are up-to-date, processing them in the order in which they appear in the rule. The *make* utility shall use the modification times of files to determine whether the corresponding targets are out-of-date.

After *make* has ensured that all of the prerequisites of a target are up-to-date and if the target is out-of-date, the commands associated with the target entry shall be executed. If there are no commands listed for the target, the target shall be treated as up-to-date.

Makefile Syntax

A makefile can contain rules, macro definitions (see **Macros** (on page 616)), and comments. There are two kinds of rules: *inference rules* and *target rules*. The *make* utility shall contain a set of built-in inference rules. If the **–r** option is present, the built-in rules shall not be used and the suffix list shall be cleared. Additional rules of both types can be specified in a makefile. If a rule is defined more than once, the value of the rule shall be that of the last one specified. Macros can also be defined more than once, and the value of the macro is specified in **Macros** (on page 616). Comments start with a number sign ('#') and continue until an unescaped <newline> is reached.

By default, the following files shall be tried in sequence: ./makefile and ./Makefile. If neither ./makefile or ./Makefile are found, other implementation-defined files may also be tried. On XSI-conformant systems, the additional files ./s.makefile, SCCS/s.makefile, ./s.Makefile, and SCCS/s.Makefile shall also be tried.

The **-f** option shall direct *make* to ignore any of these default files and use the specified argument as a makefile instead. If the '-' argument is specified, standard input shall be used.

The term *makefile* is used to refer to any rules provided by the user, whether in ./makefile or its variants, or specified by the –f option.

The rules in makefiles shall consist of the following types of lines: target rules, including special targets (see **Target Rules** (on page 614)), inference rules (see **Inference Rules** (on page 617)), macro definitions (see **Macros** (on page 616)), empty lines, and comments.

When an escaped <newline> (one preceded by a backslash) is found anywhere in the makefile except in a command line, it shall be replaced, along with any leading white space on the following line, with a single <space>. When an escaped <newline> is found in a command line

in a makefile, the command line shall contain the backslash, the <newline>, and the next line, except that the first character of the next line shall not be included if it is a <tab>.

Makefile Execution

 Makefile command lines shall be processed one at a time by writing the makefile command line to the standard output (unless one of the conditions listed under '@' suppresses the writing) and executing the command(s) in the line. A <tab> may precede the command to standard output. Command execution shall be as if the makefile command line were the argument to the *system()* function. The environment for the command being executed shall contain all of the variables in the environment of *make*.

By default, when *make* receives a non-zero status from the execution of a command, it shall terminate with an error message to standard error.

Makefile command lines can have one or more of the following prefixes: a hyphen ('-'), an at sign ('@'), or a plus sign ('+'). These shall modify the way in which *make* processes the command. When a command is written to standard output, the prefix shall not be included in the output.

- If the command prefix contains a hyphen, or the -i option is present, or the special target
 .IGNORE has either the current target as a prerequisite or has no prerequisites, any error found while executing the command shall be ignored.
- @ If the command prefix contains an at sign and the *make* utility command line -n option is not specified, or the -s option is present, or the special target .SILENT has either the current target as a prerequisite or has no prerequisites, the command shall not be written to standard output before it is executed.
- + If the command prefix contains a plus sign, this indicates a makefile command line that shall be executed even if $-\mathbf{n}$, $-\mathbf{q}$, or $-\mathbf{t}$ is specified.

Target Rules

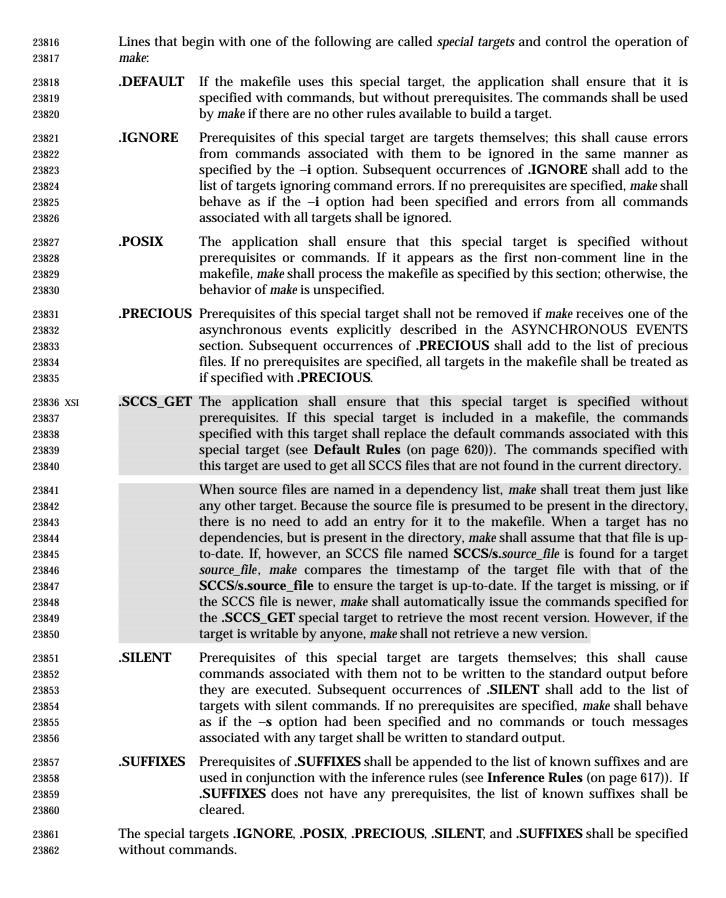
Target rules are formatted as follows:

```
target [target...]: [prerequisite...][;command]
[<tab>command
ctab>command
```

Target entries are specified by a <black>-separated, non-null list of targets, then a colon, then a
 <black>-separated, possibly empty list of prerequisites. Text following a semicolon, if any, and all following lines that begin with a <tab>, are makefile command lines to be executed to update the target. The first non-empty line that does not begin with a <tab> or '#' shall begin a new entry. An empty or blank line, or a line beginning with '#', may begin a new entry.

Applications shall select target names from the set of characters consisting solely of periods, underscores, digits, and alphabetics from the portable character set (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.1, Portable Character Set). Implementations may allow other characters in target names as extensions. The interpretation of targets containing the characters '%' and '"' is implementation-defined.

A target that has prerequisites, but does not have any commands, can be used to add to the prerequisite list for that target. Only one target rule for any given target can contain commands.



Targets with names consisting of a leading period followed by the uppercase letters "POSIX" and then any other characters are reserved for future standardization. Targets with names consisting of a leading period followed by one or more uppercase letters are reserved for implementation extensions.

Macros

 Macro definitions are in the form:

string1 = [string2]

Applications shall select macro names from the set of characters consisting solely of periods, underscores, digits, and alphabetics from the portable character set (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.1, Portable Character Set). A macro name shall not contain an equals sign. Implementations may allow other characters in macro names as extensions.

Macros can appear anywhere in the makefile. Macro expansions using the forms \$(string1) or \${string1} shall be replaced by string2, as follows:

- Macros in target lines shall be evaluated when the target line is read.
- Macros in makefile command lines shall be evaluated when the command is executed.
- Macros in the string before the equals sign in a macro definition shall be evaluated when the macro assignment is made.
- Macros after the equals sign in a macro definition shall not be evaluated until the defined macro is used in a rule or command, or before the equals sign in a macro definition.

The parentheses or braces are optional if *string1* is a single character. The macro \$\$ shall be replaced by the single character '\$'. If *string1* in a macro expansion contains a macro expansion, the results are unspecified.

Macro expansions using the forms \$(string1[:subst1=[subst2]]) or \${string1[:subst1=[subst2]]} can be used to replace all occurrences of subst1 with subst2 when the macro substitution is performed. The subst1 to be replaced shall be recognized when it is a suffix at the end of a word in string1 (where a word, in this context, is defined to be a string delimited by the beginning of the line, a <black>, or a <newline>). If string1 in a macro expansion contains a macro expansion, the results are unspecified.

Macro expansions in *string1* of macro definition lines shall be evaluated when read. Macro expansions in *string2* of macro definition lines shall be performed when the macro identified by *string1* is expanded in a rule or command.

Macro definitions shall be taken from the following sources, in the following logical order, before the makefile(s) are read.

- 1. Macros specified on the *make* utility command line, in the order specified on the command line. It is unspecified whether the internal macros defined in **Internal Macros** (on page 619) are accepted from this source.
- 2. Macros defined by the *MAKEFLAGS* environment variable, in the order specified in the environment variable. It is unspecified whether the internal macros defined in **Internal Macros** (on page 619) are accepted from this source.

3. The contents of the environment, excluding the *MAKEFLAGS* and *SHELL* variables and including the variables with null values.

4. Macros defined in the inference rules built into *make*.

Macro definitions from these sources shall not override macro definitions from a lowernumbered source. Macro definitions from a single source (for example, the *make* utility command line, the *MAKEFLAGS* environment variable, or the other environment variables) shall override previous macro definitions from the same source.

Macros defined in the makefile(s) shall override macro definitions that occur before them in the makefile(s) and macro definitions from source 4. If the –e option is not specified, macros defined in the makefile(s) shall override macro definitions from source 3. Macros defined in the makefile(s) shall not override macro definitions from source 1 or source 2.

Before the makefile(s) are read, all of the *make* utility command line options (except **-f** and **-p**) and *make* utility command line macro definitions (except any for the *MAKEFLAGS* macro), not already included in the *MAKEFLAGS* macro, shall be added to the *MAKEFLAGS* macro, quoted in an implementation-defined manner such that when *MAKEFLAGS* is read by another instance of the *make* command, the original macro's value is recovered. Other implementation-defined options and macros may also be added to the *MAKEFLAGS* macro. If this modifies the value of the *MAKEFLAGS* macro, or, if the *MAKEFLAGS* macro is modified at any subsequent time, the *MAKEFLAGS* environment variable shall be modified to match the new value of the *MAKEFLAGS* macro. The result of setting *MAKEFLAGS* in the Makefile is unspecified.

Before the makefile(s) are read, all of the *make* utility command line macro definitions (except the *MAKEFLAGS* macro or the *SHELL* macro) shall be added to the environment of *make*. Other implementation-defined variables may also be added to the environment of *make*.

The **SHELL** macro shall be treated specially. It shall be provided by *make* and set to the pathname of the shell command language interpreter (see *sh*). The *SHELL* environment variable shall not affect the value of the **SHELL** macro. If **SHELL** is defined in the makefile or is specified on the command line, it shall replace the original value of the **SHELL** macro, but shall not affect the *SHELL* environment variable. Other effects of defining **SHELL** in the makefile or on the command line are implementation-defined.

Inference Rules

Inference rules are formatted as follows:

```
23937 target:
23938 <tab>command
23939 [<tab>command]
23940 ...
```

line that does not begin with <tab> or #

The application shall ensure that the *target* portion is a valid target name (see **Target Rules** (on page 614)) of the form **.s2** or **.s1.s2** (where **.s1** and **.s2** are suffixes that have been given as prerequisites of the **.SUFFIXES** special target and *s1* and *s2* do not contain any slashes or periods.) If there is only one period in the target, it is a single-suffix inference rule. Targets with two periods are double-suffix inference rules. Inference rules can have only one target before the colon.

23947 colon

The application shall ensure that the makefile does not specify prerequisites for inference rules; no characters other than white space shall follow the colon in the first line, except when creating the *empty rule*, described below. Prerequisites are inferred, as described below.

Inference rules can be redefined. A target that matches an existing inference rule shall overwrite the old inference rule. An empty rule can be created with a command consisting of simply a semicolon (that is, the rule still exists and is found during inference rule search, but since it is empty, execution has no effect). The empty rule can also be formatted as follows:

23955 rule:

23980 XSI

where zero or more

| 23956 | where zero or more

| blank > s separate the colon and semicolon.

The *make* utility uses the suffixes of targets and their prerequisites to infer how a target can be made up-to-date. A list of inference rules defines the commands to be executed. By default, *make* contains a built-in set of inference rules. Additional rules can be specified in the makefile.

The special target .SUFFIXES contains as its prerequisites a list of suffixes that shall be used by the inference rules. The order in which the suffixes are specified defines the order in which the inference rules for the suffixes are used. New suffixes shall be appended to the current list by specifying a .SUFFIXES special target in the makefile. A .SUFFIXES target with no prerequisites shall clear the list of suffixes. An empty .SUFFIXES target followed by a new .SUFFIXES list is required to change the order of the suffixes.

Normally, the user would provide an inference rule for each suffix. The inference rule to update a target with a suffix .s1 from a prerequisite with a suffix .s2 is specified as a target .s2.s1. The internal macros provide the means to specify general inference rules (see Internal Macros (on page 619)).

When no target rule is found to update a target, the inference rules shall be checked. The suffix of the target (.s1) to be built is compared to the list of suffixes specified by the .SUFFIXES special targets. If the .s1 suffix is found in .SUFFIXES, the inference rules shall be searched in the order defined for the first .s2.s1 rule whose prerequisite file (\$*.s2) exists. If the target is out-of-date with respect to this prerequisite, the commands for that inference rule shall be executed.

If the target to be built does not contain a suffix and there is no rule for the target, the single suffix inference rules shall be checked. The single-suffix inference rules define how to build a target if a file is found with a name that matches the target name with one of the single suffixes appended. A rule with one suffix .s2 is the definition of how to build *target* from target.s2. The other suffix (.s1) is treated as null.

A tilde ($'^{\ \prime}$) in the above rules refers to an SCCS file in the current directory. Thus, the rule **.c~.o** would transform an SCCS C-language source file into an object file (**.o**). Because the **s.** of the SCCS files is a prefix, it is incompatible with *make*'s suffix point of view. Hence, the $'^{\ \prime}$ ' is a way of changing any file reference into an SCCS file reference.

Libraries

If a target or prerequisite contains parentheses, it shall be treated as a member of an archive library. For the *lib(member.o)* expression *lib* refers to the name of the archive library and *member.o* to the member name. The application shall ensure that the member is an object file with the .o suffix. The modification time of the expression is the modification time for the member as kept in the archive library; see *ar*. The .a suffix shall refer to an archive library. The .s2.a rule shall be used to update a member in the library from a file with a suffix .s2.

23991 Internal Macros

The *make* utility shall maintain five internal macros that can be used in target and inference rules. In order to clearly define the meaning of these macros, some clarification of the terms *target rule*, *inference rule*, *target*, and *prerequisite* is necessary.

Target rules are specified by the user in a makefile for a particular target. Inference rules are user-specified or *make*-specified rules for a particular class of target name. Explicit prerequisites are those prerequisites specified in a makefile on target lines. Implicit prerequisites are those prerequisites that are generated when inference rules are used. Inference rules are applied to implicit prerequisites or to explicit prerequisites that do not have target rules defined for them in the makefile. Target rules are applied to targets specified in the makefile.

Before any target in the makefile is updated, each of its prerequisites (both explicit and implicit) shall be updated. This shall be accomplished by recursively processing each prerequisite. Upon recursion, each prerequisite shall become a target itself. Its prerequisites in turn shall be processed recursively until a target is found that has no prerequisites, at which point the recursion stops. The recursion shall then back up, updating each target as it goes.

In the definitions that follow, the word *target* refers to one of:

- A target specified in the makefile
- An explicit prerequisite specified in the makefile that becomes the target when *make* processes it during recursion
- An implicit prerequisite that becomes a target when *make* processes it during recursion

In the definitions that follow, the word *prerequisite* refers to one of the following:

- An explicit prerequisite specified in the makefile for a particular target
- An implicit prerequisite generated as a result of locating an appropriate inference rule and corresponding file that matches the suffix of the target

The five internal macros are:

\$@ The \$@ shall evaluate to the full target name of the current target, or the archive filename part of a library archive target. It shall be evaluated for both target and inference rules.

For example, in the .c.a inference rule, \$@ represents the out-of-date .a file to be built. Similarly, in a makefile target rule to build lib.a from file.c, \$@ represents the out-of-date lib.a.

The \$% macro shall be evaluated only when the current target is an archive library member of the form *libname*(*member.o*). In these cases, \$@ shall evaluate to *libname* and \$% shall evaluate to *member.o*. The \$% macro shall be evaluated for both target and inference rules.

For example, in a makefile target rule to build **lib.a**(**file.o**), \$% represents **file.o**, as opposed to \$@, which represents **lib.a**.

\$? The \$? macro shall evaluate to the list of prerequisites that are newer than the current target. It shall be evaluated for both target and inference rules.

For example, in a makefile target rule to build *prog* from **file1.0**, **file2.0**, and **file3.0**, and where *prog* is not out-of-date with respect to **file1.0**, but is out-of-date with respect to **file2.0** and **file3.0**, \$? represents **file2.0** and **file3.0**.

\$ <	In an inference rule, the \$< macro shall evaluate to the filename whose existence allowed the inference rule to be chosen for the target. In the .DEFAULT rule, the \$< macro shall evaluate to the current target name. The meaning of the \$< macro shall be otherwise unspecified.	
	For example, in the .c.a inference rule, \$< represents the prerequisite .c file.	
\$*	The \$* macro shall evaluate to the current target name with its suffix deleted. It shall be evaluated at least for inference rules.	
	For example, in the $.c.a$ inference rule, $$^*.o$ represents the out-of-date $.o$ file that corresponds to the prerequisite $.c$ file.	
to any of for 'F' director filenam	the internal macros has an alternative form. When an uppercase 'D' or 'F' is appended of the macros, the meaning shall be changed to the <i>directory part</i> for 'D' and <i>filename part</i> . The directory part is the path prefix of the file without a trailing slash; for the current ry, the directory part is '.'. When the \$? macro contains more than one prerequisite e, the (PD) and (PF) (or PP) and PP) macros expand to a list of directory name parts name parts respectively.	
For the	For the target <i>lib</i> (<i>member.o</i>) and the s2.a rule, the internal macros shall be defined as:	
\$<	member.s2	
\$*	member	
\$@	lib	
\$?	member.s2	
\$%	member.o	
Default	Rules	
The default rules for <i>make</i> shall achieve results that are the same as if the following were used. Implementations that do not support the C-Language Development Utilities option may omit CC, CFLAGS, YACC, YFLAGS, LEX, LFLAGS, LDFLAGS, and the .c, .y, and .l inference rules. Implementations that do not support FORTRAN may omit FC, FFLAGS, and the .f inference rules. Implementations may provide additional macros and rules.		
SPECIA	AL TARGETS	
.SCCS_	_GET: sccs \$(SCCSFLAGS) get \$(SCCSGETFLAGS) \$@	
.SUFFI	IXES: .o .c .y .l .a .sh .f .c~ .y~ .l~ .sh~ .f~	
MACRO	OS .	
AR=ar ARFLAC YACC=y YFLAGS LEX=1e LFLAGS LDFLAC CC=c99 CFLAGS	GS=-rv vacc S= ex G= GS= O	
	S* Each of to any of for 'F' director filenam and file For the S< S* S@ S? S% Default The def Implem CC, CF, Implem rules. Ir SPECIA .SCCSSUFFI MACRO MAKE=n	

```
24076
            FFLAGS=-0 1
24077 XSI
            GET=get
24078
            GFLAGS=
24079
            SCCSFLAGS=
24080
            SCCSGETFLAGS=-s
24081
            SINGLE SUFFIX RULES
24082
24083
                 $(CC) $(CFLAGS) $(LDFLAGS) -0 $@ $<
24084
            .f:
24085
24086
                 $(FC) $(FFLAGS) $(LDFLAGS) -0 $@ $<
24087
            .sh:
24088
                 cp $< $@
24089
                 chmod a+x $@
24090 XSI
24091
                 (GET) (GFLAGS) -p << > *.c
24092
                 $(CC) $(CFLAGS) $(LDFLAGS) -0 $@ $*.c
24093
            .f~:
                 $(GET) $(GFLAGS) -p $< > $*.f
24094
                 $(FC) $(FFLAGS) $(LDFLAGS) -o $@ $*.f
24095
24096
            .sh~:
24097
                 $(GET) $(GFLAGS) -p $< > $*.sh
24098
                 cp $*.sh $@
24099
                 chmod a+x $@
24100
            DOUBLE SUFFIX RULES
24101
24102
            .c.o:
24103
                 $(CC) $(CFLAGS) -c $<
            .f.o:
24104
24105
                 $(FC) $(FFLAGS) -c $<
24106
            .y.o:
24107
                 $(YACC) $(YFLAGS) $<
24108
                 $(CC) $(CFLAGS) -c y.tab.c
                 rm -f y.tab.c
24109
24110
                 mv y.tab.o $@
24111
            .1.0:
                 $(LEX) $(LFLAGS) $<
24112
                 $(CC) $(CFLAGS) -c lex.yy.c
24113
                 rm -f lex.yy.c
24114
24115
                 mv lex.yy.o $@
24116
            .y.c:
24117
                 $(YACC) $(YFLAGS) $<
24118
                 mv y.tab.c $@
24119
            .1.c:
                 $(LEX) $(LFLAGS) $<
24120
```

```
24121
                  mv lex.yy.c $@
             .c~.o:
24122 XSI
24123
                  \$(GET) \$(GFLAGS) -p \$< > \$*.c
                  $(CC) $(CFLAGS) -c $*.c
24124
24125
             .f~.o:
24126
                  (GET) (GFLAGS) -p << > *.f
                  $(FC) $(FFLAGS) -c $*.f
24127
24128
             .y~.o:
24129
                  \$(GET) \$(GFLAGS) -p \$< > \$*.y
                  $(YACC) $(YFLAGS) $*.y
24130
24131
                  $(CC) $(CFLAGS) -c y.tab.c
24132
                  rm -f y.tab.c
24133
                  mv y.tab.o $@
24134
             .1~.0:
24135
                  $(GET) $(GFLAGS) -p $< > $*.1
24136
                  $(LEX) $(LFLAGS) $*.1
                  $(CC) $(CFLAGS) -c lex.yy.c
24137
                  rm -f lex.yy.c
24138
                  mv lex.yy.o $@
24139
             .y~.c:
24140
24141
                  \$(GET) \$(GFLAGS) -p \$< > \$*.y
                  $(YACC) $(YFLAGS) $*.y
24142
24143
                  mv y.tab.c $@
             .1~.c:
24144
                  $(GET) $(GFLAGS) -p $< > $*.1
24145
                  $(LEX) $(LFLAGS) $*.1
24146
24147
                  mv lex.yy.c $@
24148
24149
             .c.a:
24150
                  $(CC) -c $(CFLAGS) $<
24151
                  $(AR) $(ARFLAGS) $@ $*.0
                  rm -f $*.o
24152
             .f.a:
24153
24154
                  $(FC) -c $(FFLAGS) $<
24155
                  $(AR) $(ARFLAGS) $@ $*.0
24156
                  rm -f $*.o
24157 EXIT STATUS
             When the -\mathbf{q} option is specified, the make utility shall exit with one of the following values:
24158
              0 Successful completion.
24159
24160
                The target was not up-to-date.
24161
             >1 An error occurred.
             When the -\mathbf{q} option is not specified, the make utility shall exit with one of the following values:
24162
24163
                 Successful completion.
             >0 An error occurred.
24164
```

24165 CONSEQUENCES OF ERRORS

24166 Default.

24175 24176

24182

2418324184

24186 24187

24188

24189 24190

2419124192

24193

24194

24195

24201

24202 24203

24204

24205 24206

24167 APPLICATION USAGE

If there is a source file (such as ./source.c) and there are two SCCS files corresponding to it (./s.source.c and ./SCCS/s.source.c), on XSI-conformant systems *make* uses the SCCS file in the current directory. However, users are advised to use the underlying SCCS utilities (*admin*, *delta*, *get*, and so on) or the *sccs* utility for all source files in a given directory. If both forms are used for a given source file, future developers are very likely to be confused.

It is incumbent upon portable makefiles to specify the **.POSIX** special target in order to guarantee that they are not affected by local extensions.

The **-k** and **-S** options are both present so that the relationship between the command line, the *MAKEFLAGS* variable, and the makefile can be controlled precisely. If the **k** flag is passed in *MAKEFLAGS* and a command is of the form:

24177 *MAKEFLAGS* and a command is o

```
24178 $ (MAKE) -S foo
```

24179 then the default behavior is restored for the child *make*.

When the **-n** option is specified, it is always added to *MAKEFLAGS*. This allows a recursive make **-n** target to be used to see all of the action that would be taken to update target.

Because of widespread historical practice, interpreting a '#' number sign inside a variable as the start of a comment has the unfortunate side effect of making it impossible to place a number sign in a variable, thus forbidding something like:

```
24185 CFLAGS = "-D COMMENT CHAR='#'"
```

Many historical *make* utilities stop chaining together inference rules when an intermediate target is nonexistent. For example, it might be possible for a *make* to determine that both .y.c and .c.o could be used to convert a .y to a .o. Instead, in this case, *make* requires the use of a .y.o rule.

The best way to provide portable makefiles is to include all of the rules needed in the makefile itself. The rules provided use only features provided by other parts of this volume of IEEE Std 1003.1-2001. The default rules include rules for optional commands in this volume of IEEE Std 1003.1-2001. Only rules pertaining to commands that are provided are needed in an implementation's default set.

Macros used within other macros are evaluated when the new macro is used rather than when the new macro is defined. Therefore:

```
24196 MACRO = value1

24197 NEW = $ (MACRO)

24198 MACRO = value2

24199 target:

24200 echo $ (NEW)
```

would produce *value2* and not *value1* since **NEW** was not expanded until it was needed in the *echo* command line.

Some historical applications have been known to intermix *target_name* and *macro=name* operands on the command line, expecting that all of the macros are processed before any of the targets are dealt with. Conforming applications do not do this, although some backwards-compatibility support may be included in some implementations.

The following characters in filenames may give trouble: '=', ':', '':', and '@'. For inference rules, the description of < and <? seem similar. However, an example shows the

```
24209
              minor difference. In a makefile containing:
24210
              foo.o: foo.h
24211
              if foo.h is newer than foo.o, yet foo.c is older than foo.o, the built-in rule to make foo.o from
24212
              foo.c is used, with $< equal to foo.c and $? equal to foo.h. If foo.c is also newer than foo.o, $< is
24213
              equal to foo.c and $? is equal to foo.h foo.c.
24214 EXAMPLES
                  The following command:
24215
24216
                    makes the first target found in the makefile.
24217
24218
                2. The following command:
24219
                   make junk
24220
                   makes the target junk.
24221
                  The following makefile says that pgm depends on two files, a.o and b.o, and that they in
24222
                    turn depend on their corresponding source files (a.c and b.c), and a common file incl.h:
24223
                   pgm: a.o b.o
                         c99 a.o b.o -o pgm
24224
24225
                    a.o: incl.h a.c
24226
                         c99 -c a.c
                   b.o: incl.h b.c
24227
                         c99 -c b.c
24228
24229
                4. An example for making optimized .o files from .c files is:
24230
24231
                         c99 -c -0 $*.c
24232
                   or:
24233
                    .c.o:
24234
                         c99 -c -O $<
24235
                5. The most common use of the archive interface follows. Here, it is assumed that the source
24236
                   files are all C-language source:
24237
                    lib: lib(file1.o) lib(file2.o) lib(file3.o)
24238
                         @echo lib is now up-to-date
24239
                   The .c.a rule is used to make file1.o, file2.o, and file3.o and insert them into lib.
24240
                   The treatment of escaped <newline>s throughout the makefile is historical practice. For
24241
                   example, the inference rule:
24242
                    .c.o\
24243
                   works, and the macro:
24244
24245
                        bar baz\
                    f=
24246
                         biz
24247
                   a :
24248
                         echo == f==
```

```
24249
                  echoes "==bar baz biz==".
24250
                  If $? were:
24251
                  /usr/include/stdio.h /usr/include/unistd.h foo.h
                  then $(?D) would be:
24252
                  /usr/include /usr/include .
24253
24254
                  and $(?F) would be:
                  stdio.h unistd.h foo.h
24255
24256
              6. The contents of the built-in rules can be viewed by running:
24257
                  make -p -f /dev/null 2>/dev/null
```

24258 RATIONALE

24259 24260

24261 24262

24263

24264

2426524266

24267 24268

24269 24270

2427124272

24273

24276

24278

24279

24280

2428124282

24283

24284

24285

24286 24287

24288

24289

The *make* utility described in this volume of IEEE Std 1003.1-2001 is intended to provide the means for changing portable source code into executables that can be run on an IEEE Std 1003.1-2001-conforming system. It reflects the most common features present in System V and BSD *makes*.

Historically, the *make* utility has been an especially fertile ground for vendor and research organization-specific syntax modifications and extensions. Examples include:

- Syntax supporting parallel execution (such as from various multi-processor vendors, GNU, and others)
- Additional "operators" separating targets and their prerequisites (System V, BSD, and others)
- Specifying that command lines containing the strings "\${MAKE}" and "\$(MAKE)" are executed when the -n option is specified (GNU and System V)
- Modifications of the meaning of internal macros when referencing libraries (BSD and others)
- Using a single instance of the shell for all of the command lines of the target (BSD and others)
- Allowing spaces as well as tabs to delimit command lines (BSD)
- Adding C preprocessor-style "include" and "ifdef" constructs (System V, GNU, BSD, and others)
 - Remote execution of command lines (Sprite and others)
- Specifying additional special targets (BSD, System V, and most others)

Additionally, many vendors and research organizations have rethought the basic concepts of *make*, creating vastly extended, as well as completely new, syntaxes. Each of these versions of *make* fulfills the needs of a different community of users; it is unreasonable for this volume of IEEE Std 1003.1-2001 to require behavior that would be incompatible (and probably inferior) to historical practice for such a community.

In similar circumstances, when the industry has enough sufficiently incompatible formats as to make them irreconcilable, this volume of IEEE Std 1003.1-2001 has followed one or both of two courses of action. Commands have been renamed (*cksum*, *echo*, and *pax*) and/or command line options have been provided to select the desired behavior (*grep*, *od*, and *pax*).

Because the syntax specified for the *make* utility is, by and large, a subset of the syntaxes accepted by almost all versions of *make*, it was decided that it would be counter-productive to change the name. And since the makefile itself is a basic unit of portability, it would not be

completely effective to reserve a new option letter, such as *make* –**P**, to achieve the portable behavior. Therefore, the special target .**POSIX** was added to the makefile, allowing users to specify "standard" behavior. This special target does not preclude extensions in the *make* utility, nor does it preclude such extensions being used by the makefile specifying the target; it does, however, preclude any extensions from being applied that could alter the behavior of previously valid syntax; such extensions must be controlled via command line options or new special targets. It is incumbent upon portable makefiles to specify the .**POSIX** special target in order to guarantee that they are not affected by local extensions.

The portable version of *make* described in this reference page is not intended to be the state-of-the-art software generation tool and, as such, some newer and more leading-edge features have not been included. An attempt has been made to describe the portable makefile in a manner that does not preclude such extensions as long as they do not disturb the portable behavior described here.

When the $-\mathbf{n}$ option is specified, it is always added to *MAKEFLAGS*. This allows a recursive $make - \mathbf{n}$ target to be used to see all of the action that would be taken to update target.

The definition of *MAKEFLAGS* allows both the System V letter string and the BSD command line formats. The two formats are sufficiently different to allow implementations to support both without ambiguity.

Early proposals stated that an "unquoted" number sign was treated as the start of a comment. The *make* utility does not pay any attention to quotes. A number sign starts a comment regardless of its surroundings.

The text about "other implementation-defined pathnames may also be tried" in addition to ./makefile and ./Makefile is to allow such extensions as SCCS/s.Makefile and other variations. It was made an implementation-defined requirement (as opposed to unspecified behavior) to highlight surprising implementations that might select something unexpected like /etc/Makefile. XSI-conformant systems also try ./s.makefile, SCCS/s.makefile, ./s.Makefile, and SCCS/s.Makefile.

Early proposals contained the macro **NPROC** as a means of specifying that *make* should use *n* processes to do the work required. While this feature is a valuable extension for many systems, it is not common usage and could require other non-trivial extensions to makefile syntax. This extension is not required by this volume of IEEE Std 1003.1-2001, but could be provided as a compatible extension. The macro **PARALLEL** is used by some historical systems with essentially the same meaning (but without using a name that is a common system limit value). It is suggested that implementors recognize the existing use of **NPROC** and/or **PARALLEL** as extensions to *make*.

The default rules are based on System V. The default **CC**= value is *c99* instead of *cc* because this volume of IEEE Std 1003.1-2001 does not standardize the utility named *cc*. Thus, every conforming application would be required to define **CC**=*c99* to expect to run. There is no advantage conferred by the hope that the makefile might hit the "preferred" compiler because this cannot be guaranteed to work. Also, since the portable makescript can only use the *c99* options, no advantage is conferred in terms of what the script can do. It is a quality-of-implementation issue as to whether *c99* is as valuable as *cc*.

The **-d** option to *make* is frequently used to produce debugging information, but is too implementation-defined to add to this volume of IEEE Std 1003.1-2001.

The **–p** option is not passed in *MAKEFLAGS* on most historical implementations and to change this would cause many implementations to break without sufficiently increased portability.

Commands that begin with a plus sign ('+') are executed even if the $-\mathbf{n}$ option is present. Based on the GNU version of *make*, the behavior of $-\mathbf{n}$ when the plus-sign prefix is encountered has been extended to apply to $-\mathbf{q}$ and $-\mathbf{t}$ as well. However, the System V convention of forcing command execution with $-\mathbf{n}$ when the command line of a target contains either of the strings "\$(MAKE)" or "\${MAKE}" has not been adopted. This functionality appeared in early proposals, but the danger of this approach was pointed out with the following example of a portion of a makefile:

```
subdir:
   cd subdir; rm all the files; $(MAKE)
```

The loss of the System V behavior in this case is well-balanced by the safety afforded to other makefiles that were not aware of this situation. In any event, the command line plus-sign prefix can provide the desired functionality.

The double colon in the target rule format is supported in BSD systems to allow more than one target line containing the same target name to have commands associated with it. Since this is not functionality described in the SVID or XPG3 it has been allowed as an extension, but not mandated.

The default rules are provided with text specifying that the built-in rules shall be the same as if the listed set were used. The intent is that implementations should be able to use the rules without change, but will be allowed to alter them in ways that do not affect the primary behavior.

The best way to provide portable makefiles is to include all of the rules needed in the makefile itself. The rules provided use only features provided by other portions of this volume of IEEE Std 1003.1-2001. The default rules include rules for optional commands in this volume of IEEE Std 1003.1-2001. Only rules pertaining to commands that are provided are needed in the default set of an implementation.

One point of discussion was whether to drop the default rules list from this volume of IEEE Std 1003.1-2001. They provide convenience, but do not enhance portability of applications. The prime benefit is in portability of users who wish to type *make command* and have the command build from a **command.c** file.

The historical *MAKESHELL* feature was omitted. In some implementations it is used to let a user override the shell to be used to run *make* commands. This was confusing; for a portable *make*, the shell should be chosen by the makefile writer or specified on the *make* command line and not by a user running *make*.

The *make* utilities in most historical implementations process the prerequisites of a target in left-to-right order, and the makefile format requires this. It supports the standard idiom used in many makefiles that produce *yacc* programs; for example:

```
foo: y.tab.o lex.o main.o $(CC) $(CFLAGS) -o $@ t.tab.o lex.o main.o
```

In this example, if *make* chose any arbitrary order, the **lex.o** might not be made with the correct **y.tab.h**. Although there may be better ways to express this relationship, it is widely used historically. Implementations that desire to update prerequisites in parallel should require an explicit extension to *make* or the makefile format to accomplish it, as described previously.

The algorithm for determining a new entry for target rules is partially unspecified. Some historical *make*s allow blank, empty, or comment lines within the collection of commands marked by leading <tab>s. A conforming makefile must ensure that each command starts with a <tab>, but implementations are free to ignore blank, empty, and comment lines without triggering the start of a new entry.

The ASYNCHRONOUS EVENTS section includes having SIGTERM and SIGHUP, along with the more traditional SIGINT and SIGQUIT, remove the current target unless directed not to do so. SIGTERM and SIGHUP were added to parallel other utilities that have historically cleaned up their work as a result of these signals. When *make* receives any signal other than SIGQUIT, it is required to resend itself the signal it received so that it exits with a status that reflects the signal. The results from SIGQUIT are partially unspecified because, on systems that create **core** files upon receipt of SIGQUIT, the **core** from *make* would conflict with a **core** file from the command that was running when the SIGQUIT arrived. The main concern was to prevent damaged files from appearing up-to-date when *make* is rerun.

The .PRECIOUS special target was extended to affect all targets globally (by specifying no prerequisites). The .IGNORE and .SILENT special targets were extended to allow prerequisites; it was judged to be more useful in some cases to be able to turn off errors or echoing for a list of targets than for the entire makefile. These extensions to *make* in System V were made to match historical practice from the BSD *make*.

Macros are not exported to the environment of commands to be run. This was never the case in any historical *make* and would have serious consequences. The environment is the same as the environment to *make* except that *MAKEFLAGS* and macros defined on the *make* command line are added.

Some implementations do not use <code>system()</code> for all command lines, as required by the portable makefile format; as a performance enhancement, they select lines without shell metacharacters for direct execution by <code>execve()</code>. There is no requirement that <code>system()</code> be used specifically, but merely that the same results be achieved. The metacharacters typically used to bypass the direct <code>execve()</code> execution have been any of:

= | ^ () ; & < > * ? [] : \$ ' ' " \ \n

The default in some advanced versions of *make* is to group all the command lines for a target and execute them using a single shell invocation; the System V method is to pass each line individually to a separate shell. The single-shell method has the advantages in performance and the lack of a requirement for many continued lines. However, converting to this newer method has caused portability problems with many historical makefiles, so the behavior with the POSIX makefile is specified to be the same as that of System V. It is suggested that the special target .ONESHELL be used as an implementation extension to achieve the single-shell grouping for a target or group of targets.

Novice users of *make* have had difficulty with the historical need to start commands with a <tab>. Since it is often difficult to discern differences between <tab>s and <space>s on terminals or printed listings, confusing bugs can arise. In early proposals, an attempt was made to correct this problem by allowing leading <blank>s instead of <tab>s. However, implementors reported many makefiles that failed in subtle ways following this change, and it is difficult to implement a *make* that unambiguously can differentiate between macro and command lines. There is extensive historical practice of allowing leading spaces before macro definitions. Forcing macro lines into column 1 would be a significant backwards-compatibility problem for some makefiles. Therefore, historical practice was restored.

The System V INCLUDE feature was considered, but not included. This would treat a line that began in the first column and contained INCLUDE < filename > as an indication to read < filename > at that point in the makefile. This is difficult to use in a portable way, and it raises concerns about nesting levels and diagnostics. System V, BSD, GNU, and others have used different methods for including files.

The System V dynamic dependency feature was not included. It would support:

24430	cat: \$\$@.c
24431	that would expand to;
24432	cat: cat.c
24433 24434 24435	This feature exists only in the new version of System V <i>make</i> and, while useful, is not in wide usage. This means that macros are expanded twice for prerequisites: once at makefile parse time and once at target update time.
24436 24437 24438 24439 24440 24441	Consideration was given to adding metarules to the POSIX <i>make</i> . This would make %.o: %.c the same as .c.o:. This is quite useful and available from some vendors, but it would cause too many changes to this <i>make</i> to support. It would have introduced rule chaining and new substitution rules. However, the rules for target names have been set to reserve the '%' and '"' characters. These are traditionally used to implement metarules and quoting of target names, respectively. Implementors are strongly encouraged to use these characters only for these purposes.
24442 24443 24444	A request was made to extend the suffix delimiter character from a period to any character. The metarules feature in newer <i>make</i> s solves this problem in a more general way. This volume of IEEE Std 1003.1-2001 is staying with the more conservative historical definition.
24445 24446 24447 24448 24449	The standard output format for the $-\mathbf{p}$ option is not described because it is primarily a debugging option and because the format is not generally useful to programs. In historical implementations the output is not suitable for use in generating makefiles. The $-\mathbf{p}$ format has been variable across historical implementations. Therefore, the definition of $-\mathbf{p}$ was only to provide a consistently named option for obtaining <i>make</i> script debugging information.
24450	Some historical implementations have not cleared the suffix list with $-\mathbf{r}$.
24451 24452 24453 24454	Implementations should be aware that some historical applications have intermixed <i>target_name</i> and <i>macro=value</i> operands on the command line, expecting that all of the macros are processed before any of the targets are dealt with. Conforming applications do not do this, but some backwards-compatibility support may be warranted.
24455 24456 24457	Empty inference rules are specified with a semicolon command rather than omitting all commands, as described in an early proposal. The latter case has no traditional meaning and is reserved for implementation extensions, such as in GNU <i>make</i> .
24458 FUTUF 24459	RE DIRECTIONS None.
24460 SEE AI 24461 24462	CSO Chapter 2 (on page 29), ar, c99, get, lex, sccs, sh, yacc, the System Interfaces volume of IEEE Std 1003.1-2001, exec, system()
24463 CHAN 24464	GE HISTORY First released in Issue 2.
24465 Issue 5 24466	The FUTURE DIRECTIONS section is added.
24467 Issue 6	
24468	This utility is marked as part of the Software Development Utilities option.
24469 24470	The Open Group Corrigendum U029/1 is applied, correcting a typographical error in the SPECIAL TARGETS section.
24471 24472 24473	In the ENVIRONMENT VARIABLES section, the <i>PROJECTDIR</i> description is updated from "otherwise, the home directory of a user of that name is examined" to "otherwise, the value of <i>PROJECTDIR</i> is treated as a user name and that user's initial working directory is examined".

24474 24475	It is specified whether the command line is related to the makefile or to the <i>make</i> command, and the macro processing rules are updated to align with the IEEE P1003.2b draft standard.
24476	The normative text is reworded to avoid use of the term "must" for application requirements.
24477	PASC Interpretation 1003.2 #193 is applied.

Utilities man

24478 **NAME**24479 man — display system documentation
24480 **SYNOPSIS**24481 man [-k] name...

24482 DESCRIPTION

The *man* utility shall write information about each of the *name* operands. If *name* is the name of a standard utility, *man* at a minimum shall write a message describing the syntax used by the standard utility, its options, and operands. If more information is available, the *man* utility shall provide it in an implementation-defined manner.

An implementation may provide information for values of *name* other than the standard utilities. Standard utilities that are listed as optional and that are not supported by the implementation either shall cause a brief message indicating that fact to be displayed or shall cause a full display of information as described previously.

24491 OPTIONS

24487

24488 24489

24490

24494

24504

24505

24506

24507

The *man* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following option shall be supported:

24495 -k Interpret *name* operands as keywords to be used in searching a utilities summary database that contains a brief purpose entry for each standard utility and write lines from the summary database that match any of the keywords. The keyword search shall produce results that are the equivalent of the output of the following command:

 24499
 grep -Ei '

 24500
 name

 24501
 name

 24502
 ...

 24503
 ' summary-database

This assumes that the *summary-database* is a text file with a single entry per line; this organization is not required and the example using *grep*—**Ei** is merely illustrative of the type of search intended. The purpose entry to be included in the database shall consist of a terse description of the purpose of the utility.

24508 OPERANDS

24509 The following operand shall be supported:

A keyword or the name of a standard utility. When **-k** is not specified and *name* does not represent one of the standard utilities, the results are unspecified.

24512 **STDIN**

Not used.

24514 INPUT FILES

24515 None.

24516 ENVIRONMENT VARIABLES

24517 The following environment variables shall affect the execution of *man*:

24518 LANG Provide a default value for the internationalization variables that are unset or null.
24519 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
24520 Internationalization Variables for the precedence of internationalization variables
24521 used to determine the values of locale categories.)

man Utilities

24522 24523	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
24524 24525 24526 24527	LC_CTYPE	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 in the summary database). The value of <i>LC_CTYPE</i> need not affect the format of the information written about the <i>name</i> operands.
24528 24529 24530 24531	LC_MESSA(Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.
24532 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
24533 24534 24535 24536 24537 24538	PAGER	Determine an output filtering command for writing the output to a terminal. Any string acceptable as a <i>command_string</i> operand to the sh – c command shall be valid. When standard output is a terminal device, the reference page output shall be piped through the command. If the <i>PAGER</i> variable is null or not set, the command shall be either <i>more</i> or another paginator utility documented in the system documentation.
ACTIVIC	TIPONIOTIC I	ON ADDITION

24539 ASYNCHRONOUS EVENTS

24540 Default.

24541 STDOUT

The man utility shall write text describing the syntax of the utility name, its options and its operands, or, when $-\mathbf{k}$ is specified, lines from the summary database. The format of this text is implementation-defined.

24545 STDERR

24546 The standard error shall be used only for diagnostic messages.

24547 OUTPUT FILES

24548 None.

24549 EXTENDED DESCRIPTION

24550 None.

24551 EXIT STATUS

24552 The following exit values shall be returned:

24553 0 Successful completion.

24554 >0 An error occurred.

24555 CONSEQUENCES OF ERRORS

24556 Default.

24557 APPLICATION USAGE

24558 None.
 24559 **EXAMPLES** 24560 None.

24561 RATIONALE

It is recognized that the *man* utility is only of minimal usefulness as specified. The opinion of the standard developers was strongly divided as to how much or how little information *man* should be required to provide. They considered, however, that the provision of some portable way of accessing documentation would aid user portability. The arguments against a fuller

Utilities man

24566 specification were:

24569

24570

24571

24572

24575 24576

24577

24578

24579

2458024581

2458224583

24584

24585

24586

24587 24588

24589

24590 24591

24592 24593

24594

2459524596

24597 24598

• Large quantities of documentation should not be required on a system that does not have excess disk space.

- The current manual system does not present information in a manner that greatly aids user portability.
- A "better help system" is currently an area in which vendors feel that they can add value to their POSIX implementations.

The –f option was considered, but due to implementation differences, it was not included in this volume of IEEE Std 1003.1-2001.

The description was changed to be more specific about what has to be displayed for a utility. The standard developers considered it insufficient to allow a display of only the synopsis without giving a short description of what each option and operand does.

The "purpose" entry to be included in the database can be similar to the section title (less the numeric prefix) from this volume of IEEE Std 1003.1-2001 for each utility. These titles are similar to those used in historical systems for this purpose.

See *mailx* for rationale concerning the default paginator.

The caveat in the *LC_CTYPE* description was added because it is not a requirement that an implementation provide reference pages for all of its supported locales on each system; changing *LC_CTYPE* does not necessarily translate the reference page into another language. This is equivalent to the current state of *LC_MESSAGES* in IEEE Std 1003.1-2001—locale-specific messages are not yet a requirement.

The historical *MANPATH* variable is not included in POSIX because no attempt is made to specify naming conventions for reference page files, nor even to mandate that they are files at all. On some implementations they could be a true database, a hypertext file, or even fixed strings within the *man* executable. The standard developers considered the portability of reference pages to be outside their scope of work. However, users should be aware that *MANPATH* is implemented on a number of historical systems and that it can be used to tailor the search pattern for reference pages from the various categories (utilities, functions, file formats, and so on) when the system administrator reveals the location and conventions for reference pages on the system.

The keyword search can rely on at least the text of the section titles from these utility descriptions, and the implementation may add more keywords. The term "section titles" refers to the strings such as:

```
24599 man - Display system documentation
24600 ps - Report process status
```

24601 FUTURE DIRECTIONS

24602 None.
 24603 SEE ALSO
 24604 more

24605 CHANGE HISTORY

First released in Issue 4.

man Utilities

24607 **Issue 5**

24608 The FUTURE DIRECTIONS section is added.

Utilities mesg

24609 NAME

24610 mesg — permit or deny messages

24611 SYNOPSIS

24612 UP mesg [y|n]

24613

24614 **DESCRIPTION**

The mesg utility shall control whether other users are allowed to send messages via write, talk, or 24615 other utilities to a terminal device. The terminal device affected shall be determined by searching 24616 for the first terminal in the sequence of devices associated with standard input, standard output, 24617 24618 and standard error, respectively. With no arguments, mesg shall report the current state without changing it. Processes with appropriate privileges may be able to send messages to the terminal 24619 24620 independent of the current state.

24621 OPTIONS

None. 24622

24623 OPERANDS

The following operands shall be supported in the POSIX locale: 24624

24625 Grant permission to other users to send messages to the terminal device. y

Deny permission to other users to send messages to the terminal device. 24626 n

24627 **STDIN**

24628 Not used.

24629 INPUT FILES

24630 None.

24631 ENVIRONMENT VARIABLES

24632 The following environment variables shall affect the execution of *mesg*:

LANG Provide a default value for the internationalization variables that are unset or null. 24633 24634 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables 24635 used to determine the values of locale categories.) 24636

LC ALL If set to a non-empty string value, override the values of all the other 24637

internationalization variables. 24638

24639 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 24640 characters (for example, single-byte as opposed to multi-byte characters in

24641 arguments).

24642 LC_MESSAGES

Determine the locale that should be used to affect the format and contents of 24643 24644

diagnostic messages written (by mesg) to standard error.

NLSPATH 24645 XSI Determine the location of message catalogs for the processing of *LC_MESSAGES*.

24646 ASYNCHRONOUS EVENTS

Default. 24647

24648 STDOUT

24649 If no operand is specified, *mesg* shall display the current terminal state in an unspecified format. mesg Utilities

24650 STDERR

The standard error shall be used only for diagnostic messages.

24652 OUTPUT FILES

24653 None.

24654 EXTENDED DESCRIPTION

24655 None.

24656 EXIT STATUS

24657 The following exit values shall be returned:

24658 0 Receiving messages is allowed.

24659 1 Receiving messages is not allowed.

24660 >1 An error occurred.

24661 CONSEQUENCES OF ERRORS

24662 Default.

24663 APPLICATION USAGE

The mechanism by which the message status of the terminal is changed is unspecified. Therefore, unspecified actions may cause the status of the terminal to change after *mesg* has successfully completed. These actions may include, but are not limited to: another invocation of the *mesg* utility, login procedures; invocation of the *stty* utility, invocation of the *chmod* utility or *chmod*() function, and so on.

24669 EXAMPLES

24670 None.

24671 RATIONALE

24672

24673

2467424675

24676

24679

24680

24681 24682

24683

The terminal changed by *mesg* is that associated with the standard input, output, or error, rather than the controlling terminal for the session. This is because users logged in more than once should be able to change any of their login terminals without having to stop the job running in those sessions. This is not a security problem involving the terminals of other users because appropriate privileges would be required to affect the terminal of another user.

The method of checking each of the first three file descriptors in sequence until a terminal is found was adopted from System V.

The file /dev/tty is not specified for the terminal device because it was thought to be too restrictive. Typical environment changes for the *n* operand are that write permissions are removed for *others* and *group* from the appropriate device. It was decided to leave the actual description of what is done as unspecified because of potential differences between implementations.

The format for standard output is unspecified because of differences between historical implementations. This output is generally not useful to shell scripts (they can use the exit status), so exact parsing of the output is unnecessary.

24687 FUTURE DIRECTIONS

24688 None.

24689 SEE ALSO

24690 talk, write

Utilities mesg

24691 CHANGE HISTORY

First released in Issue 2.

24693 **Issue 6**

24694 This utility is marked as part of the User Portability Utilities option.

mkdir Utilities

```
24695 NAME
24696
              mkdir — make directories
24697 SYNOPSIS
24698
              mkdir [-p] [-m mode] dir...
24699 DESCRIPTION
              The mkdir utility shall create the directories specified by the operands, in the order specified.
24700
24701
              For each dir operand, the mkdir utility shall perform actions equivalent to the mkdir() function
24702
              defined in the System Interfaces volume of IEEE Std 1003.1-2001, called with the following
24703
              arguments:
               1. The dir operand is used as the path argument.
24704
                   The value of the bitwise-inclusive OR of S_IRWXU, S_IRWXG, and S_IRWXO is used as
24705
                   the mode argument. (If the -\mathbf{m} option is specified, the mode option-argument overrides this
24706
                   default.)
24707
24708 OPTIONS
              The mkdir utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
24709
              12.2, Utility Syntax Guidelines.
24710
              The following options shall be supported:
24711
              -m mode
                           Set the file permission bits of the newly-created directory to the specified mode
24712
                           value. The mode option-argument shall be the same as the mode operand defined
24713
                           for the chmod utility. In the symbolic_mode strings, the op characters '+' and '-'
24714
                           shall be interpreted relative to an assumed initial mode of a=rwx; '+' shall add
24715
                           permissions to the default mode, '-' shall delete permissions from the default
24716
                           mode.
24717
24718
                           Create any missing intermediate pathname components.
              -p
                           For each dir operand that does not name an existing directory, effects equivalent to
24719
24720
                           those caused by the following command shall occur:
24721
                           mkdir -p -m $(umask -S),u+wx $(dirname dir) &&
24799
                           mkdir [-m mode] dir
                           where the -m mode option represents that option supplied to the original
24723
24724
                           invocation of mkdir, if any.
                           Each dir operand that names an existing directory shall be ignored without error.
24725
24726 OPERANDS
              The following operand shall be supported:
              dir
24728
                           A pathname of a directory to be created.
24729 STDIN
              Not used.
24730
24731 INPUT FILES
              None.
24732
24733 ENVIRONMENT VARIABLES
              The following environment variables shall affect the execution of mkdir:
24734
              LANG
                           Provide a default value for the internationalization variables that are unset or null.
24735
```

24736

24737

(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,

Internationalization Variables for the precedence of internationalization variables

Utilities mkdir

24738		used to determine the values of locale categories.)	
24739 24740	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.	
24741 24742 24743	LC_CTYPE	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).	
24744 24745 24746	LC_MESSA	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.	
24747 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .	
24748 ASYNO 24749	CHRONOUS I Default.	EVENTS	
24750 STDOU 24751	J T Not used.		
24752 STDER 24753		d error shall be used only for diagnostic messages.	
24754 OUTPU 24755	J T FILES None.		
24756 EXTEN 24757	DED DESCR None.	IPTION	
24758 EXIT S ' 24759		ng exit values shall be returned:	
24760 24761		specified directories were created successfully or the $-\mathbf{p}$ option was specified and all cified directories now exist.	
24762	>0 An erro	r occurred.	
24763 CONSI 24764	E QUENCES O Default.	OF ERRORS	
24765 APPLIC 24766 24767 24768 24769 24770 24771	CATION USAGE The default file mode for directories is <i>a=rwx</i> (777 on most systems) with selected permissions removed in accordance with the file mode creation mask. For intermediate pathname components created by <i>mkdir</i> , the mode is the default modified by <i>u+wx</i> so that the subdirectories can always be created regardless of the file mode creation mask; if different ultimate permissions are desired for the intermediate directories, they can be changed afterwards with <i>chmod</i> .		
24772	Note that so	me of the requested directories may have been created even if an error occurs.	
24773 EXAMI 24774	PLES None.		
24775 RATIO 24776		V – m option was included to control the file mode.	
24777	The System V – p option was included to create any needed intermediate directories and to complement the functionality provided by <i>rmdir</i> for removing directories in the path prefix as		

option is also useful to ensure that a particular directory exists.

24778

24779

24780

complement the functionality provided by rmdir for removing directories in the path prefix as

they become empty. Because no error is produced if any path component already exists, the $-\mathbf{p}$

mkdir Utilities

The functionality of *mkdir* is described substantially through a reference to the *mkdir*() function in the System Interfaces volume of IEEE Std 1003.1-2001. For example, by default, the mode of the directory is affected by the file mode creation mask in accordance with the specified behavior of the *mkdir*() function. In this way, there is less duplication of effort required for describing details of the directory creation.

FUTURE DIRECTIONS

None.

24788 **SEE ALSO**

24789 chmod, rm, rmdir, umask, the System Interfaces volume of IEEE Std 1003.1-2001, mkdir()

24790 CHANGE HISTORY

First released in Issue 2.

24792 **Issue 5**

24793 The FUTURE DIRECTIONS section is added.

Utilities mkfifo

24794 NAME 24795 mkfifo — make FIFO special files 24796 SYNOPSIS 24797 mkfifo [-m mode] file...

24798 DESCRIPTION

The *mkfifo* utility shall create the FIFO special files specified by the operands, in the order specified.

For each *file* operand, the *mkfifo* utility shall perform actions equivalent to the *mkfifo*() function defined in the System Interfaces volume of IEEE Std 1003.1-2001, called with the following arguments:

- 1. The *file* operand is used as the *path* argument.
- 2. The value of the bitwise-inclusive OR of S_IRUSR, S_IWUSR, S_IRGRP, S_IWGRP, S_IROTH, and S_IWOTH is used as the *mode* argument. (If the -**m** option is specified, the value of the *mkfifo*() *mode* argument is unspecified, but the FIFO shall at no time have permissions less restrictive than the -**m** *mode* option-argument.)

24809 OPTIONS

24804

24805

24806

24807

24808

The *mkfifo* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following option shall be supported:

24813 — m mode Set the file permission bits of the newly-created FIFO to the specified mode value.

The mode option-argument shall be the same as the mode operand defined for the chmod utility. In the symbolic_mode strings, the op characters '+' and '-' shall be interpreted relative to an assumed initial mode of a=rw.

24817 OPERANDS

24818 The following operand shall be supported:

24819 *file* A pathname of the FIFO special file to be created.

24820 **STDIN**

Not used.

24822 INPUT FILES

24823 None.

24824 ENVIRONMENT VARIABLES

24825 The following environment variables shall affect the execution of *mkfifo*:

24826 LANG Provide a default value for the internationalization variables that are unset or null.
24827 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
24828 Internationalization Variables for the precedence of internationalization variables
24829 used to determine the values of locale categories.)

24830 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

24832 LC_CTYPE 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).

24835 LC_MESSAGES

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

mkfifo Utilities

24838 XSI **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 24839 ASYNCHRONOUS EVENTS 24840 Default. **24841 STDOUT** 24842 Not used. 24843 STDERR The standard error shall be used only for diagnostic messages. 24844 24845 OUTPUT FILES 24846 None. 24847 EXTENDED DESCRIPTION 24848 None. 24849 EXIT STATUS 24850 The following exit values shall be returned: 0 All the specified FIFO special files were created successfully. 24851 >0 An error occurred. 24852 24853 CONSEQUENCES OF ERRORS 24854 Default. 24855 APPLICATION USAGE 24856 None. 24857 EXAMPLES None. 24858 24859 RATIONALE This utility was added to permit shell applications to create FIFO special files. 24860 The -m option was added to control the file mode, for consistency with the similar functionality 24861 24862 provided by the *mkdir* utility. Early proposals included a $-\mathbf{p}$ option similar to the $mkdir -\mathbf{p}$ option that created intermediate 24863 directories leading up to the FIFO specified by the final component. This was removed because 24864 24865 it is not commonly needed and is not common practice with similar utilities. The functionality of *mkfifo* is described substantially through a reference to the *mkfifo*() function 24866 in the System Interfaces volume of IEEE Std 1003.1-2001. For example, by default, the mode of 24867 24868 the FIFO file is affected by the file mode creation mask in accordance with the specified behavior of the *mkfifo()* function. In this way, there is less duplication of effort required for describing 24869 details of the file creation. 24870 24871 FUTURE DIRECTIONS None. 24872 **24873 SEE ALSO** chmod, umask, the System Interfaces volume of IEEE Std 1003.1-2001, mkfifo() 24874 24875 CHANGE HISTORY

24876

First released in Issue 3.

Utilities more

24877 NAME

24883

24884

24885 24886

24887

24888

24889

24890

24891

24892

24893

24894

24895

24896

24897

24901

24902

24903

24909

24910

2491124912

24913

24914 24915

24916

24917

2491824919

24878 more — display files on a page-by-page basis

24879 SYNOPSIS

24880 UP more [-ceisu] [-n number] [-p command] [-t tagstring] [file ...]
24881

24882 **DESCRIPTION**

The *more* utility shall read files and either write them to the terminal on a page-by-page basis or filter them to standard output. If standard output is not a terminal device, all input files shall be copied to standard output in their entirety, without modification, except as specified for the –s option. If standard output is a terminal device, the files shall be written a number of lines (one screenful) at a time under the control of user commands. See the EXTENDED DESCRIPTION section.

Certain block-mode terminals do not have all the capabilities necessary to support the complete *more* definition; they are incapable of accepting commands that are not terminated with a <newline>. Implementations that support such terminals shall provide an operating mode to *more* in which all commands can be terminated with a <newline> on those terminals. This mode:

- Shall be documented in the system documentation
- Shall, at invocation, inform the user of the terminal deficiency that requires the <newline>
 usage and provide instructions on how this warning can be suppressed in future invocations
 - Shall not be required for implementations supporting only fully capable terminals
- Shall not affect commands already requiring <newline>s
- Shall not affect users on the capable terminals from using *more* as described in this volume of IEEE Std 1003.1-2001

24900 OPTIONS

-е

-i

The *more* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported:

24904 —c If a screen is to be written that has no lines in common with the current screen, or
24905 more is writing its first screen, more shall not scroll the screen, but instead shall
24906 redraw each line of the screen in turn, from the top of the screen to the bottom. In
24907 addition, if more is writing its first screen, the screen shall be cleared. This option
24908 may be silently ignored on devices with insufficient terminal capabilities.

By default, *more* shall exit immediately after writing the last line of the last file in the argument list. If the **–e** option is specified:

- 1. If there is only a single file in the argument list and that file was completely displayed on a single screen, *more* shall exit immediately after writing the last line of that file.
- 2. Otherwise, *more* shall exit only after reaching end-of-file on the last file in the argument list twice without an intervening operation. See the EXTENDED DESCRIPTION section.

Perform pattern matching in searches without regard to case; see the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.2, Regular Expression General Requirements.

more Utilities

24920 24921	- n number	Specify the number of lines per screenful. The <i>number</i> argument is a positive decimal integer. The $-\mathbf{n}$ option shall override any values obtained from any other
24922		source.
24923	- p command	Each time a screen from a new file is displayed or redisplayed (including as a
24924	F	result of <i>more</i> commands; for example, :p), execute the <i>more</i> command(s) in the
24925		command arguments in the order specified, as if entered by the user after the first
24926		screen has been displayed. No intermediate results shall be displayed (that is, if the
24927		command is a movement to a screen different from the normal first screen, only
24928		the screen resulting from the command shall be displayed.) If any of the
24929		commands fail for any reason, an informational message to this effect shall be
24930		written, and no further commands specified using the $-\mathbf{p}$ option shall be executed for this file.
24931 24932	- s	Behave as if consecutive empty lines were a single empty line.
24933	–t tagstring	Write the screenful of the file containing the tag named by the <i>tagstring</i> argument.
24934		See the <i>ctags</i> utility. The tags feature represented by -t <i>tagstring</i> and the :t command is optional. It shall be provided on any system that also provides a
24935 24936		conforming implementation of <i>ctags</i> ; otherwise, the use of –t produces undefined
24937		results.
24938 24939		The filename resulting from the –t option shall be logically added as a prefix to the list of command line files, as if specified by the user. If the tag named by the
24940		tagstring argument is not found, it shall be an error, and more shall take no further
24941		action.
24942 24943		If the tag specifies a line number, the first line of the display shall contain the beginning of that line. If the tag specifies a pattern, the first line of the display shall
24944		contain the beginning of the matching text from the first line of the file that
24945		contains that pattern. If the line does not exist in the file or matching text is not
24946		found, an informational message to this effect shall be displayed, and <i>more</i> shall
24947		display the default screen as if –t had not been specified.
24948		If both the $-t$ <i>tagstring</i> and $-p$ <i>command</i> options are given, the $-t$ <i>tagstring</i> shall be
24949		processed first; that is, the file and starting line for the display shall be as specified
24950		by -t, and then the - p <i>more</i> command shall be executed. If the line (matching text)
24951		specified by the -t command does not exist (is not found), no -p more command
24952		shall be executed for this file at any time.
24953	-u	Treat a backspace> as a printable control character, displayed as an
24954		implementation-defined character sequence (see the EXTENDED DESCRIPTION
24955		section), suppressing backspacing and the special handling that produces
24956 24957		underlined or standout mode text on some terminal types. Also, do not ignore a <carriage-return> at the end of a line.</carriage-return>
	NIDC	carriage returns at the end of a line.
24958 OPERA I		ng operand shall be supported:
24959		· · · · · · · · · · · · · · · · · · ·
24960	file	A pathname of an input file. If no <i>file</i> operands are specified, the standard input shall be used if a <i>file</i> is the transfer of the standard input shall be used at that point in the
24961 24962		shall be used. If a <i>file</i> is $'-'$, the standard input shall be read at that point in the sequence.
		sequence.
24963 STDIN	The stands	dimput shall be used only if no file anomandsifiedifie file
24964	ine standard	d input shall be used only if no <i>file</i> operands are specified, or if a <i>file</i> operand is $'-'$.

Utilities more

24965 INPUT FILES

24973

24992

24993

24994 24995

24997

24998

24999

25000

25001

24966 The input files being examined shall be text files. If standard output is a terminal, standard error shall be used to read commands from the user. If standard output is a terminal, standard error is 24967 not readable, and command input is needed, *more* may attempt to obtain user commands from 24968 the controlling terminal (for example, /dev/tty); otherwise, more shall terminate with an error 24969 24970 indicating that it was unable to read user commands. If standard output is not a terminal, no error shall result if standard error cannot be opened for reading. 24971

24972 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *more*:

24974 COLUMNS Override the system-selected horizontal display line size. See the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables for valid values 24975 24976 and results when it is unset or null.

EDITOR Used by the v command to select an editor. See the EXTENDED DESCRIPTION 24977 section. 24978

LANG Provide a default value for the internationalization variables that are unset or null. 24979 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 24980 Internationalization Variables for the precedence of internationalization variables 24981 used to determine the values of locale categories.) 24982

LC ALL If set to a non-empty string value, override the values of all the other 24983 internationalization variables. 24984

24985 LC_COLLATE

Determine the locale for the behavior of ranges, equivalence classes, and multi-24986 24987 character collating elements within regular expressions.

Determine the locale for the interpretation of sequences of bytes of text data as 24988 LC_CTYPE 24989 characters (for example, single-byte as opposed to multi-byte characters in 24990 arguments and input files) and the behavior of character classes within regular expressions. 24991

LC MESSAGES

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

NLSPATH Determine the location of message catalogs for the processing of LC MESSAGES. 24996 XSI LINES

Override the system-selected vertical screen size, used as the number of lines in a screenful. See the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables for valid values and results when it is unset or null. The -n option shall take precedence over the LINES variable for determining the number of lines in a screenful.

MORE Determine a string containing options described in the OPTIONS section preceded 25002 with hyphens and <black>-separated as on the command line. Any command line 25003 options shall be processed after those in the MORE variable, as if the command 25004 line were: 25005

more \$MORE options operands 25006

The MORE variable shall take precedence over the TERM and LINES variables for 25007 determining the number of lines in a screenful. 25008

more Utilities

Determine the name of the terminal type. If this variable is unset or null, an unspecified default terminal type is used.

25011 ASYNCHRONOUS EVENTS

25012 Default.

STDOUT

25014 The standard output shall be used to write the contents of the input files.

25015 STDERR

 The standard error shall be used for diagnostic messages and user commands (see the INPUT FILES section), and, if standard output is a terminal device, to write a prompting string. The prompting string shall appear on the screen line below the last line of the file displayed in the current screenful. The prompt shall contain the name of the file currently being examined and shall contain an end-of-file indication and the name of the next file, if any, when prompting at the end-of-file. If an error or informational message is displayed, it is unspecified whether it is contained in the prompt. If it is not contained in the prompt, it shall be displayed and then the user shall be prompted for a continuation character, at which point another message or the user prompt may be displayed. The prompt is otherwise unspecified. It is unspecified whether informational messages are written for other user commands.

25026 OUTPUT FILES

25027 None.

25028 EXTENDED DESCRIPTION

The following section describes the behavior of *more* when the standard output is a terminal device. If the standard output is not a terminal device, no options other than —s shall have any effect, and all input files shall be copied to standard output otherwise unmodified, at which time *more* shall exit without further action.

The number of lines available per screen shall be determined by the -n option, if present, or by examining values in the environment (see the ENVIRONMENT VARIABLES section). If neither method yields a number, an unspecified number of lines shall be used.

The maximum number of lines written shall be one less than this number, because the screen line after the last line written shall be used to write a user prompt and user input. If the number of lines in the screen is less than two, the results are undefined. It is unspecified whether user input is permitted to be longer than the remainder of the single line where the prompt has been written.

The number of columns available per line shall be determined by examining values in the environment (see the ENVIRONMENT VARIABLES section), with a default value as described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables.

Lines that are longer than the display shall be folded; the length at which folding occurs is unspecified, but should be appropriate for the output device. Folding may occur between glyphs of single characters that take up multiple display columns.

When standard output is a terminal and $-\mathbf{u}$ is not specified, *more* shall treat
 <code>carriage-return>s</code> specially:

• A character, followed first by a sequence of *n*
backspace>s (where *n* is the same as the number of column positions that the character occupies), then by *n* underscore characters ('_'), shall cause that character to be written as underlined text, if the terminal type supports that. The *n* underscore characters, followed first by *n*
backspace>s, then any character with *n* column positions, shall also cause that character to be written as underlined text, if the terminal type supports that.

Utilities more

• A sequence of *n*

- backspace>s (where *n* is the same as the number of column positions that the previous character occupies) that appears between two identical printable characters shall cause the first of those two characters to be written as emboldened text (that is, visually brighter, standout mode, or inverse-video mode), if the terminal type supports that, and the second to be discarded. Immediately subsequent occurrences of

-character pairs for that same character shall also be discarded. (For example, the sequence "a\ba\ba\ba\ba" is interpreted as a single emboldened 'a'.)

- The *more* utility shall logically discard all other <backspace>s from the line as well as the character which precedes them, if any.
- A <carriage-return> at the end of a line shall be ignored, rather than being written as a non-printable character, as described in the next paragraph.

It is implementation-defined how other non-printable characters are written. Implementations should use the same format that they use for the *ex* **print** command; see the OPTIONS section within the *ed* utility. It is unspecified whether a multi-column character shall be separated if it crosses a display line boundary; it shall not be discarded. The behavior is unspecified if the number of columns on the display is less than the number of columns any single character in the line being displayed would occupy.

When each new file is displayed (or redisplayed), *more* shall write the first screen of the file. Once the initial screen has been written, *more* shall prompt for a user command. If the execution of the user command results in a screen that has lines in common with the current screen, and the device has sufficient terminal capabilities, *more* shall scroll the screen; otherwise, it is unspecified whether the screen is scrolled or redrawn.

For all files but the last (including standard input if no file was specified, and for the last file as well, if the —e option was not specified), when *more* has written the last line in the file, *more* shall prompt for a user command. This prompt shall contain the name of the next file as well as an indication that *more* has reached end-of-file. If the user command is **f**, <control>-F, <space>, **j**, <newline>, **d**, <control>-D, or **s**, *more* shall display the next file. Otherwise, if displaying the last file, *more* shall exit. Otherwise, *more* shall execute the user command specified.

Several of the commands described in this section display a previous screen from the input stream. In the case that text is being taken from a non-rewindable stream, such as a pipe, it is implementation-defined how much backwards motion is supported. If a command cannot be executed because of a limitation on backwards motion, an error message to this effect shall be displayed, the current screen shall not change, and the user shall be prompted for another command.

If a command cannot be performed because there are insufficient lines to display, *more* shall alert the terminal. If a command cannot be performed because there are insufficient lines to display or a / command fails: if the input is the standard input, the last screen in the file may be displayed; otherwise, the current file and screen shall not change, and the user shall be prompted for another command.

The interactive commands in the following sections shall be supported. Some commands can be preceded by a decimal integer, called *count* in the following descriptions. If not specified with the command, *count* shall default to 1. In the following descriptions, *pattern* is a basic regular expression, as described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.3, Basic Regular Expressions. The term "examine" is historical usage meaning "open the file for viewing"; for example, *more* **foo** would be expressed as examining file **foo**.

In the following descriptions, unless otherwise specified, *line* is a line in the *more* display, not a line from the file being examined.

Utilities more

In the following descriptions, the *current position* refers to two things: 1. The position of the current line on the screen 25103 The line number (in the file) of the current line on the screen 25104 Usually, the line on the screen corresponding to the current position is the third line on the 25105 screen. If this is not possible (there are fewer than three lines to display or this is the first page of 25106 25107 the file, or it is the last page of the file), then the current position is either the first or last line on the screen as described later. 25108 25109 Help Synopsis: h 25110 25111 Write a summary of these commands and other implementation-defined commands. The behavior shall be as if the *more* utility were executed with the -e option on a file that contained 25112 the summary information. The user shall be prompted as described earlier in this section when 25113 end-of-file is reached. If the user command is one of those specified to continue to the next file, 25114 more shall return to the file and screen state from which the h command was executed. 25115 **Scroll Forward One Screenful** 25116 Synopsis: [count]f 25117 25118 [count] < control > - F 25119 Scroll forward *count* lines, with a default of one screenful. If *count* is more than the screen size, only the final screenful shall be written. 25120 Scroll Backward One Screenful 25121 25122 Synopsis: [count]b [count] < control > -B 25123 Scroll backward *count* lines, with a default of one screenful (see the –**n** option). If *count* is more 25124 25125 than the screen size, only the final screenful shall be written. 25126 **Scroll Forward One Line** 25127 Synopsis: [count] < space > 25128 [count]j [count] < newline > 25129 25130 Scroll forward *count* lines. The default *count* for the <space> shall be one screenful; for j and <newline>, one line. The entire *count* lines shall be written, even if *count* is more than the screen 25131 25132 size. **Scroll Backward One Line** 25133 25134 Synopsis: [count]k Scroll backward *count* lines. The entire *count* lines shall be written, even if *count* is more than the 25135

25136

screen size.

25102

Utilities more

25137	Scroll Forward One Half Screenful
25138 25139	Synopsis: [count]d [count] <control>-D</control>
25140 25141	Scroll forward <i>count</i> lines, with a default of one half of the screen size. If <i>count</i> is specified, it shall become the new default for subsequent \mathbf{d} , <control>-D, and \mathbf{u} commands.</control>
25142	Skip Forward One Line
25143	Synopsis: [count]s
25144 25145 25146	Display the screenful beginning with the line <i>count</i> lines after the last line on the current screen. If <i>count</i> would cause the current position to be such that less than one screenful would be written, the last screenful in the file shall be written.
25147	Scroll Backward One Half Screenful
25148 25149	Synopsis: [count]u [count] <control>-U</control>
25150 25151 25152	Scroll backward <i>count</i> lines, with a default of one half of the screen size. If <i>count</i> is specified, it shall become the new default for subsequent d , <control>–D, u, and <control>–U commands. The entire <i>count</i> lines shall be written, even if <i>count</i> is more than the screen size.</control></control>
25153	Go to Beginning of File
25154	Synopsis: [count]g
25155	Display the screenful beginning with line <i>count</i> .
25156	Go to End-of-File
25157	Synopsis: [count]G
25158 25159	If <i>count</i> is specified, display the screenful beginning with the line <i>count</i> . Otherwise, display the last screenful of the file.
25160	Refresh the Screen
25161 25162	Synopsis: r <control>-L</control>
25163	Refresh the screen.
25164	Discard and Refresh
25165	Synopsis: R
25166 25167	Refresh the screen, discarding any buffered input. If the current file is non-seekable, buffered input shall not be discarded and the $\bf R$ command shall be equivalent to the $\bf r$ command.

Utilities more

25168	Mark Position
25169	Synopsis: mletter
25170	Mark the current position with the letter named by letter, where letter represents the name of one
25171	of the lowercase letters of the portable character set. When a new file is examined, all marks may
25172	be lost.
25173	Return to Mark
25174	Synopsis: 'letter
25175 25176	Return to the position that was previously marked with the letter named by <i>letter</i> , making that line the current position.
25177	Return to Previous Position
25178	Synopsis:
25179	Return to the position from which the last large movement command was executed (where a
25180	"large movement" is defined as any movement of more than a screenful of lines). If no such
25181	movements have been made, return to the beginning of the file.
25182	Search Forward for Pattern
25183	Synopsis: [count]/[!]pattern <newline></newline>
25184	Display the screenful beginning with the countth line containing the pattern. The search shall
25185	start after the first line currently displayed. The null regular expression ('/' followed by a
25186	<newline>) shall repeat the search using the previous regular expression, with a default count. If</newline>
25187	the character '!' is included, the matching lines shall be those that do not contain the pattern. If
25188	no match is found for the <i>pattern</i> , a message to that effect shall be displayed.
25189	Search Backward for Pattern
25190	Synopsis: [count]?[!]pattern <newline></newline>
25191	Display the screenful beginning with the countth previous line containing the pattern. The
25192	search shall start on the last line before the first line currently displayed. The null regular
25193	expression ('?' followed by a <newline>) shall repeat the search using the previous regular</newline>
25194	expression, with a default <i>count</i> . If the character '!' is included, matching lines shall be those
25195	that do not contain the <i>pattern</i> . If no match is found for the <i>pattern</i> , a message to that effect shall
25196	be displayed.
25197	Repeat Search
25198	Synopsis: [count]n
25199 25200	Repeat the previous search for <i>count</i> th line containing the last <i>pattern</i> (or not containing the last <i>pattern</i> , if the previous search was "/!" or "?!").

Utilities more

25201 Repeat Search in Reverse 25202 Synopsis: [count] N Repeat the search in the opposite direction of the previous search for the *count*th line containing 25203 25204 the last *pattern* (or not containing the last *pattern*, if the previous search was "/!" or "?!"). **Examine New File** 25205 25206 Synopsis: :e [filename] < newline > Examine a new file. If the filename argument is not specified, the current file (see the :n and :p 25207 25208 commands below) shall be re-examined. The filename shall be subjected to the process of shell word expansions (see Section 2.6 (on page 36)); if more than a single pathname results, the 25209 25210 effects are unspecified. If filename is a number sign ('#'), the previously examined file shall be re-examined. If *filename* is not accessible for any reason (including that it is a non-seekable file), 25211 an error message to this effect shall be displayed and the current file and screen shall not change. 25212 25213 **Examine Next File** 25214 Synopsis: [count]:n Examine the next file. If a number *count* is specified, the *count*th next file shall be examined. If 25215 25216 filename refers to a non-seekable file, the results are unspecified. **Examine Previous File** 25217 25218 Synopsis: [count]:p 25219 Examine the previous file. If a number count is specified, the countth previous file shall be examined. If *filename* refers to a non-seekable file, the results are unspecified. 25220 Go to Tag 25221 25222 Synopsis: :t tagstring<newline> If the file containing the tag named by the tagstring argument is not the current file, examine the 25223 file, as if the :e command was executed with that file as the argument. Otherwise, or in addition, 25224 display the screenful beginning with the tag, as described for the -t option (see the OPTIONS 25225 25226 section). If the ctags utility is not supported by the system, the use of :t produces undefined 25227 results. **Invoke Editor** 25228 25229 Synopsis: Invoke an editor to edit the current file being examined. If standard input is being examined, the 25230 results are unspecified. The name of the editor shall be taken from the environment variable 25231 EDITOR, or shall default to vi. If the last pathname component in EDITOR is either vi or ex, the 25232 editor shall be invoked with a -c linenumber command line argument, where linenumber is the 25233 line number of the file line containing the display line currently displayed as the first line of the 25234 25235 screen. It is implementation-defined whether line-setting options are passed to editors other 25236 than *vi* and *ex*. When the editor exits, *more* shall resume with the same file and screen as when the editor was 25237

invoked.

25238

more Utilities

25239	Display Position
25240	Synopsis: =
25241	<control>-G</control>
25242	Write a message for which the information references the first byte of the line after the last line of
25243	the file on the screen. This message shall include the name of the file currently being examined,
25244	its number relative to the total number of files there are to examine, the line number in the file,
25245	the byte number and the total bytes in the file, and what percentage of the file precedes the
25246	current position. If <i>more</i> is reading from standard input, or the file is shorter than a single screen,
25247	the line number, the byte number, the total bytes, and the percentage need not be written.
	P
25248	Quit
25249	Synopsis: q
25250	: q
25251	ZZ
25252	Exit more.
25253 EXIT S	
25254	The following exit values shall be returned:
25255	0 Successful completion.
25256	>0 An error occurred.
25257 CONSI	EQUENCES OF ERRORS
25258	If an error is encountered accessing a file when using the :n command, more shall attempt to
25259	examine the next file in the argument list, but the final exit status shall be affected. If an error is
25260	encountered accessing a file via the :p command, more shall attempt to examine the previous file
25261	in the argument list, but the final exit status shall be affected. If an error is encountered accessing
25262	a file via the :e command, more shall remain in the current file and the final exit status shall not
25263	be affected.
25264 APPLIC	CATION USAGE
25265	When the standard output is not a terminal, only the $-s$ filter-modification option is effective.
25266	This is based on historical practice. For example, a typical implementation of man pipes its
25267	output through <i>more</i> – s to squeeze excess white space for terminal users. When <i>man</i> is piped to
25268	<i>lp</i> , however, it is undesirable for this squeezing to happen.
25269 EXAM l	PLES
25270	The $-\mathbf{p}$ allows arbitrary commands to be executed at the start of each file. Examples are:
25271	more -p G file1 file2
25272	Examine each file starting with its last screenful.
25273	more - p 100 file1 file2
25274	Examine each file starting with line 100 in the current position (usually the third line, so line
25275	98 would be the first line written).
25276	more -p /100 file1 file2
25277	Examine each file starting with the first line containing the string "100" in the current
25278	position
25279 RATIO	NALE
25280	The <i>more</i> utility, available in BSD and BSD-derived systems, was chosen as the prototype for the
05001	DOSIV file display program since it is more widely available then either the public domain

25281

25282

POSIX file display program since it is more widely available than either the public-domain

program less or than pg, a pager provided in System V. The 4.4 BSD more is the model for the

Utilities more

features selected; it is almost fully upwards-compatible from the 4.3 BSD version in wide use and has become more amenable for *vi* users. Several features originally derived from various file editors, found in both *less* and *pg*, have been added to this volume of IEEE Std 1003.1-2001 as they have proved extremely popular with users.

There are inconsistencies between *more* and *vi* that result from historical practice. For example, the single-character commands **h**, **f**, **b**, and <space> are screen movers in *more*, but cursor movers in *vi*. These inconsistencies were maintained because the cursor movements are not applicable to *more* and the powerful functionality achieved without the use of the control key justifies the differences.

The tags interface has been included in a program that is not a text editor because it promotes another degree of consistent operation with *vi*. It is conceivable that the paging environment of *more* would be superior for browsing source code files in some circumstances.

The operating mode referred to for block-mode terminals effectively adds a <newline> to each Synopsis line that currently has none. So, for example, **d**<newline> would page one screenful. The mode could be triggered by a command line option, environment variable, or some other method. The details are not imposed by this volume of IEEE Std 1003.1-2001 because there are so few systems known to support such terminals. Nevertheless, it was considered that all systems should be able to support *more* given the exception cited for this small community of terminals because, in comparison to *vi*, the cursor movements are few and the command set relatively amenable to the optional <newline>s.

Some versions of *more* provide a shell escaping mechanism similar to the *ex*! command. The standard developers did not consider that this was necessary in a paginator, particularly given the wide acceptance of multiple window terminals and job control features. (They chose to retain such features in the editors and *mailx* because the shell interaction also gives an opportunity to modify the editing buffer, which is not applicable to *more*.)

The $-\mathbf{p}$ (position) option replaces the + command because of the Utility Syntax Guidelines. In early proposals, it took a *pattern* argument, but historical *less* provided the *more* general facility of a command. It would have been desirable to use the same $-\mathbf{c}$ as ex and vi, but the letter was already in use.

The text stating "from a non-rewindable stream ... implementations may limit the amount of backwards motion supported" would allow an implementation that permitted no backwards motion beyond text already on the screen. It was not possible to require a minimum amount of backwards motion that would be effective for all conceivable device types. The implementation should allow the user to back up as far as possible, within device and reasonable memory allocation constraints.

Historically, non-printable characters were displayed using the ARPA standard mappings, which are as follows:

1. Printable characters are left alone.

- 2. Control characters less than \177 are represented as followed by the character offset from the '@' character in the ASCII map; for example, \007 is represented as 'G'.
- 3. \177 is represented as followed by '?'.

The display of characters having their eighth bit set was less standard. Existing implementations use hex (0x00), octal (\setminus 000), and a meta-bit display. (The latter displayed characters with their eighth bit set as the two characters "M-", followed by the seven-bit display as described previously.) The latter probably has the best claim to historical practice because it was used with the $-\mathbf{v}$ option of 4 BSD and 4 BSD-derived versions of the cat utility since 1980.

more Utilities

25329 No specific display format is required by IEEE Std 1003.1-2001. Implementations are encouraged 25330 to conform to historic practice in the absence of any strong reason to diverge. **25331 FUTURE DIRECTIONS** 25332 None. 25333 SEE ALSO 25334 Chapter 2 (on page 29), ctags, ed, ex, vi 25335 CHANGE HISTORY 25336 First released in Issue 4. 25337 Issue 5 The FUTURE DIRECTIONS section is added. 25338 25339 Issue 6 This utility is marked as part of the User Portability Utilities option. 25340 25341 The obsolescent SYNOPSIS is removed. 25342 The utility has been extensively reworked for alignment with the IEEE P1003.2b draft standard: • Changes have been made as a result of IEEE PASC Interpretations 1003.2 #37 and #109. 25343 • The more utility should be able to handle underlined and emboldened displays of characters 25344 25345 that are wider than a single column position.

Utilities mv

```
25346 NAME
25347 mv — move files
25348 SYNOPSIS
25349 mv [-fi] source_file target_file
25350 mv [-fi] source_file... target_file
```

DESCRIPTION

In the first synopsis form, the *mv* utility shall move the file named by the *source_file* operand to the destination specified by the *target_file*. This first synopsis form is assumed when the final operand does not name an existing directory and is not a symbolic link referring to an existing directory.

In the second synopsis form, *mv* shall move each file named by a *source_file* operand to a destination file in the existing directory named by the *target_dir* operand, or referenced if *target_dir* is a symbolic link referring to an existing directory. The destination path for each *source_file* shall be the concatenation of the target directory, a single slash character, and the last pathname component of the *source_file*. This second form is assumed when the final operand names an existing directory.

If any operand specifies an existing file of a type not specified by the System Interfaces volume of IEEE Std 1003.1-2001, the behavior is implementation-defined.

For each *source_file* the following steps shall be taken:

- 1. If the destination path exists, the **-f** option is not specified, and either of the following conditions is true:
 - a. The permissions of the destination path do not permit writing and the standard input is a terminal.
 - b. The -i option is specified.

the *mv* utility shall write a prompt to standard error and read a line from standard input. If the response is not affirmative, *mv* shall do nothing more with the current *source_file* and go on to any remaining *source_files*.

- 2. The *mv* utility shall perform actions equivalent to the *rename*() function defined in the System Interfaces volume of IEEE Std 1003.1-2001, called with the following arguments:
 - a. The source_file operand is used as the old argument.
 - b. The destination path is used as the *new* argument.

If this succeeds, *mv* shall do nothing more with the current *source_file* and go on to any remaining *source_files*. If this fails for any reasons other than those described for the *errno* [EXDEV] in the System Interfaces volume of IEEE Std 1003.1-2001, *mv* shall write a diagnostic message to standard error, do nothing more with the current *source_file*, and go on to any remaining *source_files*.

- 3. If the destination path exists, and it is a file of type directory and *source_file* is not a file of type directory, or it is a file not of type directory and *source_file* is a file of type directory, *mv* shall write a diagnostic message to standard error, do nothing more with the current *source_file*, and go on to any remaining *source_files*.
- 4. If the destination path exists, *mv* shall attempt to remove it. If this fails for any reason, *mv* shall write a diagnostic message to standard error, do nothing more with the current *source_file*, and go on to any remaining *source_files*.

mv Utilities

25389 5. The file hierarchy rooted in source_file shall be duplicated as a file hierarchy rooted in the 25390 destination path. If source_file or any of the files below it in the hierarchy are symbolic links, the links themselves shall be duplicated, including their contents, rather than any 25391 files to which they refer. The following characteristics of each file in the file hierarchy shall 25392 25393 be duplicated: The time of last data modification and time of last access 25394 The user ID and group ID 25395 The file mode 25396 25397 If the user ID, group ID, or file mode of a regular file cannot be duplicated, the file mode bits S_ISUID and S_ISGID shall not be duplicated. 25398 When files are duplicated to another file system, the implementation may require that the 25399 process invoking *mv* has read access to each file being duplicated. 25400 If the duplication of the file hierarchy fails for any reason, mv shall write a diagnostic 25401 25402 message to standard error, do nothing more with the current source_file, and go on to any remaining source_files. 25403 If the duplication of the file characteristics fails for any reason, mv shall write a diagnostic 25404 25405 message to standard error, but this failure shall not cause *mv* to modify its exit status. 6. The file hierarchy rooted in *source_file* shall be removed. If this fails for any reason, *mv* shall 25406 write a diagnostic message to the standard error, do nothing more with the current 25407 25408 *source_file*, and go on to any remaining *source_files*. 25409 OPTIONS The mv utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, 25410 Utility Syntax Guidelines. 25411 The following options shall be supported: 25412 $-\mathbf{f}$ Do not prompt for confirmation if the destination path exists. Any previous 25413 25414 occurrence of the -i option is ignored. $-\mathbf{i}$ Prompt for confirmation if the destination path exists. Any previous occurrence of 25415 the -f option is ignored. 25416 Specifying more than one of the -f or -i options shall not be considered an error. The last option 25417 specified shall determine the behavior of *mv*. 25418 25419 OPERANDS The following operands shall be supported: 25420 A pathname of a file or directory to be moved. 25421 source_file target file A new pathname for the file or directory being moved. 25422 25423 target_dir A pathname of an existing directory into which to move the input files. 25424 **STDIN**

the STDERR section. Otherwise, the standard input shall not be used. 25427 **INPUT FILES**

25425

25428

The input files specified by each *source_file* operand can be of any file type.

The standard input shall be used to read an input line in response to each prompt specified in

Utilities mv

25429 ENVIRONMENT VARIABLES 25430 The following environment variables shall affect the execution of *mv*: 25431 Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 25432 Internationalization Variables for the precedence of internationalization variables 25433 used to determine the values of locale categories.) 25434 LC_ALL If set to a non-empty string value, override the values of all the other 25435 internationalization variables. 25436 25437 LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multi-25438 character collating elements used in the extended regular expression defined for 25439 the **yesexpr** locale keyword in the *LC_MESSAGES* category. 25440 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 25441 25442 characters (for example, single-byte as opposed to multi-byte characters in 25443 arguments and input files), the behavior of character classes used in the extended regular expression defined for the **yesexpr** locale keyword in the *LC_MESSAGES* 25444 category. 25445 LC MESSAGES 25446 Determine the locale for the processing of affirmative responses that should be 25447 used to affect the format and contents of diagnostic messages written to standard 25448 25449 error. **NLSPATH** 25450 XSI Determine the location of message catalogs for the processing of *LC_MESSAGES*. 25451 ASYNCHRONOUS EVENTS Default. 25452 25453 STDOUT Not used. 25454 25455 STDERR Prompts shall be written to the standard error under the conditions specified in the 25456 25457 DESCRIPTION section. The prompts shall contain the destination pathname, but their format is otherwise unspecified. Otherwise, the standard error shall be used only for diagnostic messages. 25458 25459 OUTPUT FILES The output files may be of any file type. 25460 25461 EXTENDED DESCRIPTION None. 25462 25463 EXIT STATUS The following exit values shall be returned: 25464 All input files were moved successfully. 25465 >0 An error occurred. 25466 25467 CONSEQUENCES OF ERRORS If the copying or removal of source file is prematurely terminated by a signal or error, mv may 25468 leave a partial copy of source_file at the source or destination. The mv utility shall not modify 25469 both source_file and the destination path simultaneously; termination at any point shall leave 25470 either *source_file* or the destination path complete. 25471

mv Utilities

25472 APPLICATION USAGE

Some implementations mark for update the *st_ctime* field of renamed files and some do not.

Applications which make use of the *st_ctime* field may behave differently with respect to renamed files unless they are designed to allow for either behavior.

25476 EXAMPLES

If the current directory contains only files **a** (of any type defined by the System Interfaces volume of IEEE Std 1003.1-2001), **b** (also of any type), and a directory **c**:

25479 mv a b c 25480 mv c d

results with the original files \mathbf{a} and \mathbf{b} residing in the directory \mathbf{d} in the current directory.

25482 RATIONALE

Early proposals diverged from the SVID and BSD historical practice in that they required that when the destination path exists, the –f option is not specified, and input is not a terminal, *mv* fails. This was done for compatibility with *cp*. The current text returns to historical practice. It should be noted that this is consistent with the *rename()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001, which does not require write permission on the target.

For absolute clarity, paragraph (1), describing the behavior of *mv* when prompting for confirmation, should be interpreted in the following manner:

```
if (exists AND (NOT f_option) AND
      ((not_writable AND input_is_terminal) OR i_option))
```

The $-\mathbf{i}$ option exists on BSD systems, giving applications and users a way to avoid accidentally unlinking files when moving others. When the standard input is not a terminal, the 4.3 BSD mv deletes all existing destination paths without prompting, even when $-\mathbf{i}$ is specified; this is inconsistent with the behavior of the 4.3 BSD cp utility, which always generates an error when the file is unwritable and the standard input is not a terminal. The standard developers decided that use of $-\mathbf{i}$ is a request for interaction, so when the destination path exists, the utility takes instructions from whatever responds to standard input.

The rename() function is able to move directories within the same file system. Some historical versions of mv have been able to move directories, but not to a different file system. The standard developers considered that this was an annoying inconsistency, so this volume of IEEE Std 1003.1-2001 requires directories to be able to be moved even across file systems. There is no $-\mathbf{R}$ option to confirm that moving a directory is actually intended, since such an option was not required for moving directories in historical practice. Requiring the application to specify it sometimes, depending on the destination, seemed just as inconsistent. The semantics of the rename() function were preserved as much as possible. For example, mv is not permitted to "rename" files to or from directories, even though they might be empty and removable.

Historic implementations of *mv* did not exit with a non-zero exit status if they were unable to duplicate any file characteristics when moving a file across file systems, nor did they write a diagnostic message for the user. The former behavior has been preserved to prevent scripts from breaking; a diagnostic message is now required, however, so that users are alerted that the file characteristics have changed.

The exact format of the interactive prompts is unspecified. Only the general nature of the contents of prompts are specified because implementations may desire more descriptive prompts than those used on historical implementations. Therefore, an application not using the $-\mathbf{f}$ option or using the $-\mathbf{i}$ option relies on the system to provide the most suitable dialog directly with the user, based on the behavior specified.

Utilities mv

When mv is dealing with a single file system and source_file is a symbolic link, the link itself is 25519 25520 moved as a consequence of the dependence on the rename() functionality, per the DESCRIPTION. Across file systems, this has to be made explicit. 25521 25522 FUTURE DIRECTIONS 25523 None. 25524 SEE ALSO *cp*, *ln*, the System Interfaces volume of IEEE Std 1003.1-2001, *rename*() 25525 25526 CHANGE HISTORY First released in Issue 2. 25527 25528 **Issue 6** The mv utility is changed to describe processing of symbolic links as specified in the 25529 IEEE P1003.2b draft standard. 25530 The APPLICATION USAGE section is added. 25531

newgrp Utilities

```
25532 NAME
```

25533 newgrp — change to a new group

25534 SYNOPSIS

25535 UP newgrp [-1] [group]

DESCRIPTION

The *newgrp* utility shall create a new shell execution environment with a new real and effective group identification. Of the attributes listed in Section 2.12 (on page 61), the new shell execution environment shall retain the working directory, file creation mask, and exported variables from the previous environment (that is, open files, traps, unexported variables, alias definitions, shell functions, and *set* options may be lost). All other aspects of the process environment that are preserved by the *exec* family of functions defined in the System Interfaces volume of IEEE Std 1003.1-2001 shall also be preserved by *newgrp*; whether other aspects are preserved is unspecified.

A failure to assign the new group identifications (for example, for security or password-related reasons) shall not prevent the new shell execution environment from being created.

The *newgrp* utility shall affect the supplemental groups for the process as follows:

- On systems where the effective group ID is normally in the supplementary group list (or whenever the old effective group ID actually is in the supplementary group list):
 - If the new effective group ID is also in the supplementary group list, *newgrp* shall change the effective group ID.
 - If the new effective group ID is not in the supplementary group list, *newgrp* shall add the new effective group ID to the list, if there is room to add it.
- On systems where the effective group ID is not normally in the supplementary group list (or whenever the old effective group ID is not in the supplementary group list):
 - If the new effective group ID is in the supplementary group list, *newgrp* shall delete it.
 - If the old effective group ID is not in the supplementary list, *newgrp* shall add it if there is room.

Note: The System Interfaces volume of IEEE Std 1003.1-2001 does not specify whether the effective group ID of a process is included in its supplementary group list.

With no operands, *newgrp* shall change the effective group back to the groups identified in the user's user entry, and shall set the list of supplementary groups to that set in the user's group database entries.

If a password is required for the specified group, and the user is not listed as a member of that group in the group database, the user shall be prompted to enter the correct password for that group. If the user is listed as a member of that group, no password shall be requested. If no password is required for the specified group, it is implementation-defined whether users not listed as members of that group can change to that group. Whether or not a password is required, implementation-defined system accounting or security mechanisms may impose additional authorization restrictions that may cause *newgrp* to write a diagnostic message and suppress the changing of the group identification.

OPTIONS

The *newgrp* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

Utilities newgrp

25576 The following option shall be supported: $-\mathbf{l}$ (The letter ell.) Change the environment to what would be expected if the user 25577 25578 actually logged in again. 25579 OPERANDS 25580 The following operand shall be supported: A group name from the group database or a non-negative numeric group ID. 25581 group 25582 Specifies the group ID to which the real and effective group IDs shall be set. If group is a non-negative numeric string and exists in the group database as a group 25583 name (see getgrnam()), the numeric group ID associated with that group name 25584 shall be used as the group ID. 25585 25586 **STDIN** Not used. 25587 25588 INPUT FILES The file /dev/tty shall be used to read a single line of text for password checking, when one is 25589 25590 required. 25591 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *newgrp*: 25592 Provide a default value for the internationalization variables that are unset or null. LANG 25593 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 25594 25595 Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 25596 LC ALL If set to a non-empty string value, override the values of all the other 25597 internationalization variables. 25598 25599 LC CTYPE 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 25600 arguments). 25601 LC MESSAGES 25602 Determine the locale that should be used to affect the format and contents of 25603 diagnostic messages written to standard error. 25604 **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 25605 XSI 25606 ASYNCHRONOUS EVENTS Default. 25607 25608 **STDOUT** Not used. 25609 25610 STDERR 25611 The standard error shall be used for diagnostic messages and a prompt string for a password, if one is required. Diagnostic messages may be written in cases where the exit status is not 25612 available. See the EXIT STATUS section. 25613 25614 OUTPUT FILES None. 25615 25616 EXTENDED DESCRIPTION 25617 None.

Utilities newgrp

25618 EXIT STATUS

25624

25629

25630

25631

25632 25633

25634

25635 25636

25637

25640

25641

25642

25643 25644

25648

25649

25650

25651

25652

25653

25654

25655

25656 25657

25658

25659

25660

25661 25662

25663

25619 If newgrp succeeds in creating a new shell execution environment, whether or not the group 25620 identification was changed successfully, the exit status shall be the exit status of the shell. 25621

Otherwise, the following exit value shall be returned:

25622 >0 An error occurred.

25623 CONSEQUENCES OF ERRORS

The invoking shell may terminate.

25625 APPLICATION USAGE

25626 There is no convenient way to enter a password into the group database. Use of group 25627 passwords is not encouraged, because by their very nature they encourage poor security practices. Group passwords may disappear in the future. 25628

> A common implementation of *newgrp* is that the current shell uses *exec* to overlay itself with newgrp, which in turn overlays itself with a new shell after changing group. On some implementations, however, this may not occur and *newgrp* may be invoked as a subprocess.

> The newgrp command is intended only for use from an interactive terminal. It does not offer a useful interface for the support of applications.

> The exit status of *newgrp* is generally inapplicable. If *newgrp* is used in a script, in most cases it successfully invokes a new shell and the rest of the original shell script is bypassed when the new shell exits. Used interactively, newgrp displays diagnostic messages to indicate problems. But usage such as:

newgrp foo 25638 25639 echo \$?

> is not useful because the new shell might not have access to any status newgrp may have generated (and most historical systems do not provide this status). A zero status echoed here does not necessarily indicate that the user has changed to the new group successfully. Following newgrp with the id command provides a portable means of determining whether the group change was successful or not.

25645 EXAMPLES

25646 None

25647 RATIONALE

Most historical implementations use one of the *exec* functions to implement the behavior of newgrp. Errors detected before the exec leave the environment unchanged, while errors detected after the *exec* leave the user in a changed environment. While it would be useful to have *newgrp* issue a diagnostic message to tell the user that the environment changed, it would be inappropriate to require this change to some historical implementations.

The password mechanism is allowed in the group database, but how this would be implemented is not specified.

The newgrp utility was retained in this volume of IEEE Std 1003.1-2001, even given the existence of the multiple group permissions feature in the System Interfaces volume of IEEE Std 1003.1-2001, for several reasons. First, in some implementations, the group ownership of a newly created file is determined by the group of the directory in which the file is created, as allowed by the System Interfaces volume of IEEE Std 1003.1-2001; on other implementations, the group ownership of a newly created file is determined by the effective group ID. On implementations of the latter type, newgrp allows files to be created with a specific group ownership. Finally, many implementations use the real group ID in accounting, and on such systems, *newgrp* allows the accounting identity of the user to be changed.

Utilities newgrp

25664 FUTURE DIRECTIONS

25665 None.

25666 SEE ALSO

25667 Chapter 2 (on page 29), sh, the System Interfaces volume of IEEE Std 1003.1-2001, exec,

25668 getgrnam()

25669 CHANGE HISTORY

First released in Issue 2.

25671 **Issue 6**

25672 This utility is marked as part of the User Portability Utilities option.

The obsolescent SYNOPSIS is removed.

25674 The text describing supplemental groups is no longer conditional on {NGROUPS_MAX} being

greater than 1. This is because {NGROUPS_MAX} now has a minimum value of 8. This is a FIPS

requirement.

nice **Utilities**

25677 **NAME**

25678 nice — invoke a utility with an altered nice value

25679 SYNOPSIS

nice [-n increment] utility [argument...] 25680 UP

25681

25682 **DESCRIPTION**

The nice utility shall invoke a utility, requesting that it be run with a different nice value (see the 25683 Base Definitions volume of IEEE Std 1003.1-2001, Section 3.239, Nice Value). With no options 25684 and only if the user has appropriate privileges, the executed utility shall be run with a nice value 25685 25686 that is some implementation-defined quantity less than or equal to the nice value of the current process. If the user lacks appropriate privileges to affect the nice value in the requested manner, 25687 the nice utility shall not affect the nice value; in this case, a warning message may be written to 25688 standard error, but this shall not prevent the invocation of *utility* or affect the exit status. 25689

25690 OPTIONS

The *nice* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 25691 12.2, Utility Syntax Guidelines. 25692

The following option is supported: 25693

-n increment A positive or negative decimal integer which shall have the same effect on the 25694 execution of the utility as if the utility had called the nice() function with the 25695 numeric value of the *increment* option-argument. 25696

25697 OPERANDS

The following operands shall be supported: 25698

The name of a utility that is to be invoked. If the *utility* operand names any of the utility 25699 special built-in utilities in Section 2.14 (on page 64), the results are undefined. 25700

Any string to be supplied as an argument when invoking the utility named by the 25701 argument 25702

utility operand.

25703 **STDIN**

25704 Not used.

25705 INPUT FILES

25706 None.

25707 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *nice*: 25708

LANG Provide a default value for the internationalization variables that are unset or null. 25709 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 25710 Internationalization Variables for the precedence of internationalization variables 25711 used to determine the values of locale categories.) 25712

LC ALL If set to a non-empty string value, override the values of all the other 25713 internationalization variables. 25714

Determine the locale for the interpretation of sequences of bytes of text data as 25715 LC_CTYPE 25716 characters (for example, single-byte as opposed to multi-byte characters in arguments). 25717

LC_MESSAGES 25718

Determine the locale that should be used to affect the format and contents of 25719 25720 diagnostic messages written to standard error.

Utilities nice

PATH Determine the search path used to locate the utility to be invoked. See the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables. Path Str.	25721 XSI	<i>NLSPATH</i> Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .
25728 STDOUT 25728 STDERR 25729 The standard error shall be used only for diagnostic messages. 25730 OUTPUT FILES 25731 None. 25732 EXTENDED DESCRIPTION 25733 None. 25732 EXTENDED DISCRIPTION 25733 None. 25735 If utility is invoked, the exit status of nice shall be the exit status of utility; otherwise, the nice utility shall exit with one of the following values: 25736 If utility shall exit with one of the following values: 25737 1-125 An error occurred in the nice utility. 25738 126 The utility specified by utility was found but could not be invoked. 25739 127 The utility specified by utility could not be found. 25740 CONSEQUENCES OF ERRORS 25741 Default. 25742 APPLICATION USAGE 25743 The only guaranteed portable uses of this utility are: 25744 nice utility 25745 Run utility with the default lower nice value. 25746 nice - ¬ < positive integer> utility 25747 Run utility with a lower nice value. 25748 On some implementations they have no discernible effect on the invoked utility and on some others they are exactly equivalent. 25750 Historical systems have frequently supported the < positive integer> up to 20. Since there is no error penalty associated with guessing a number that is too high, users without access to the system conformance document (to see what limits are actually in place) could use the historical 25753 1 to 20 range or attempt to use very large numbers if the job should be truly low priority. 25756 The command. env. nice. nohup. time. 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 comitions" and the values above 128 can be confused with termination due to receipt of a signal. The value 126 was cripts produce meaningful error messages differentiating the 128 and 127 casses. The distinction of the confused with terminat		ı J
25727 Not used. 25728 STDERR 25739 OUTPUT FILES 25731 None. 25732 EXTENDED DESCRIPTION 25733 None. 25734 EXT STATUS 25735 If utility is invoked, the exit status of nice shall be the exit status of utility; otherwise, the nice utility shall exit with one of the following values: 25736 utility shall exit with one of the following values: 25737 1-125 An error occurred in the nice utility. 25738 126 The utility specified by utility was found but could not be invoked. 25740 CONSEQUENCES OF ERRORS 25741 Default. 25742 APPLICATION USAGE 25743 The only guaranteed portable uses of this utility are: 25744 nice utility 25745 Run utility with the default lower nice value. 25746 nice — - positive integer> utility 25747 Run utility with a lower nice value. 25748 On some implementations they have no discernible effect on the invoked utility and on some others they are exactly equivalent. 25750 Historical systems have frequently supported the <positive integer=""> up to 20. Since there is no others they are exactly equivalent. 25751 The nice value of a process can be displayed using the command: 25752 The ince walue of a process can be displayed using the command: 25753 The command, env, nice, nohup, time, 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 128 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 scripts produce meaningful error messages differentiating the 126 and 127 cases. The distinction scripts produce meaningful error messages differentiating the 126 and 127 cases. Th</positive>		
25729 The standard error shall be used only for diagnostic messages. 25730 OUTPUT FILES 25731 None. 25732 EXTENDED DESCRIPTION 25733 None. 25734 EXIT STATUS 25735 If utility is invoked, the exit status of nice shall be the exit status of utility; otherwise, the nice utility shall exit with one of the following values: 25736 12 The utility specified by utility was found but could not be invoked. 25738 126 The utility specified by utility was found but could not be invoked. 25739 127 The utility specified by utility could not be found. 25740 CONSEQUENCES OF ERRORS 25741 Default. 25742 APPLICATION USAGE 25743 The only guaranteed portable uses of this utility are: 25744 nice utility 25745 Run utility with the default lower nice value. 25746 nice		
25732 EXTENDED DESCRIPTION 25733 None. 25734 EXIT STATUS 25735 If utility is invoked, the exit status of nice shall be the exit status of utility; otherwise, the nice utility shall exit with one of the following values: 25737 1-125 An error occurred in the nice utility. 25738 126 The utility specified by utility was found but could not be invoked. 25739 127 The utility specified by utility could not be found. 25740 CONSEQUENCES OF ERRORS 25741 Default. 25742 APPLICATION USAGE 25743 The only guaranteed portable uses of this utility are: 25744 nice utility 25744 nice utility 25745 Run utility with the default lower nice value. 25746 nice <pre>25747</pre>		
25734 EXIT STATUS 25735 If utility is invoked, the exit status of nice shall be the exit status of utility; otherwise, the nice utility shall exit with one of the following values: 25736 1-125 An error occurred in the nice utility. 25738 126 The utility specified by utility was found but could not be invoked. 25739 127 The utility specified by utility could not be found. 25740 CONSEQUENCES OF ERRORS 25741 Default. 25742 APPLICATION USAGE 25743 The only guaranteed portable uses of this utility are: 25744 nice utility 25745 Run utility with the default lower nice value. 25746 nice -n <pre>- positive integer> utility 25747 Run utility with a lower nice value. 25748 On some implementations they have no discernible effect on the invoked utility and on some others they are exactly equivalent. 25750 Historical systems have frequently supported the <pre>- positive integer> up to 20</pre> Since there is no error penalty associated with guessing a number that is too high, users without access to the system conformance document (to see what limits are actually in place) could use the historical 1 to 20 range or attempt to use very large numbers if the job should be truly low priority. 25754 The nice value of a process can be displayed using the command: 25755 ps -o nice 25756 The command, env, nice, nohup, time, 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</pre>		
If utility is invoked, the exit status of nice shall be the exit status of utility; otherwise, the nice utility shall exit with one of the following values: 1-125		
25738 126 The utility specified by utility was found but could not be invoked. 25740 CONSEQUENCES OF ERRORS 25741 Default. 25742 APPLICATION USAGE 25742 The only guaranteed portable uses of this utility are: 25744 nice utility 25745 Run utility with the default lower nice value. 25746 nice –n <pre> 25747 run utility with a lower nice value. 25748 On some implementations they have no discernible effect on the invoked utility and on some others they are exactly equivalent. 25750 Historical systems have frequently supported the <pre> 25750 rerror penalty associated with guessing a number that is too high, users without access to the system conformance document (to see what limits are actually in place) could use the historical 1 to 20 range or attempt to use very large numbers if the job should be truly low priority. 25754 The nice value of a process can be displayed using the command: 25755 ps -o nice 25756 The command, env, nice, nohup, time, 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 are recovered as a construction of the receipt of a signal. The value 126 was chosen because 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 of the receipt of a signal to the script produce meaningful error messages differentiating the 126 and 127 cases. The distinction of the receipt of a signal to the script of a signal to the scr</pre></pre>	25735	If utility is invoked, the exit status of nice shall be the exit status of utility; otherwise, the nice
25739 127 The utility specified by utility could not be found. 25740 CONSEQUENCES OF ERRORS 25741 Default. 25742 APPLICATION USAGE 25743 The only guaranteed portable uses of this utility are: 25744 nice utility 25745 Run utility with the default lower nice value. 25746 nice -n <pre> 25747 run utility with a lower nice value. 25748 On some implementations they have no discernible effect on the invoked utility and on some others they are exactly equivalent. 25750 Historical systems have frequently supported the <pre> 25750 error penalty associated with guessing a number that is too high, users without access to the system conformance document (to see what limits are actually in place) could use the historical I to 20 range or attempt to use very large numbers if the job should be truly low priority. 25754 The nice value of a process can be displayed using the command: 25755 ps -o nice 25756 The command, env, nice, nohup, time, 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</pre></pre>	25737	1-125 An error occurred in the <i>nice</i> utility.
25740 CONSEQUENCES OF ERRORS 25741 Default. 25742 APPLICATION USAGE 25743 The only guaranteed portable uses of this utility are: 25744 nice utility 25745 Run utility with the default lower nice value. 25746 nice -n <pre> 25747 Run utility with a lower nice value. 25748 On some implementations they have no discernible effect on the invoked utility and on some others they are exactly equivalent. 25750 Historical systems have frequently supported the <pre> 25750 error penalty associated with guessing a number that is too high, users without access to the system conformance document (to see what limits are actually in place) could use the historical I to 20 range or attempt to use very large numbers if the job should be truly low priority. 25754 The nice value of a process can be displayed using the command: 25755 ps -o nice 25756 The command, env, nice, nohup, time, 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</pre></pre>	25738	126 The utility specified by <i>utility</i> was found but could not be invoked.
Default. Default. Default. Default. Default. Default. Default. Default. Default. Default. Default. Default. Default. The only guaranteed portable uses of this utility are: nice utility Run utility with the default lower nice value. Default. Nunce -n < positive integer> utility Run utility with a lower nice value. Default. Defautility Run utility with the default lower nice value. Defautility with a lower nice value. Defautility and on some others they integer> up to 20. Since there is no error penalty associated with guessing a number that is too high, users without access to the system conformance document (to see what limits are actually in place) could use the historical 1 to 20 range or attempt to use very large numbers if the job should be truly low priority. Defautility and on some others they integer> up to 20. Since there is no error penalty associated with guessing a number that is too high, users without access to the system conformance document (to see what limits are actually in place) could use the historical 1 to 20 range or attempt to use very large numbers if the job should be truly low priority. Defautility and on some others they are exactly equivalent. Defautility and on some others they are exactly equivalent. Defautility and on some others they are exactly equivalent. Defautility and on some in line in they are exactly equivalent. Defautility and on some	25739	127 The utility specified by <i>utility</i> could not be found.
The only guaranteed portable uses of this utility are: nice utility Run utility with the default lower nice value. nice —n <positive integer=""> utility Run utility with a lower nice value. On some implementations they have no discernible effect on the invoked utility and on some others they are exactly equivalent. Historical systems have frequently supported the <positive integer=""> up to 20. Since there is no error penalty associated with guessing a number that is too high, users without access to the system conformance document (to see what limits are actually in place) could use the historical 1 to 20 range or attempt to use very large numbers if the job should be truly low priority. The nice value of a process can be displayed using the command: ps —o nice The command, env, nice, nohup, time, 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</positive></positive>		·
Run utility with the default lower nice value. nice — n < positive integer> utility Run utility with a lower nice value. On some implementations they have no discernible effect on the invoked utility and on some others they are exactly equivalent. Historical systems have frequently supported the < positive integer> up to 20. Since there is no error penalty associated with guessing a number that is too high, users without access to the system conformance document (to see what limits are actually in place) could use the historical 1 to 20 range or attempt to use very large numbers if the job should be truly low priority. The nice value of a process can be displayed using the command: ps —o nice The command, env, nice, nohup, time, 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		
Run utility with a lower nice value. On some implementations they have no discernible effect on the invoked utility and on some others they are exactly equivalent. Historical systems have frequently supported the <positive integer=""> up to 20. Since there is no error penalty associated with guessing a number that is too high, users without access to the system conformance document (to see what limits are actually in place) could use the historical 1 to 20 range or attempt to use very large numbers if the job should be truly low priority. The nice value of a process can be displayed using the command: ps -o nice The command, env, nice, nohup, time, 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</positive>		·
others they are exactly equivalent. Historical systems have frequently supported the <i><positive integer=""></positive></i> up to 20. Since there is no error penalty associated with guessing a number that is too high, users without access to the system conformance document (to see what limits are actually in place) could use the historical 1 to 20 range or attempt to use very large numbers if the job should be truly low priority. The nice value of a process can be displayed using the command: ps -o nice The command, env, nice, nohup, time, 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		
error penalty associated with guessing a number that is too high, users without access to the system conformance document (to see what limits are actually in place) could use the historical 1 to 20 range or attempt to use very large numbers if the job should be truly low priority. The nice value of a process can be displayed using the command: ps -o nice The command, env, nice, nohup, time, 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		·
25755 ps -o nice The command, env, nice, nohup, time, 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	25751 25752	error penalty associated with guessing a number that is too high, users without access to the system conformance document (to see what limits are actually in place) could use the historical
The <i>command, env, nice, nohup, time,</i> and <i>xargs</i> 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	25754	The nice value of a process can be displayed using the command:
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	25755	ps -o nice
1 1	25757 25758 25759 25760 25761 25762	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

nice Utilities

exec the utility fail with [ENOENT], and uses 126 when any attempt to *exec* the utility fails for any other reason.

25766 EXAMPLES

25767 None.

25768 RATIONALE

25773

25774 25775

2577625777

25778

25779

25782

25783

25784

25785

25786

2578725788

2578925790

25791 25792

25793

25794

25795

25796

25797

25798

25799

Due to the text about the limits of the nice value being implementation-defined, *nice* is not actually required to change the nice value of the executed command; the limits could be zero differences from the system default, although the implementor is required to document this fact in the conformance document.

The 4.3 BSD version of *nice* does not check whether *increment* is a valid decimal integer. The command *nice* $-\mathbf{x}$ *utility*, for example, would be treated the same as the command *nice* $-\mathbf{1}$ *utility*. If the user does not have appropriate privileges, this results in a "permission denied" error. This is considered a bug.

When a user without appropriate privileges gives a negative *increment*, System V treats it like the command *nice* **–0** *utility*, while 4.3 BSD writes a "permission denied" message and does not run the utility. Neither was considered clearly superior, so the behavior was left unspecified.

25780 The C shell has a built-in version of *nice* that has a different interface from the one described in this volume of IEEE Std 1003.1-2001.

The term "utility" is used, rather than "command", to highlight the fact that shell compound commands, pipelines, and so on, cannot be used. Special built-ins also cannot be used. However, "utility" includes user application programs and shell scripts, not just utilities defined in this volume of IEEE Std 1003.1-2001.

Historical implementations of *nice* provide a nice value range of 40 or 41 discrete steps, with the default nice value being the midpoint of that range. By default, they lower the nice value of the executed utility by 10.

Some historical documentation states that the *increment* value must be within a fixed range. This is misleading; the valid *increment* values on any invocation are determined by the current process nice value, which is not always the default.

The definition of nice value is not intended to suggest that all processes in a system have priorities that are comparable. Scheduling policy extensions such as the realtime priorities in the System Interfaces volume of IEEE Std 1003.1-2001 make the notion of a single underlying priority for all scheduling policies problematic. Some implementations may implement the *nice*-related features to affect all processes on the system, others to affect just the general time-sharing activities implied by this volume of IEEE Std 1003.1-2001, and others may have no effect at all. Because of the use of "implementation-defined" in *nice* and *renice*, a wide range of implementation strategies are possible.

25800 FUTURE DIRECTIONS

25801 None.

25802 SEE ALSO

25803 Chapter 2 (on page 29), renice, the System Interfaces volume of IEEE Std 1003.1-2001, nice()

25804 CHANGE HISTORY

First released in Issue 4.

Utilities nice

25806 **Issue 6**

25807 This utility is marked as part of the User Portability Utilities option.

25808 The obsolescent SYNOPSIS is removed.

nl **Utilities**

25809 NAME

25816

25817 25818

25819

25820

25821

25822

25823

25824

25830

25832

25833

25839

25843 25844

25845

25810 nl — line numbering filter

25811 SYNOPSIS

nl [-p] [-b type] [-d delim] [-f type] [-h type] [-i incr] [-l num] [-n format] 25812 XSI [-s sep] [-v startnum] [-w width] [file] 25813 25814

25815 **DESCRIPTION**

The *nl* utility shall read lines from the named *file* or the standard input if no *file* is named and shall reproduce the lines to standard output. Lines shall be numbered on the left. Additional functionality may be provided in accordance with the command options in effect.

The *nl* utility views the text it reads in terms of logical pages. Line numbering shall be reset at the start of each logical page. A logical page consists of a header, a body, and a footer section. Empty sections are valid. Different line numbering options are independently available for header, body, and footer (for example, no numbering of header and footer lines while numbering blank lines only in the body).

The starts of logical page sections shall be signaled by input lines containing nothing but the following delimiter characters:

Line	Start of
\:\:\:	Header
\:\:	Body
\:	Footer

Unless otherwise specified, *nl* shall assume the text being read is in a single logical page body.

25831 OPTIONS

The *nl* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines. Only one file can be named.

The following options shall be supported: 25834

−**b** type Specify which logical page body lines shall be numbered. Recognized types and 25835 25836 their meaning are:

Number all lines. 25837

-d delim

Number only non-empty lines. t 25838

No line numbering.

Number only lines that contain the basic regular expression specified in **p**string 25840 25841

The default *type* for logical page body shall be **t** (text lines numbered). 25842

characters. If only one character is entered, the second character shall remain the default character ':'. 25846 Specify the same as **b** type except for footer. The default for logical page footer 25847 -f type shall be **n** (no lines numbered). 25848 Specify the same as **b** type except for header. The default type for logical page 25849 -h type 25850 header shall be **n** (no lines numbered).

Specify the delimiter characters that indicate the start of a logical page section.

These can be changed from the default characters "\:" to two user-specified

Utilities nl

25851 25852	− i incr	Specify the increment value used to number logical page lines. The default shall be 1.
25853 25854 25855	–l num	Specify the number of blank lines to be considered as one. For example, $-\mathbf{l}\ 2$ results in only the second adjacent blank line being numbered (if the appropriate $-\mathbf{h}\ \mathbf{a}$, $-\mathbf{b}\ \mathbf{a}$, or $-\mathbf{f}\ \mathbf{a}$ option is set). The default shall be 1.
25856 25857 25858	− n format	Specify the line numbering format. Recognized values are: In , left justified, leading zeros suppressed; rn , right justified, leading zeros suppressed; rz , right justified, leading zeros kept. The default <i>format</i> shall be rn (right justified).
25859	- p	Specify that numbering should not be restarted at logical page delimiters.
25860 25861	−s <i>sep</i>	Specify the characters used in separating the line number and the corresponding text line. The default <i>sep</i> shall be a <tab>.</tab>
25862	–v startnum	Specify the initial value used to number logical page lines. The default shall be 1.
25863 25864	− w width	Specify the number of characters to be used for the line number. The default <i>width</i> shall be 6.
25865 OPERA 25866		ng operand shall be supported:
25867	file	A pathname of a text file to be line-numbered.
25868 STDIN 25869	The standard	d input is a text file that is used if no <i>file</i> operand is given.
25870 INPUT	FILES	
95071	The input fil	a named by the file approach is a tayt file
25871 25872 FNVIR	•	e named by the <i>file</i> operand is a text file.
	ONMENT VA	•
25872 ENVIR	ONMENT VA	ARIABLES
25872 ENVIR 25873 25874 25875 25876	ONMENT VA	ARIABLES ag environment variables shall affect the execution of <i>nl</i> : Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables
25872 ENVIR 25873 25874 25875 25876 25877 25878	ONMENT VA The followin	ARIABLES ag environment variables shall affect the execution of <i>nl</i> : Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables.
25872 ENVIR 25873 25874 25875 25876 25877 25878 25879 25880 25881	ONMENT VA The followin LANG LC_ALL	ARIABLES ag environment variables shall affect the execution of <i>nl</i> : Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. E Determine the locale for the behavior of ranges, equivalence classes, and multi-
25872 ENVIR 25873 25874 25875 25876 25877 25878 25879 25880 25881 25882 25883 25884 25885 25886	ONMENT VA The followin LANG LC_ALL LC_COLLAT	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. E Determine the locale for the behavior of ranges, equivalence classes, and multicharacter collating elements within regular expressions. 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), the behavior of character classes within regular expressions, and for deciding which characters are in character class graph (for the -b t, -f t, and -h t options).
25872 ENVIR 25873 25874 25875 25876 25877 25878 25879 25880 25881 25882 25883 25884 25885 25886 25887 25888	ONMENT VA The followin LANG LC_ALL LC_COLLAT LC_CTYPE	RIABLES In genvironment variables shall affect the execution of nl: Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. E Determine the locale for the behavior of ranges, equivalence classes, and multicharacter collating elements within regular expressions. 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), the behavior of character classes within regular expressions, and for deciding which characters are in character class graph (for the -b t, -f t, and -h t options). GES Determine the locale that should be used to affect the format and contents of

nl Utilities

25892 ASYNCHRONOUS EVENTS 25893 Default. 25894 **STDOUT** The standard output shall be a text file in the following format: 25895 "%s%s%s", <line number>, <separator>, <input line> 25896 25897 where *< line number>* is one of the following numeric formats: %6d When the **rn** format is used (the default; see $-\mathbf{n}$). 25898 When the **rz** format is used. %06d 25899 %-6d When the **ln** format is used. 25900 25901 <empty> When line numbers are suppressed for a portion of the page; the *<separator>* is also 25902 suppressed. 25903 In the preceding list, the number 6 is the default width; the –w option can change this value. 25904 STDERR 25905 The standard error shall be used only for diagnostic messages. 25906 OUTPUT FILES 25907 None. 25908 EXTENDED DESCRIPTION 25909 None. 25910 EXIT STATUS 25911 The following exit values shall be returned: Successful completion. 25912 25913 >0 An error occurred. 25914 CONSEQUENCES OF ERRORS 25915 Default. 25916 APPLICATION USAGE 25917 In using the -d delim option, care should be taken to escape characters that have special meaning 25918 to the command interpreter. 25919 EXAMPLES The command: 25920 25921 nl -v 10 -i 10 -d \!+ file1 numbers file1 starting at line number 10 with an increment of 10. The logical page delimiter is 25922 "!+". Note that the '!' has to be escaped when using csh as a command interpreter because of 25923 25924 its history substitution syntax. For ksh and sh the escape is not necessary, but does not do any 25925 harm. 25926 RATIONALE None. 25927 25928 FUTURE DIRECTIONS

25929

None.

Utilities nl

25930 SEE ALSO
25931 pr

25932 CHANGE HISTORY
25933 First released in Issue 2.

25934 Issue 5
25935 The option [-f type] is added to the SYNOPSIS. The option descriptions are presented in alphabetic order. The description of -bt is changed to "Number only non-empty lines".

25937 Issue 6
25938 The obsolescent behavior allowing the options to be intermingled with the optional file operand is removed.

nm Utilities

```
25940 NAME
25941
               nm — write the name list of an object file (DEVELOPMENT)
25942 SYNOPSIS
25943 UP SD XSI nm [-APv][-efox][-g] -u] [-t format] file...
25944
25945 DESCRIPTION
               This utility shall be provided on systems that support both the User Portability Utilities option
25946
               and the Software Development Utilities option. On other systems it is optional. Certain options
25947
               are only available on XSI-conformant systems.
25948
               The nm utility shall display symbolic information appearing in the object file, executable file, or
25949
               object-file library named by file. If no symbolic information is available for a valid input file, the
25950
               nm utility shall report that fact, but not consider it an error condition.
25951
               The default base used when numeric values are written is unspecified. On XSI-conformant
25952 XSI
25953
               systems, it shall be decimal.
25954 OPTIONS
               The nm utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
25955
               12.2, Utility Syntax Guidelines.
25956
               The following options shall be supported:
25957
25958
               -\mathbf{A}
                             Write the full pathname or library name of an object on each line.
               -е
                             Write only external (global) and static symbol information.
25959 XSI
                             Produce full output. Write redundant symbols (.text, .data, and .bss), normally
               -\mathbf{f}
25960 XSI
                             suppressed.
25961
                             Write only external (global) symbol information.
25962
               -g
                             Write numeric values in octal (equivalent to -\mathbf{t} \mathbf{o}).
25963 XSI
               -o
               -\mathbf{P}
25964
                             Write information in a portable output format, as specified in the STDOUT section.
               -t format
                             Write each numeric value in the specified format. The format shall be dependent
25965
25966
                             on the single character used as the format option-argument:
                             d
                                  The offset is written in decimal (default).
25967 XSI
                                  The offset is written in octal.
25968
                             0
                                  The offset is written in hexadecimal.
25969
                             Write only undefined symbols.
25970
               -u
25971
                             Sort output by value instead of alphabetically.
               -v
                             Write numeric values in hexadecimal (equivalent to -\mathbf{t} \mathbf{x}).
25972 XSI
               -\mathbf{x}
25973 OPERANDS
               The following operand shall be supported:
25974
               file
                             A pathname of an object file, executable file, or object-file library.
25975
25976 STDIN
               See the INPUT FILES section.
25977
```

Utilities nm

25978 INPUT FILES

25983

The input file shall be an object file, an object-file library whose format is the same as those produced by the *ar* utility for link editing, or an executable file. The *nm* utility may accept additional implementation-defined object library formats for the input file.

25982 ENVIRONMENT VARIABLES

25984	LANG	Provide a default value for the internationalization variables that are unset or null.
25985		(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
25986		Internationalization Variables for the precedence of internationalization variables
25987		used to determine the values of locale categories.)
	- ~	

The following environment variables shall affect the execution of *nm*:

25988 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

25990 LC_COLLATE

Determine the locale for character collation information for the symbol-name and symbol-value collation sequences.

25993 *LC_CTYPE* 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).

25996 LC_MESSAGES

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

25999 XSI NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

26000 ASYNCHRONOUS EVENTS

26001 Default.

26002 STDOUT

26006

If symbolic information is present in the input files, then for each file or for each member of an archive, the *nm* utility shall write the following information to standard output. By default, the format is unspecified, but the output shall be sorted alphabetically by symbol name:

- Library or object name, if –A is specified
- Symbol name
- Symbol type, which shall either be one of the following single characters or an implementation-defined type represented by a single character:
- 26010 A Global absolute symbol.
- 26011 a Local absolute symbol.
- 26012 B Global "bss" (that is, uninitialized data space) symbol.
- 26013 b Local bss symbol.
- 26014 D Global data symbol.
- 26015 d Local data symbol.
- 26016 T Global text symbol.
- 26017 t Local text symbol.
- U Undefined symbol.

nm Utilities

- 26019
- · Value of the symbol
- 26020

26033

26034

26036

26037

26039 26040

2604126042

26044 26045 • The size associated with the symbol, if applicable

This information may be supplemented by additional information specific to the implementation.

26023 If the $-\mathbf{P}$ option is specified, the previous information shall be displayed using the following portable format. The three versions differ depending on whether $-\mathbf{t}$ \mathbf{d} , $-\mathbf{t}$ \mathbf{o} , or $-\mathbf{t}$ \mathbf{x} was specified, respectively:

```
26026 "%s%s %s %d %d\n", <library/object name>, <name>, <type>,
26027 < value>, <size>

26028 "%s%s %s %o %o\n", <library/object name>, <name>, <type>,
26029 < value>, <size>

26030 "%s%s %s %x %x\n", <library/object name>, <name>, <type>,
26031 < value>, <size>
```

26032 where *< library/object name>* shall be formatted as follows:

- If –A is not specified, < library/object name> shall be an empty string.
- If –A is specified and the corresponding *file* operand does not name a library:

```
26035 "%s: ", <file>
```

• If **–A** is specified and the corresponding *file* operand names a library. In this case, *<object file>* shall name the object file in the library containing the symbol being described:

```
26038 "%s[%s]: ", <file>, <object file>
```

If –A is not specified, then if more than one *file* operand is specified or if only one *file* operand is specified and it names a library, *nm* shall write a line identifying the object containing the following symbols before the lines containing those symbols, in the form:

• If the corresponding *file* operand does not name a library:

```
26043 "%s:\n", <file>
```

• If the corresponding *file* operand names a library; in this case, *<object file>* shall be the name of the file in the library containing the following symbols:

```
26046 "%s[%s]:\n", <file>, <object file>
```

26047 If $-\mathbf{P}$ is specified, but $-\mathbf{t}$ is not, the format shall be as if $-\mathbf{t} \times \mathbf{x}$ had been specified.

26048 STDERR

The standard error shall be used only for diagnostic messages.

26050 OUTPUT FILES

26051 None.

26052 EXTENDED DESCRIPTION

26053 None.

26054 EXIT STATUS

26055 The following exit values shall be returned:

26056 0 Successful completion.

26057 >0 An error occurred.

Utilities nm

26058 CONSEQUENCES OF ERRORS

26059 Default.

26060 APPLICATION USAGE

Mechanisms for dynamic linking make this utility less meaningful when applied to an executable file because a dynamically linked executable may omit numerous library routines that would be found in a statically linked executable.

26064 EXAMPLES

26065 None.

26066 RATIONALE

26071

26072

26073

26074

26078

26079

26080

26081

26082

26083 26084

26085

26086

26087

26088

26089

26090 26091

26092

Historical implementations of *nm* have used different bases for numeric output and supplied different default types of symbols that were reported. The -t *format* option, similar to that used in *od* and *strings*, can be used to specify the numeric base; -g and -u can be used to restrict the amount of output or the types of symbols included in the output.

The compromise of using $-\mathbf{t}$ *format versus* using $-\mathbf{d}$, $-\mathbf{o}$, and other similar options was necessary because of differences in the meaning of $-\mathbf{o}$ between implementations. The $-\mathbf{o}$ option from BSD has been provided here as $-\mathbf{A}$ to avoid confusion with the $-\mathbf{o}$ from System V (which has been provided here as $-\mathbf{t}$ and as $-\mathbf{o}$ on XSI-conformant systems).

The option list was significantly reduced from that provided by historical implementations.

The *nm* description is a subset of both the System V and BSD *nm* utilities with no specified default output.

It was recognized that mechanisms for dynamic linking make this utility less meaningful when applied to an executable file (because a dynamically linked executable file may omit numerous library routines that would be found in a statically linked executable file), but the value of *nm* during software development was judged to outweigh other limitations.

The default output format of nm is not specified because of differences in historical implementations. The $-\mathbf{P}$ option was added to allow some type of portable output format. After a comparison of the different formats used in SunOS, BSD, SVR3, and SVR4, it was decided to create one that did not match the current format of any of these four systems. The format devised is easy to parse by humans, easy to parse in shell scripts, and does not need to vary depending on locale (because no English descriptions are included). All of the systems currently have the information available to use this format.

The format given in *nm* STDOUT uses spaces between the fields, which may be any number of

blank>s required to align the columns. The single-character types were selected to match historical practice, and the requirement that implementation additions also be single characters made parsing the information easier for shell scripts.

26093 FUTURE DIRECTIONS

26094 None.

26095 SEE ALSO

26096 ar, c99

26097 CHANGE HISTORY

26098 First released in Issue 2.

26099 Issue 6

This utility is marked as supported when both the User Portability Utilities option and the Software Development Utilities option are supported.

nohup Utilities

26102	NAME		
26103		nohup — inv	oke a utility immune to hangups
26104	SYNOP		
26105		nohup util	lity [argument]
26106 26107 26108 26109	DESCR.	The <i>nohup</i> ut	tility shall invoke the utility named by the <i>utility</i> operand with arguments supplied <i>ent</i> operands. At the time the named <i>utility</i> is invoked, the SIGHUP signal shall be ored.
26110 26111 26112 26113 26114 26115		shall be appe be created or nohup.out in created or o	rd output is a terminal, all output written by the named <i>utility</i> to its standard output ended to the end of the file nohup.out in the current directory. If nohup.out cannot or opened for appending, the output shall be appended to the end of the file in the directory specified by the <i>HOME</i> environment variable. If neither file can be spened for appending, <i>utility</i> shall not be invoked. If a file is created, the file's points shall be set to S_IRUSR S_IWUSR.
26116 26117			ard error is a terminal, all output written by the named <i>utility</i> to its standard error rected to the same file descriptor as the standard output.
	OPTIO		
26119		None.	
	OPERA		a anamanda ahall ba ayunnantad.
26121			g operands shall be supported:
26122 26123		utility	The name of a utility that is to be invoked. If the <i>utility</i> operand names any of the special built-in utilities in Section 2.14 (on page 64), the results are undefined.
26124 26125		argument	Any string to be supplied as an argument when invoking the utility named by the <i>utility</i> operand.
26126 26127	STDIN	Not used.	
26128 26129	INPUT	FILES None.	
26130	ENVIR	ONMENT VA	RIABLES
26131		The followin	g environment variables shall affect the execution of <i>nohup</i> :
26132 26133 26134		НОМЕ	Determine the pathname of the user's home directory: if the output file nohup.out cannot be created in the current directory, the <i>nohup</i> utility shall use the directory named by <i>HOME</i> to create the file.
26135 26136 26137 26138		LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
26139 26140		LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
26141 26142 26143		LC_CTYPE	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).
26144 26145		LC_MESSAG	GES Determine the locale that should be used to affect the format and contents of

nohup **Utilities**

26146		diagnostic messages written to standard error.
26147 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
26148 26149 26150	PATH	Determine the search path that is used to locate the utility to be invoked. See the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables.

26151 ASYNCHRONOUS EVENTS

26152 The nohup utility shall take the standard action for all signals except that SIGHUP shall be 26153 ignored.

26154 STDOUT

If the standard output is not a terminal, the standard output of nohup shall be the standard 26155 26156 output generated by the execution of the utility specified by the operands. Otherwise, nothing 26157 shall be written to the standard output.

26158 STDERR

If the standard output is a terminal, a message shall be written to the standard error, indicating 26159 the name of the file to which the output is being appended. The name of the file shall be either 26160 26161 nohup.out or \$HOME/nohup.out.

26162 OUTPUT FILES

If the standard output is a terminal, all output written by the named utility to the standard 26163 output and standard error is appended to the file nohup.out, which is created if it does not 26164 already exist. 26165

26166 EXTENDED DESCRIPTION

None. 26167

26168 EXIT STATUS

26169 The following exit values shall be returned:

26170 126 The utility specified by *utility* was found but could not be invoked.

26171 127 An error occurred in the *nohup* utility or the utility specified by *utility* could not be

found. 26172

26173 Otherwise, the exit status of *nohup* shall be that of the utility specified by the *utility* operand.

26174 CONSEQUENCES OF ERRORS

Default. 26175

26177

26178 26179

26180

26181 26182

26183

26184

26185

26186

26176 APPLICATION USAGE

The command, env, nice, nohup, time, 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.

26187 EXAMPLES

26188 It is frequently desirable to apply *nohup* to pipelines or lists of commands. This can be done by 26189 placing pipelines and command lists in a single file; this file can then be invoked as a utility, and the *nohup* applies to everything in the file. 26190

nohup Utilities

26191	Alternatively, the following command can be used to apply <i>nohup</i> to a complex command:
26192	nohup sh -c 'complex-command-line'
26193 RATIO 26194 26195	NALE The 4.3 BSD version ignores SIGTERM and SIGHUP, and if ./nohup.out cannot be used, it fails instead of trying to use \$HOME/nohup.out.
26196 26197	The <i>csh</i> utility has a built-in version of <i>nohup</i> that acts differently from the <i>nohup</i> defined in this volume of IEEE Std 1003.1-2001.
26198 26199 26200	The term <i>utility</i> is used, rather than <i>command</i> , to highlight the fact that shell compound commands, pipelines, special built-ins, and so on, cannot be used directly. However, <i>utility</i> includes user application programs and shell scripts, not just the standard utilities.
26201 26202	Historical versions of the <i>nohup</i> utility use default file creation semantics. Some more recent versions use the permissions specified here as an added security precaution.
26203 26204 26205 26206	Some historical implementations ignore SIGQUIT in addition to SIGHUP; others ignore SIGTERM. An early proposal allowed, but did not require, SIGQUIT to be ignored. Several reviewers objected that <i>nohup</i> should only modify the handling of SIGHUP as required by this volume of IEEE Std 1003.1-2001.
26207 FUTUI	RE DIRECTIONS
26208	None.
26209 SEE Al 26210	LSO Chapter 2 (on page 29), <i>sh</i> , the System Interfaces volume of IEEE Std 1003.1-2001, <i>signal</i> ()
26211 CHAN 26212	GE HISTORY First released in Issue 2.

Utilities od

26213 **NAME** 26214 od — dump files in various formats 26215 SYNOPSIS 26216 od [-v] [-A address base] [-j skip] [-N count] [-t type string]... 26217 [file...] 26218 XSI od [-bcdosx][file] [[+]offset[.][b]] 26219 26220 DESCRIPTION The od utility shall write the contents of its input files to standard output in a user-specified 26221 26222 format. 26223 OPTIONS The od utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, 26224 Utility Syntax Guidelines, except that the order of presentation of the -t options and the 26225 XSI 26226 **-bcdosx** options is significant. The following options shall be supported: 26227 -A address base 26228 Specify the input offset base. See the EXTENDED DESCRIPTION section. The 26229 application shall ensure that the address_base option-argument is a character. The 26230 characters 'd', 'o', and 'x' specify that the offset base shall be written in 26231 decimal, octal, or hexadecimal, respectively. The character 'n' specifies that the 26232 offset shall not be written. 26233 −b Interpret bytes in octal. This shall be equivalent to **-t o1**. 26234 XSI Interpret bytes as characters specified by the current setting of the LC CTYPE 26235 XSI -с category. Certain non-graphic characters appear as C escapes: "NUL=\0", 26236 "BS=\b", "FF=\f", "NL=\n", "CR=\r", "HT=\t"; others appear as 3-digit octal 26237 numbers. 26238 $-\mathbf{d}$ Interpret words (two-byte units) in unsigned decimal. This shall be equivalent to 26239 XSI 26240 −t u2. Jump over *skip* bytes from the beginning of the input. The *od* utility shall read or 26241 −**j** skip seek past the first skip bytes in the concatenated input files. If the combined input 26242 26243 is not at least skip bytes long, the od utility shall write a diagnostic message to standard error and exit with a non-zero exit status. 26244 By default, the skip option-argument shall be interpreted as a decimal number. 26245 With a leading 0x or 0X, the offset shall be interpreted as a hexadecimal number; 26246 otherwise, with a leading '0', the offset shall be interpreted as an octal number. 26247 Appending the character 'b', 'k', or 'm' to offset shall cause it to be interpreted 26248 as a multiple of 512, 1024, or 1048 576 bytes, respectively. If the skip number is 26249 26250 hexadecimal, any appended 'b' shall be considered to be the final hexadecimal digit. 26251 -N count 26252 Format no more than *count* bytes of input. By default, *count* shall be interpreted as a decimal number. With a leading 0x or 0X, count shall be interpreted as a 26253 hexadecimal number; otherwise, with a leading '0', it shall be interpreted as an 26254 octal number. If count bytes of input (after successfully skipping, if -i skip is 26255 specified) are not available, it shall not be considered an error; the od utility shall 26256 26257 format the input that is available.

od Utilities

26258 XSI	-o	Interpret $words$ (two-byte units) in octal. This shall be equivalent to $-\mathbf{t}$ o2 .
26259 XSI 26260	-s	Interpret $words$ (two-byte units) in signed decimal. This shall be equivalent to $-\mathbf{t} \ \mathbf{d2}$.
26261 26262 26263 26264 26265 26266 26267 26268 26269 26270 26271 26272 26273 26274 26275 26276	-t type_string	Specify one or more output types. See the EXTENDED DESCRIPTION section. The application shall ensure that the $type_string$ option-argument is a string specifying the types to be used when writing the input data. The string shall consist of the type specification characters a, c, d, f, o, u, and x, specifying named character, character, signed decimal, floating point, octal, unsigned decimal, and hexadecimal, respectively. The type specification characters d, f, o, u, and x can be followed by an optional unsigned decimal integer that specifies the number of bytes to be transformed by each instance of the output type. The type specification character f can be followed by an optional F, D, or L indicating that the conversion should be applied to an item of type float , double , or long double , respectively. The type specification characters d, o, u, and x can be followed by an optional C, S, I, or L indicating that the conversion should be applied to an item of type char , short , int , or long , respectively. Multiple types can be concatenated within the same $type_string$ and multiple $-t$ options can be specified. Output lines shall be written for each type specified in the order in which the type specification
26277 26278 26279 26280 26281	- v	characters are specified. Write all input data. Without the $-\mathbf{v}$ option, any number of groups of output lines, which would be identical to the immediately preceding group of output lines (except for the byte offsets), shall be replaced with a line containing only an asterisk ('*').
26282 XSI	-x	Interpret $words$ (two-byte units) in hexadecimal. This shall be equivalent to $-\mathbf{t} \mathbf{x} 2$.
26283 XSI 26284		es can be specified by using multiple – bcdostx options. Output lines are written for ecified in the order in which the types are specified.
26285 OPERA 26286		ng operands shall be supported:
26287 26288	file	A pathname of a file to be read. If no <i>file</i> operands are specified, the standard input shall be used.
26289 26290 26291 26292 XSI 26293 26294		If there are no more than two operands, none of the $-A$, $-j$, $-N$, or $-t$ options is specified, and either of the following is true: the first character of the last operand is a plus sign $('+')$, or there are two operands and the first character of the last operand is numeric; the last operand shall be interpreted as an offset operand on XSI-conformant systems. Under these conditions, the results are unspecified on systems that are not XSI-conformant systems.
26295 XSI 26296 26297 26298	[+] <i>offset</i> [.][b]	The <i>offset</i> operand specifies the offset in the file where dumping is to commence. This operand is normally interpreted as octal bytes. If '.' is appended, the offset shall be interpreted in decimal. If 'b' is appended, the offset shall be interpreted in units of 512 bytes.
26299 STDIN 26300 26301	The standar section.	d input shall be used only if no file operands are specified. See the INPUT FILES

Utilities od

26302 INPUT FILES

26303 The input files can be any file type.

26304 ENVIRONMENT VARIABLES

26305 The following environment variables shall affect the execution of *od*:

26306 LANG Provide a default value for the internationalization variables that are unset or null.
26307 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
26308 Internationalization Variables for the precedence of internationalization variables
26309 used to determine the values of locale categories.)

26310 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

26312 *LC_CTYPE* 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).

26315 LC_MESSAGES

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

26318 LC_NUMERIC

Determine the locale for selecting the radix character used when writing floating-point formatted output.

26321 XSI NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

26322 ASYNCHRONOUS EVENTS

26323 Default.

26324 STDOUT

See the EXTENDED DESCRIPTION section.

26326 STDERR

26336

26337 26338

26339

26340

26341

26342

26343

26344 26345

26346

26327 The standard error shall be used only for diagnostic messages.

26328 OUTPUT FILES

26329 None.

26330 EXTENDED DESCRIPTION

The *od* utility shall copy sequentially each input file to standard output, transforming the input data according to the output types specified by the **-t** option or the **-bcdosx** options. If no output type is specified, the default output shall be as if **-t oS** had been specified.

The number of bytes transformed by the output type specifier c may be variable depending on the LC_CTYPE category.

The default number of bytes transformed by output type specifiers d, f, o, u, and x corresponds to the various C-language types as follows. If the c99 compiler is present on the system, these specifiers shall correspond to the sizes used by default in that compiler. Otherwise, these sizes may vary among systems that conform to IEEE Std 1003.1-2001.

• For the type specifier characters d, o, u, and x, the default number of bytes shall correspond to the size of the underlying implementation's basic integer type. For these specifier characters, the implementation shall support values of the optional number of bytes to be converted corresponding to the number of bytes in the C-language types **char**, **short**, **int**, and **long**. These numbers can also be specified by an application as the characters 'C', 'S', 'I', and 'L', respectively. The implementation shall also support the values 1, 2, 4, and 8, even if it provides no C-Language types of those sizes. The implementation shall support the

od **Utilities**

26347 26348 26349

decimal value corresponding to the C-language type long long. The byte order used when interpreting numeric values is implementation-defined, but shall correspond to the order in which a constant of the corresponding type is stored in memory on the system.

• For the type specifier character £, the default number of bytes shall correspond to the number of bytes in the underlying implementation's basic double precision floating-point data type. The implementation shall support values of the optional number of bytes to be converted corresponding to the number of bytes in the C-language types float, double, and long **double**. These numbers can also be specified by an application as the characters 'F', 'D', and 'L', respectively.

26355

The type specifier character a specifies that bytes shall be interpreted as named characters from the International Reference Version (IRV) of the ISO/IEC 646:1991 standard. Only the least significant seven bits of each byte shall be used for this type specification. Bytes with the values listed in the following table shall be written using the corresponding names for those characters.

26360

Table 4-12 Named Characters in *od*

26371

Value	Name	Value	Name	Value	Name	Value	Name
\000	nul	∖001	soh	∖002	stx	∖003	etx
\004	eot	∖005	enq	\006	ack	∖007	bel
\010	bs	\011	ht	\012	\mathbf{lf} or \mathbf{nl}^*	∖013	vt
\014	ff	\015	cr	\016	so	\017	si
∖020	dle	\021	dc1	\022	dc2	\023	dc3
\024	dc4	∖025	nak	\026	syn	\027	etb
\030	can	\031	em	\032	sub	∖033	esc
\034	fs	∖035	gs	∖036	rs	\037	us
∖040	sp	\177	del				

26372 Note: The "\012" value may be written either as **lf** or **nl**.

26386

26387

26388

26389

26390

26391

26392

The type specifier character c specifies that bytes shall be interpreted as characters specified by the current setting of the *LC_CTYPE* locale category. Characters listed in the table in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5, File Format Notation ('\\', '\a', '\b', '\f', '\n', '\r', '\t', '\v') shall be written as the corresponding escape sequences, except that backslash shall be written as a single backslash and a NUL shall be written as '\0'. Other non-printable characters shall be written as one three-digit octal number for each byte in the character. If the size of a byte on the system is greater than nine bits, the format used for nonprintable characters is implementation-defined. Printable multi-byte characters shall be written in the area corresponding to the first byte of the character; the two-character sequence "**" shall be written in the area corresponding to each remaining byte in the character, as an indication that the character is continued. When either the $-\mathbf{j}$ skip or $-\mathbf{N}$ count option is specified along with the c type specifier, and this results in an attempt to start or finish in the middle of a multi-byte character, the result is implementation-defined.

The input data shall be manipulated in blocks, where a block is defined as a multiple of the least common multiple of the number of bytes transformed by the specified output types. If the least common multiple is greater than 16, the results are unspecified. Each input block shall be written as transformed by each output type, one per written line, in the order that the output types were specified. If the input block size is larger than the number of bytes transformed by the output type, the output type shall sequentially transform the parts of the input block, and the output from each of the transformations shall be separated by one or more

blank>s.

Utilities od

If, as a result of the specification of the -N option or end-of-file being reached on the last input file, input data only partially satisfies an output type, the input shall be extended sufficiently with null bytes to write the last byte of the input.

Unless –A n is specified, the first output line produced for each input block shall be preceded by the input offset, cumulative across input files, of the next byte to be written. The format of the input offset is unspecified; however, it shall not contain any <blank>s, shall start at the first character of the output line, and shall be followed by one or more
blank>s. In addition, the offset of the byte following the last byte written shall be written after all the input data has been processed, but shall not be followed by any
blank>s.

If no –A option is specified, the input offset base is unspecified.

26403 EXIT STATUS

26402

26410

26411

26412

26413

26414

26416

26417

26419

26420

26404 The following exit values shall be returned:

26405 0 All input files were processed successfully.

26406 >0 An error occurred.

26407 CONSEQUENCES OF ERRORS

26408 Default.

26409 APPLICATION USAGE

XSI-conformant applications are warned not to use filenames starting with '+' or a first operand starting with a numeric character so that the old functionality can be maintained by implementations, unless they specify one of the -A, -j, or -N options. To guarantee that one of these filenames is always interpreted as a filename, an application could always specify the address base format with the -A option.

26415 EXAMPLES

If a file containing 128 bytes with decimal values zero to 127, in increasing order, is supplied as standard input to the command:

```
26418 od -A d -t a
```

on an implementation using an input block size of 16 bytes, the standard output, independent of the current locale setting, would be similar to:

```
ff
26421
              0000000 nul soh stx etx eot enq ack bel
                                                                   bs
                                                                        ht
                                                                             nl
                                                                                  vt
                                                                                             cr
                                                                                                  so
                                                                                                       si
              0000016 dle dc1 dc2 dc3 dc4 nak syn etb can
26422
                                                                                        fs
                                                                        em
                                                                           sub esc
                                                                                             gs
                                                                                                  rs
                                                                                                       us
                                     11
26423
              0000032
                         sp
                                !
                                          #
                                               $
                                                     응
                                                          &
                                                                    (
                                                                         )
                                                                                    +
                                                                                                         ?
26424
              0000048
                          0
                                1
                                     2
                                          3
                                               4
                                                     5
                                                          6
                                                               7
                                                                    8
                                                                         9
                                                                                    ;
                                                                                         <
26425
              0000064
                          @
                                Α
                                     В
                                          C
                                               D
                                                    Ε
                                                          F
                                                               G
                                                                    Η
                                                                         Ι
                                                                              J
                                                                                    K
                                                                                         L
                                                                                              M
                                                                                                   Ν
                                                                                                         0
                                                                                    [
                                                                                              ]
26426
              0800000
                          Ρ
                                Q
                                     R
                                          S
                                               Т
                                                    U
                                                          V
                                                               W
                                                                    Χ
                                                                         Υ
                                                                               Ζ
                                                                                         \
              0000096
                                     b
                                               d
                                                          f
                                                                    h
                                                                         i
                                                                               j
                                                                                    k
                                                                                         1
26427
                                а
                                          С
                                                     0
                                                               g
                                                                                              m
                                                                                                   n
                                                                                                         0
                                                                                    {
                                                                                              }
26428
              0000112
                                                                                                      del
                          р
                                q
                                                                         У
26429
              0000128
```

Note that this volume of IEEE Std 1003.1-2001 allows **nl** or **lf** to be used as the name for the ISO/IEC 646: 1991 standard IRV character with decimal value 10. The IRV names this character **lf** (line feed), but traditional implementations have referred to this character as newline (**nl**) and the POSIX locale character set symbolic name for the corresponding character is a <newline>.

26434 The command:

```
26435 od -A o -t o2x2x -N 18
```

on a system with 32-bit words and an implementation using an input block size of 16 bytes could write 18 bytes in approximately the following format:

od Utilities

```
26438
            0000000 032056 031440 041123 042040 052516 044530 020043 031464
26439
                                3320
                                        4253
                                                4420
                                                         554e
                                                                 4958
                                                                         2023
                        342e
                                                                                 3334
26440
                            342e3320
                                           42534420
                                                            554e4958
                                                                             20233334
            0000020 032472
26441
26442
26443
                            353a0000
            0000022
26444
26445
            The command:
26446
            od -A d -t f -t o4 -t x4 -N 24 -j 0x15
            on a system with 64-bit doubles (for example, IEEE Std 754-1985 double precision floating-point
26447
            format) would skip 21 bytes of input data and then write 24 bytes in approximately the
26448
26449
            following format:
            000000
                         1.00000000000000e+00
                                                     1.57350000000000e+01
26450
                      07774000000 00000000000 10013674121 35341217270
26451
                         3ff00000
                                       0000000
                                                     402f3851
26452
                                                                   eb851eb8
            0000016
                         1.40668230000000e+02
26453
                      10030312542 04370303230
26454
26455
                         40619562
                                       23e18698
```

26457 RATIONALE

The *od* utility went through several names in early proposals, including *hd*, *xd*, and most recently *hexdump*. There were several objections to all of these based on the following reasons:

- The *hd* and *xd* names conflicted with historical utilities that behaved differently.
- The *hexdump* description was much more complex than needed for a simple dump utility.
- The *od* utility has been available on all historical implementations and there was no need to create a new name for a utility so similar to the historical *od* utility.

The original reasons for not standardizing historical *od* were also fairly widespread. Those reasons are given below along with rationale explaining why the standard developers believe that this version does not suffer from the indicated problem:

- The BSD and System V versions of *od* have diverged, and the intersection of features provided by both does not meet the needs of the user community. In fact, the System V version only provides a mechanism for dumping octal bytes and **shorts**, signed and unsigned decimal **shorts**, hexadecimal **shorts**, and ASCII characters. BSD added the ability to dump **floats**, **doubles**, named ASCII characters, and octal, signed decimal, unsigned decimal, and hexadecimal **longs**. The version presented here provides more normalized forms for dumping bytes, **shorts**, **ints**, and **longs** in octal, signed decimal, unsigned decimal, and hexadecimal; **float**, **double**, and **long double**; and named ASCII as well as current locale characters.
- It would not be possible to come up with a compatible superset of the BSD and System V flags that met the requirements of the standard developers. The historical default *od* output is the specified default output of this utility. None of the option letters chosen for this version of *od* conflict with any of the options to historical versions of *od*.
- On systems with different sizes for short, int, and long, there was no way to ask for dumps of ints, even in the BSD version. Because of the way options are named, the name space could not be extended to solve these problems. This is why the -t option was added (with type specifiers more closely matched to the printf() formats used in the rest of this volume of

Utilities od

 IEEE Std 1003.1-2001) and the optional field sizes were added to the d, f, o, u, and x type specifiers. It is also one of the reasons why the historical practice was not mandated as a required obsolescent form of od. (Although the old versions of od are not listed as an obsolescent form, implementations are urged to continue to recognize the older forms for several more years.) The a, c, f, o, and x types match the meaning of the corresponding format characters in the historical implementations of od except for the default sizes of the fields converted. The d format is signed in this volume of IEEE Std 1003.1-2001 to match the printf() notation. (Historical versions of od used d as a synonym for u in this version. The System V implementation uses s for signed decimal; BSD uses i for signed decimal and s for null-terminated strings.) Other than d and u, all of the type specifiers match format characters in the historical BSD version of od.

The sizes of the C-language types char, short, int, long, float, double, and long double are used even though it is recognized that there may be zero or more than one compiler for the C language on an implementation and that they may use different sizes for some of these types. (For example, one compiler might use 2 bytes **shorts**, 2 bytes **ints**, and 4 bytes **longs**, while another compiler (or an option to the same compiler) uses 2 bytes **shorts**, 4 bytes **ints**, and 4 bytes longs.) Nonetheless, there has to be a basic size known by the implementation for these types, corresponding to the values reported by invocations of the getconf utility when called with system_var operands {UCHAR_MAX}, {USHORT_MAX}, {UINT_MAX}, and {ULONG_MAX} for the types char, short, int, and long, respectively. There are similar constants required by the ISO C standard, but not required by the System Interfaces volume of IEEE Std 1003.1-2001 or this volume of IEEE Std 1003.1-2001. They are {FLT_MANT_DIG}, {DBL_MANT_DIG}, and {LDBL_MANT_DIG} for the types float, double, and long double, respectively. If the optional c99 utility is provided by the implementation and used as specified by this volume of IEEE Std 1003.1-2001, these are the sizes that would be provided. If an option is used that specifies different sizes for these types, there is no guarantee that the *od* utility is able to interpret binary data output by such a program correctly.

This volume of IEEE Std 1003.1-2001 requires that the numeric values of these lengths be recognized by the od utility and that symbolic forms also be recognized. Thus, a conforming application can always look at an array of **unsigned long** data elements using od - t uL.

- The method of specifying the format for the address field based on specifying a starting offset in a file unnecessarily tied the two together. The -A option now specifies the address base and the -S option specifies a starting offset.
- It would be difficult to break the dependence on U.S. ASCII to achieve an internationalized utility. It does not seem to be any harder for *od* to dump characters in the current locale than it is for the *ed* or *sed* I commands. The c type specifier does this without difficulty and is completely compatible with the historical implementations of the c format character when the current locale uses a superset of the ISO/IEC 646: 1991 standard as a codeset. The a type specifier (from the BSD a format character) was left as a portable means to dump ASCII (or more correctly ISO/IEC 646: 1991 standard (IRV)) so that headers produced by *pax* could be deciphered even on systems that do not use the ISO/IEC 646: 1991 standard as a subset of their base codeset.

The use of "**" as an indication of continuation of a multi-byte character in c specifier output was chosen based on seeing an implementation that uses this method. The continuation bytes have to be marked in a way that is not ambiguous with another single-byte or multi-byte character.

An early proposal used -S and -n, respectively, for the -j and -N options eventually selected. These were changed to avoid conflicts with historical implementations.

od Utilities

26532 The original standard specified -t o2 as the default when no output type was given. This was 26533 changed to -t oS (the length of a short) to accommodate a supercomputer implementation that historically used 64 bits as its default (and that defined shorts as 64 bits). This change should not 26534 affect conforming applications. The requirement to support lengths of 1, 2, and 4 was added at 26535 the same time to address an historical implementation that had no two-byte data types in its C 26536 26537 compiler. The use of a basic integer data type is intended to allow the implementation to choose a word 26538 size commonly used by applications on that architecture. 26539 **26540 FUTURE DIRECTIONS** 26541 All option and operand interfaces marked as extensions may be withdrawn in a future version. **26542 SEE ALSO** c99, sed 26543 26544 CHANGE HISTORY 26545 First released in Issue 2. 26546 Issue 5 In the description of the -c option, the phrase "This is equivalent to -t c." is deleted. 26547 The FUTURE DIRECTIONS section is modified. 26548 26549 Issue 6 26550 The *od* utility is changed to remove the assumption that **short** was a two-byte entity, as per the revisions in the IEEE P1003.2b draft standard. 26551 The normative text is reworded to avoid use of the term "must" for application requirements. 26552

IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/33 is applied, correcting the examples

which used an undefined $-\mathbf{n}$ option, which should have been $-\mathbf{N}$.

26553

Utilities paste

26555 **NAME** 26556 paste — merge corresponding or subsequent lines of files 26557 SYNOPSIS paste [-s] [-d list] file... 26558 26559 **DESCRIPTION** The paste utility shall concatenate the corresponding lines of the given input files, and write the 26560 resulting lines to standard output. 26561 The default operation of paste shall concatenate the corresponding lines of the input files. The 26562 <newline> of every line except the line from the last input file shall be replaced with a <tab>. 26563 If an end-of-file condition is detected on one or more input files, but not all input files, paste shall 26564 behave as though empty lines were read from the files on which end-of-file was detected, unless 26565 26566 the $-\mathbf{s}$ option is specified. 26567 **OPTIONS** The paste utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 26568 12.2, Utility Syntax Guidelines. 26569 The following options shall be supported: 26570 -d list Unless a backslash character appears in *list*, each character in *list* is an element 26571 specifying a delimiter character. If a backslash character appears in *list*, the 26572 backslash character and one or more characters following it are an element 26573 26574 specifying a delimiter character as described below. These elements specify one or more delimiters to use, instead of the default <tab>, to replace the <newline> of 26575 the input lines. The elements in *list* shall be used circularly; that is, when the list is 26576 exhausted the first element from the list is reused. When the -s option is specified: 26577 The last <newline> in a file shall not be modified. 26578 • The delimiter shall be reset to the first element of *list* after each *file* operand is 26579 26580 processed. When the $-\mathbf{s}$ option is not specified: 26581 26582 • The <newline>s in the file specified by the last file operand shall not be modified. 26583 26584 • The delimiter shall be reset to the first element of list each time a line is processed from each file. 26585 26586 If a backslash character appears in *list*, it and the character following it shall be used to represent the following delimiter characters: 26587 \n < newline>. 26588 \t <tab>. 26589 \\ Backslash character. 26590 \0 Empty string (not a null character). If '\0' is immediately followed by the 26591 character 'x', the character 'X', or any character defined by the LC_CTYPE 26592 digit keyword (see the Base Definitions volume of IEEE Std 1003.1-2001, 26593 Chapter 7, Locale), the results are unspecified. 26594 If any other characters follow the backslash, the results are unspecified. 26595 Concatenate all of the lines of each separate input file in command line order. The 26596 -5

<newline> of every line except the last line in each input file shall be replaced with

paste Utilities

26598 the $\langle tab \rangle$, unless otherwise specified by the $-\mathbf{d}$ option. 26599 OPERANDS 26600 The following operand shall be supported: file A pathname of an input file. If '-' is specified for one or more of the files, the 26601 26602 standard input shall be used; the standard input shall be read one line at a time, circularly, for each instance of '-'. Implementations shall support pasting of at 26603 least 12 file operands. 26604 26605 STDIN The standard input shall be used only if one or more *file* operands is '-'. See the INPUT FILES 26606 section. 26607 26608 INPUT FILES The input files shall be text files, except that line lengths shall be unlimited. 26609 **26610 ENVIRONMENT VARIABLES** The following environment variables shall affect the execution of *paste*: 26611 26612 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 26613 Internationalization Variables for the precedence of internationalization variables 26614 used to determine the values of locale categories.) 26615 LC_ALL If set to a non-empty string value, override the values of all the other 26616 26617 internationalization variables. Determine the locale for the interpretation of sequences of bytes of text data as 26618 LC_CTYPE 26619 characters (for example, single-byte as opposed to multi-byte characters in arguments and input files). 26620 26621 LC_MESSAGES 26622 Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. 26623 Determine the location of message catalogs for the processing of *LC_MESSAGES*. NLSPATH 26624 XSI 26625 ASYNCHRONOUS EVENTS Default. 26626 26627 STDOUT 26628

Concatenated lines of input files shall be separated by the <tab> (or other characters under the 26629 control of the $-\mathbf{d}$ option) and terminated by a <newline>.

26630 STDERR

26631 The standard error shall be used only for diagnostic messages.

26632 OUTPUT FILES None. 26633

26634 EXTENDED DESCRIPTION

None. 26635

26636 EXIT STATUS

The following exit values shall be returned: 26637

26638 Successful completion.

26639 An error occurred. Utilities paste

26640 CONSEQUENCES OF ERRORS

If one or more input files cannot be opened when the **-s** option is not specified, a diagnostic message shall be written to standard error, but no output is written to standard output. If the **-s** option is specified, the *paste* utility shall provide the default behavior described in Section 1.11 (on page 20).

26645 APPLICATION USAGE

26648 26649

26650

26651

26652

26653

26654

26657

26658

26659 26660

When the escape sequences of the *list* option-argument are used in a shell script, they must be quoted; otherwise, the shell treats the $' \setminus '$ as a special character.

Conforming applications should only use the specific backslash escaped delimiters presented in this volume of IEEE Std 1003.1-2001. Historical implementations treat ' \x' ', where ' \x' ' is not in this list, as ' \x' ', but future implementations are free to expand this list to recognize other common escapes similar to those accepted by *printf* and other standard utilities.

Most of the standard utilities work on text files. The *cut* utility can be used to turn files with arbitrary line lengths into a set of text files containing the same data. The *paste* utility can be used to create (or recreate) files with arbitrary line lengths. For example, if *file* contains long lines:

```
26655 cut -b 1-500 -n file > file1
26656 cut -b 501- -n file > file2
```

creates **file1** (a text file) with lines no longer than 500 bytes (plus the <newline>) and **file2** that contains the remainder of the data from *file*. Note that **file2** is not a text file if there are lines in *file* that are longer than 500 + {LINE_MAX} bytes. The original file can be recreated from **file1** and **file2** using the command:

```
26661 paste -d "\0" file1 file2 > file
```

26662 The commands:

```
26663 paste -d "\0" ...
26664 paste -d "" ...
```

are not necessarily equivalent; the latter is not specified by this volume of IEEE Std 1003.1-2001 and may result in an error. The construct $' \setminus 0'$ is used to mean "no separator" because historical versions of *paste* did not follow the syntax guidelines, and the command:

```
26668 paste -d"" ...
```

could not be handled properly by *getopt()*.

26670 EXAMPLES

2667126672

Write out a directory in four columns:

```
ls | paste - - - -
```

26673 2. Combine pairs of lines from a file into single lines:

```
26674 paste -s -d "\t\n" file
```

26675 RATIONALE

26676 None.

26677 FUTURE DIRECTIONS

26678 None.

paste Utilities

26679 SEE ALSO

26680 Section 1.11 (on page 20), *cut*, *grep*, *pr*

26681 CHANGE HISTORY

First released in Issue 2.

26683 **Issue 6**

The normative text is reworded to avoid use of the term "must" for application requirements.

Utilities patch

```
26685 NAME
26686
               patch — apply changes to files
26687 SYNOPSIS
               patch [-blNR] [ -c | -e | -n] [-d dir] [-D define] [-i patchfile]
26688 UP
26689
                     [-o outfile] [-p num] [-r rejectfile] [file]
26690
26691 DESCRIPTION
               The patch utility shall read a source (patch) file containing any of the three forms of difference
26692
26693
               (diff) listings produced by the diff utility (normal, context, or in the style of ed) and apply those
26694
               differences to a file. By default, patch shall read from the standard input.
               The patch utility shall attempt to determine the type of the diff listing, unless overruled by a -c,
26695
               -\mathbf{e}, or -\mathbf{n} option.
26696
               If the patch file contains more than one patch, patch shall attempt to apply each of them as if they
26697
               came from separate patch files. (In this case, the application shall ensure that the name of the
26698
26699
               patch file is determinable for each diff listing.)
26700 OPTIONS
               The patch utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
26701
               12.2, Utility Syntax Guidelines.
26702
26703
               The following options shall be supported:
26704
               -b
                             Save a copy of the original contents of each modified file, before the differences are
                             applied, in a file of the same name with the suffix .orig appended to it. If the file
26705
                             already exists, it shall be overwritten; if multiple patches are applied to the same
26706
                             file, the .orig file shall be written only for the first patch. When the -o outfile option
26707
                             is also specified, file.orig shall not be created but, if outfile already exists,
26708
26709
                             outfile.orig shall be created.
                             Interpret the patch file as a context difference (the output of the utility diff when
26710
               -\mathbf{c}
26711
                             the -\mathbf{c} or -\mathbf{C} options are specified).
               −d dir
26712
                             Change the current directory to dir before processing as described in the
26713
                             EXTENDED DESCRIPTION section.
26714
               −D define
                             Mark changes with one of the following C preprocessor constructs:
26715
                             #ifdef define
26716
                             . . .
26717
                             #endif
                             #ifndef define
26718
26719
```

#endif

optionally combined with the C preprocessor construct **#else**. If the patched file is processed with the C preprocessor, where the macro *define* is defined, the output shall contain the changes from the patch file; otherwise, the output shall not contain the patches specified in the patch file.

26725 — \mathbf{e} Interpret the patch file as an ed script, rather than a diff script.

26720

26721

26722

26723

26724

26726 — i patchfile Read the patch information from the file named by the pathname patchfile, rather than the standard input.

patch Utilities

26728 26729 26730	- l	(The letter ell.) Cause any sequence of <blank>s in the difference script to match any sequence of <blank>s in the input file. Other characters shall be matched exactly.</blank></blank>	
26731	-n	Interpret the script as a normal difference.	
26732 26733	- N	Ignore patches where the differences have already been applied to the file; by default, already-applied patches shall be rejected.	
26734 26735 26736 26737 26738 26739	− o outfile	Instead of modifying the files (specified by the <i>file</i> operand or the difference listings) directly, write a copy of the file referenced by each patch, with the appropriate differences applied, to <i>outfile</i> . Multiple patches for a single file shall be applied to the intermediate versions of the file created by any previous patches, and shall result in multiple, concatenated versions of the file being written to <i>outfile</i> .	
26740 26741 26742 26743 26744 26745	− p num	For all pathnames in the patch file that indicate the names of files to be patched, delete <i>num</i> pathname components from the beginning of each pathname. If the pathname in the patch file is absolute, any leading slashes shall be considered the first component (that is, $-\mathbf{p} \ 1$ shall remove the leading slashes). Specifying $-\mathbf{p} \ 0$ shall cause the full pathname to be used. If $-\mathbf{p}$ is not specified, only the basename (the final pathname component) shall be used.	
26746 26747 26748 26749 26750 26751 26752 26753	-R	Reverse the sense of the patch script; that is, assume that the difference script was created from the new version to the old version. The $-\mathbf{R}$ option cannot be used with ed scripts. The $patch$ utility shall attempt to reverse each portion of the script before applying it. Rejected differences shall be saved in swapped format. If this option is not specified, and until a portion of the patch file is successfully applied, $patch$ attempts to apply each portion in its reversed sense as well as in its normal sense. If the attempt is successful, the user shall be prompted to determine whether the $-\mathbf{R}$ option should be set.	
26754 26755 26756	–r rejectfile	Override the default reject filename. In the default case, the reject file shall have the same name as the output file, with the suffix .rej appended to it; see Patch Application (on page 694).	
26757 OPERA 26758		ng operand shall be supported:	
26759	file	A pathname of a file to patch.	
26760 STDIN 26761			
26762 INPUT 26763	162 INPUT FILES 163 Input files shall be text files.		
26764 ENVIR 26765	26764 ENVIRONMENT VARIABLES 26765 The following environment variables shall affect the execution of <i>patch</i> :		
26766 26767 26768 26769	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)	
26770 26771	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.	

Utilities patch

26772 26773 26774	LC_CTYPE	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).
26775 26776 26777 26778	LC_MESSAC	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.
26779 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
26780 26781	LC_TIME	Determine the locale for recognizing the format of file timestamps written by the <i>diff</i> utility in a context-difference input file.

26782 ASYNCHRONOUS EVENTS

26783 Default.

26784 STDOUT

Not used. 26785

26786 STDERR

26793

26794

26795

26800

26801

26802 26803

26805

26806

26807

26808

The standard error shall be used for diagnostic and informational messages. 26787

26788 OUTPUT FILES

The output of the patch utility, the save files (.orig suffixes), and the reject files (.rej suffixes) 26789 26790 shall be text files.

26791 EXTENDED DESCRIPTION

A patch file may contain patching instructions for more than one file; filenames shall be 26792 determined as specified in Filename Determination (on page 694). When the -b option is specified, for each patched file, the original shall be saved in a file of the same name with the suffix .orig appended to it.

For each patched file, a reject file may also be created as noted in Patch Application (on page 26796 694). In the absence of a -r option, the name of this file shall be formed by appending the suffix 26797 **.rej** to the original filename. 26798

26799 **Patch File Format**

The patch file shall contain zero or more lines of header information followed by one or more patches. Each patch shall contain zero or more lines of filename identification in the format produced by diff -c, and one or more sets of diff output, which are customarily called hunks.

The *patch* utility shall recognize the following expression in the header information:

Index: pathname 26804

The file to be patched is named *pathname*.

If all lines (including headers) within a patch begin with the same leading sequence of

blank>s, the patch utility shall remove this sequence before proceeding. Within each patch, if the type of difference is context, the *patch* utility shall recognize the following expressions:

*** filename timestamp 26809

The patches arose from *filename*. 26810

--- filename timestamp 26811

The patches should be applied to *filename*. 26812

26813 Each hunk within a patch shall be the *diff* output to change a line range within the original file. The line numbers for successive hunks within a patch shall occur in ascending order. 26814

patch Utilities

Filename Determination

26830 XSI

If no *file* operand is specified, *patch* shall perform the following steps to determine the filename to use:

- 1. If the type of *diff* is context, the *patch* utility shall delete pathname components (as specified by the -**p** option) from the filename on the line beginning with "***", then test for the existence of this file relative to the current directory (or the directory specified with the -**d** option). If the file exists, the *patch* utility shall use this filename.
- 2. If the type of *diff* is context, the *patch* utility shall delete the pathname components (as specified by the -**p** option) from the filename on the line beginning with "---", then test for the existence of this file relative to the current directory (or the directory specified with the -**d** option). If the file exists, the *patch* utility shall use this filename.
- 3. If the header information contains a line beginning with the string **Index**:, the *patch* utility shall delete pathname components (as specified by the **-p** option) from this line, then test for the existence of this file relative to the current directory (or the directory specified with the **-d** option). If the file exists, the *patch* utility shall use this filename.
- 4. If an **SCCS** directory exists in the current directory, *patch* shall attempt to perform a *get* –**e SCCS**/*s.filename* command to retrieve an editable version of the file. If the file exists, the *patch* utility shall use this filename.
- 5. The *patch* utility shall write a prompt to standard output and request a filename interactively from the controlling terminal (for example, /dev/tty).

Patch Application

If the -c, -e, or -n option is present, the *patch* utility shall interpret information within each hunk as a context difference, an *ed* difference, or a normal difference, respectively. In the absence of any of these options, the *patch* utility shall determine the type of difference based on the format of information within the hunk.

For each hunk, the *patch* utility shall begin to search for the place to apply the patch at the line number at the beginning of the hunk, plus or minus any offset used in applying the previous hunk. If lines matching the hunk context are not found, *patch* shall scan both forwards and backwards at least 1 000 bytes for a set of lines that match the hunk context.

If no such place is found and it is a context difference, then another scan shall take place, ignoring the first and last line of context. If that fails, the first two and last two lines of context shall be ignored and another scan shall be made. Implementations may search more extensively for installation locations.

If no location can be found, the *patch* utility shall append the hunk to the reject file. The rejected hunk shall be written in context-difference format regardless of the format of the patch file. If the input was a normal or *ed*-style difference, the reject file may contain differences with zero lines of context. The line numbers on the hunks in the reject file may be different from the line numbers in the patch file since they shall reflect the approximate locations for the failed hunks in the new file rather than the old one.

26854 If the type of patch is an *ed* diff, the implementation may accomplish the patching by invoking the *ed* utility.

26856 EXIT STATUS

The following exit values shall be returned:

26858 0 Successful completion.

Utilities patch

One or more lines were written to a reject file.

26859

26898

26860 >1 An error occurred. 26861 CONSEQUENCES OF ERRORS 26862 Patches that cannot be correctly placed in the file shall be written to a reject file. **26863 APPLICATION USAGE** 26864 The $-\mathbf{R}$ option does not work with *ed* scripts because there is too little information to reconstruct 26865 the reverse operation. 26866 The -p option makes it possible to customize a patch file to local user directory structures 26867 without manually editing the patch file. For example, if the filename in the patch file was: /curds/whey/src/blurfl/blurfl.c 26868 26869 Setting **-p 0** gives the entire pathname unmodified; **-p 1** gives: curds/whey/src/blurfl/blurfl.c 26870 without the leading slash, -**p 4** gives: 26871 26872 blurfl/blurfl.c and not specifying –**p** at all gives: 26873 26874 blurfl.c . 26875 EXAMPLES 26876 None. 26877 RATIONALE Some of the functionality in historical *patch* implementations was not specified. The following 26878 documents those features present in historical implementations that have not been specified. 26879 A deleted piece of functionality was the '+' pseudo-option allowing an additional set of options 26880 and a patch file operand to be given. This was seen as being insufficiently useful to standardize. 26881 In historical implementations, if the string "Prereq:" appeared in the header, the patch utility 26882 26883 would search for the corresponding version information (the string specified in the header, delimited by
blank>s or the beginning or end of a line or the file) anywhere in the original file. 26884 This was deleted as too simplistic and insufficiently trustworthy a mechanism to standardize. 26885 For example, if: 26886 Prereq: 1.2 26887 26888 were in the header, the presence of a delimited 1.2 anywhere in the file would satisfy the 26889 prerequisite. The following options were dropped from historical implementations of patch as insufficiently 26890 useful to standardize: 26891 26892 -b The -b option historically provided a method for changing the name extension of the backup file from the default .orig. This option has been modified and retained 26893 in this volume of IEEE Std 1003.1-2001. 26894 $-\mathbf{F}$ The -F option specified the number of lines of a context diff to ignore when 26895 searching for a place to install a patch. 26896 $-\mathbf{f}$ 26897 The –**f** option historically caused *patch* not to request additional information from

the user.

patch **Utilities**

26899 26900	- r	The $-\mathbf{r}$ option historically provided a method of overriding the extension of the reject file from the default $.\mathbf{rej}$.		
26901	-s	The -s option historically caused <i>patch</i> to work silently unless an error occurred.		
26902	-x	The -x option historically set internal debugging flags.		
26903 26904 26905 26906 26907 26908 26909	the case of 1 filenames), the limit. It was of .orig and	system implementations, the saving of a .orig file may produce unwanted results. In 2, 13, or 14-character filenames (on file systems supporting 14-character maximum he .orig file overwrites the new file. The reject file may also exceed this filename suggested, due to some historical practice, that a tilde ('~') suffix be used instead some other character instead of the .rej suffix. This was rejected because it is not the user which file is which. The suffixes .orig and .rej are clearer and more ble.		
26910 26911 26912		In has the opposite sense in some historical implementations—do not save the .orig bult case here is not to save the files, making <i>patch</i> behave more consistently with the rd utilities.		
26913	The -w option	on in early proposals was changed to $-\mathbf{l}$ to match historical practice.		
26914 26915 26916 26917	previously a	on was included because without it, a non-interactive application cannot reject pplied patches. For example, if a user is piping the output of <i>diff</i> into the <i>patch</i> the user only wants to patch a file to a newer version non-interactively, the –N uired.		
26918 26919 26920 26921	addition to ju	the –l option description were proposed to allow matching across <newline>s in ust <blank>s. Since this is not historical practice, and since some ambiguities could aggested that future developments in this area utilize another option letter, such as</blank></newline>		
	REDIRECTIONS			
26923	None.			
26924 SEE ALS 26925	sO ed, diff			
26926 CHANC 26927	GE HISTORY First released			
26928 Issue 5 26929	The FUTURE	E DIRECTIONS section is added.		
26930 Issue 6 26931	This utility is	s marked as part of the User Portability Utilities option.		
26932 26933		ion of the $-\mathbf{D}$ option and the steps in Filename Determination (on page 694) are natch historical practice as defined in the IEEE P1003.2b draft standard.		
26934	The normativ	ve text is reworded to avoid use of the term "must" for application requirements.		
26935 26936		3.1-2001/Cor 1-2002, item XCU/TC1/D6/34 is applied, clarifying the way that the performs ifdef selection for the $-\mathbf{D}$ option.		

Utilities pathchk

26937 **NAME** 26938 pathchk — check pathnames 26939 SYNOPSIS 26940 pathchk [-p] pathname... 26941 **DESCRIPTION** The pathchk utility shall check that one or more pathnames are valid (that is, they could be used 26942 to access or create a file without causing syntax errors) and portable (that is, no filename 26943 truncation results). More extensive portability checks are provided by the $-\mathbf{p}$ option. 26944 26945 By default, the *pathchk* utility shall check each component of each *pathname* operand based on the underlying file system. A diagnostic shall be written for each *pathname* operand that: 26946 • Is longer than {PATH_MAX} bytes (see **Pathname Variable Values** in the Base Definitions 26947 26948 volume of IEEE Std 1003.1-2001, Chapter 13, Headers, < limits.h >) Contains any component longer than {NAME_MAX} bytes in its containing directory 26949 Contains any component in a directory that is not searchable 26950 • Contains any character in any component that is not valid in its containing directory 26951 The format of the diagnostic message is not specified, but shall indicate the error detected and 26952 the corresponding *pathname* operand. 26953 It shall not be considered an error if one or more components of a pathname operand do not exist 26954 26955 as long as a file matching the pathname specified by the missing components could be created that does not violate any of the checks specified above. 26956 26957 OPTIONS The pathchk utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 26958 12.2, Utility Syntax Guidelines. 26959 The following option shall be supported: 26960 26961 Instead of performing checks based on the underlying file system, write a -p diagnostic for each *pathname* operand that: 26962 • Is longer than {_POSIX_PATH_MAX} bytes (see **Minimum Values** in the Base 26963 Definitions volume of IEEE Std 1003.1-2001, Chapter 13, Headers, < limits.h>) 26964 • Contains any component longer than {_POSIX_NAME_MAX} bytes 26965 26966 Contains any character in any component that is not in the portable filename 26967 character set 26968 OPERANDS The following operand shall be supported: 26969 pathname 26970 A pathname to be checked. 26971 **STDIN** Not used. 26972 26973 INPUT FILES 26974 None. 26975 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *pathchk*: 26976

Provide a default value for the internationalization variables that are unset or null.

Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,

26977

26978

LANG

pathchk Utilities

26979 Internationalization Variables for the precedence of internationalization variables 26980 used to determine the values of locale categories.) 26981 LC ALL If set to a non-empty string value, override the values of all the other internationalization variables. 26982 26983 LC_CTYPE 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 26984 arguments). 26985 LC MESSAGES 26986 Determine the locale that should be used to affect the format and contents of 26987 diagnostic messages written to standard error. 26988 **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 26989 XSI **26990 ASYNCHRONOUS EVENTS** Default. 26991 26992 **STDOUT** Not used. 26993 **26994 STDERR** The standard error shall be used only for diagnostic messages. 26995 26996 OUTPUT FILES None. 26997 26998 EXTENDED DESCRIPTION None. 26999 27000 EXIT STATUS 27001 The following exit values shall be returned: 27002 All *pathname* operands passed all of the checks. 27003 >0 An error occurred. 27004 CONSEQUENCES OF ERRORS 27005 Default. 27006 APPLICATION USAGE 27007 The test utility can be used to determine whether a given pathname names an existing file; it does not, however, give any indication of whether or not any component of the pathname was 27008 27009 truncated in a directory where the POSIX NO TRUNC feature is not in effect. The pathchk 27010 utility does not check for file existence; it performs checks to determine whether a pathname does exist or could be created with no pathname component truncation. 27011 The noclobber option in the shell (see the set special built-in) can be used to atomically create a 27012 27013 file. As with all file creation semantics in the System Interfaces volume of IEEE Std 1003.1-2001, 27014 it guarantees atomic creation, but still depends on applications to agree on conventions and cooperate on the use of files after they have been created. 27015 27016 EXAMPLES To verify that all pathnames in an imported data interchange archive are legitimate and 27017 unambiguous on the current system: 27018

pax -f archive | sed -e '/ == .*/s///' | xargs pathchk

if [\$? -eq 0]

pax -r -f archive

then

27019

27020 27021

Utilities pathchk

```
27023
             else
27024
                  echo Investigate problems before importing files.
27025
                  exit 1
             fi
27026
             To verify that all files in the current directory hierarchy could be moved to any system
27027
             conforming to the System Interfaces volume of IEEE Std 1003.1-2001 that also supports the pax
27028
27029
             utility:
27030
             find . -print | xargs pathchk -p
27031
             if [ $? -eq 0 ]
27032
             then
27033
                  pax -w -f archive .
27034
             else
27035
                  echo Portable archive cannot be created.
                  exit 1
27036
27037
             fi
             To verify that a user-supplied pathname names a readable file and that the application can create
27038
27039
             a file extending the given path without truncation and without overwriting any existing file:
             case $- in
27040
                  *C*)
27041
                            reset="";;
                  *)
27042
                            reset="set +C"
27043
                            set -C;;
27044
             esac
27045
             test -r "$path" && pathchk "$path.out" &&
27046
                  rm "$path.out" > "$path.out"
27047
             if [ $? -ne 0 ]; then
                  printf "%s: %s not found or %s.out fails \
27048
             creation checks.\n" $0 "$path" "$path"
27049
27050
                  $reset
                              # Reset the noclobber option in case a trap
                               # on EXIT depends on it.
27051
                  exit 1
27052
27053
             fi
27054
             $reset
             PROCESSING < "$path" > "$path.out"
27055
27056
             The following assumptions are made in this example:
27057
               1. PROCESSING represents the code that is used by the application to use $path once it is
27058
                  verified that $path.out works as intended.
27059
               2. The state of the noclobber option is unknown when this code is invoked and should be set
27060
                  on exit to the state it was in when this code was invoked. (The reset variable is used in this
27061
                  example to restore the initial state.)
27062
               3. Note the usage of:
```

```
rm "$path.out" > "$path.out"
```

2706327064

2706527066

- a. The *pathchk* command has already verified, at this point, that **\$path.out** is not truncated.
- b. With the *noclobber* option set, the shell verifies that **Spath.out** does not already exist before invoking *rm*.

pathchk Utilities

c. If the shell succeeded in creating **Spath.out**, *rm* removes it so that the application can create the file again in the **PROCESSING** step.

d. If the **PROCESSING** step wants the file to exist already when it is invoked, the:

rm "\$path.out" > "\$path.out"

should be replaced with:

> "\$path.out"

which verifies that the file did not already exist, but leaves **\$path.out** in place for use by **PROCESSING**.

27076 RATIONALE

2707327074

27075

27077

27078

27079

27080

27081 27082

2708327084

27085 27086

27087 27088

2708927090

27091

27092

27093

27095

27096

27097 27098 The *pathchk* utility was new for the ISO POSIX-2:1993 standard. It, along with the *set* $-\mathbf{C}(noclobber)$ option added to the shell, replaces the *mktemp*, *validfnam*, and *create* utilities that appeared in early proposals. All of these utilities were attempts to solve several common problems:

- Verify the validity (for several different definitions of "valid") of a pathname supplied by a user, generated by an application, or imported from an external source.
- Atomically create a file.
- Perform various string handling functions to generate a temporary filename.

The *create* utility, included in an early proposal, provided checking and atomic creation in a single invocation of the utility; these are orthogonal issues and need not be grouped into a single utility. Note that the *noclobber* option also provides a way of creating a lock for process synchronization; since it provides an atomic *create*, there is no race between a test for existence and the following creation if it did not exist.

Having a function like *tmpnam()* in the ISO C standard is important in many high-level languages. The shell programming language, however, has built-in string manipulation facilities, making it very easy to construct temporary filenames. The names needed obviously depend on the application, but are frequently of a form similar to:

27094 \$TMPDIR/application abbreviation\$\$.suffix

In cases where there is likely to be contention for a given suffix, a simple shell **for** or **while** loop can be used with the shell *noclobber* option to create a file without risk of collisions, as long as applications trying to use the same filename name space are cooperating on the use of files after they have been created.

27099 FUTURE DIRECTIONS

27100 None.

27101 **SEE ALSO**

27102 Section 2.7 (on page 43), set (on page 86), test

27103 CHANGE HISTORY

First released in Issue 4.

```
27105 NAME
              pax — portable archive interchange
27106
27107 SYNOPSIS
              pax [-cdnv] [-H |-L] [-f archive] [-s replstr] ... [pattern...]
27108
27109
              pax -r[-cdiknuv][-H|-L][-f archive][-o options]...[-p string]...
                    [-s replstr]...[pattern...]
27110
27111
              pax -w[-dituvX][-H|-L][-b blocksize][[-a][-f archive][-o options]...
27112
                    [-s replstr]...[-x format][file...]
27113
              pax -r -w[-diklntuvX][-H|-L][-p string]...[-s replstr]...
27114
                    [file...] directory
27115 DESCRIPTION
              The pax utility shall read, write, and write lists of the members of archive files and copy
27116
              directory hierarchies. A variety of archive formats shall be supported; see the -x format option.
27117
              The action to be taken depends on the presence of the -r and -w options. The four combinations
27118
              of -r and -w are referred to as the four modes of operation: list, read, write, and copy modes,
27119
              corresponding respectively to the four forms shown in the SYNOPSIS section.
27120
              list
27121
                            In list mode (when neither –r nor –w are specified), pax shall write the names of
27122
                            the members of the archive file read from the standard input, with pathnames
27123
                            matching the specified patterns, to standard output. If a named file is of type
27124
                            directory, the file hierarchy rooted at that file shall be listed as well.
27125
              read
                            In read mode (when -r is specified, but -w is not), pax shall extract the members of
27126
                            the archive file read from the standard input, with pathnames matching the
                            specified patterns. If an extracted file is of type directory, the file hierarchy rooted
27127
                            at that file shall be extracted as well. The extracted files shall be created performing
27128
                            pathname resolution with the directory in which pax was invoked as the current
27129
                            working directory.
27130
27131
                            If an attempt is made to extract a directory when the directory already exists, this
27132
                            shall not be considered an error. If an attempt is made to extract a FIFO when the
27133
                            FIFO already exists, this shall not be considered an error.
27134
                            The ownership, access, and modification times, and file mode of the restored files
27135
                            are discussed under the -p option.
27136
              write
                            In write mode (when -w is specified, but -r is not), pax shall write the contents of
27137
                            the file operands to the standard output in an archive format. If no file operands are
27138
                            specified, a list of files to copy, one per line, shall be read from the standard input.
27139
                            A file of type directory shall include all of the files in the file hierarchy rooted at the
27140
                            In copy mode (when both -\mathbf{r} and -\mathbf{w} are specified), pax shall copy the file operands
27141
              copy
27142
                            to the destination directory.
                            If no file operands are specified, a list of files to copy, one per line, shall be read
27143
                            from the standard input. A file of type directory shall include all of the files in the
27144
27145
                            file hierarchy rooted at the file.
27146
                            The effect of the copy shall be as if the copied files were written to an archive file
27147
                            and then subsequently extracted, except that there may be hard links between the
27148
                            original and the copied files. If the destination directory is a subdirectory of one of
                            the files to be copied, the results are unspecified. If the destination directory is a
```

file of a type not defined by the System Interfaces volume of IEEE Std 1003.1-2001, the results are implementation-defined; otherwise, it shall be an error for the file named by the *directory* operand not to exist, not be writable by the user, or not be a file of type directory.

In **read** or **copy** modes, if intermediate directories are necessary to extract an archive member, *pax* shall perform actions equivalent to the *mkdir()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001, called with the following arguments:

- The intermediate directory used as the path argument
- The value of the bitwise-inclusive OR of S_IRWXU, S_IRWXG, and S_IRWXO as the mode argument

If any specified *pattern* or *file* operands are not matched by at least one file or archive member, *pax* shall write a diagnostic message to standard error for each one that did not match and exit with a non-zero exit status.

The archive formats described in the EXTENDED DESCRIPTION section shall be automatically detected on input. The default output archive format shall be implementation-defined.

A single archive can span multiple files. The *pax* utility shall determine, in an implementation-defined manner, what file to read or write as the next file.

If the selected archive format supports the specification of linked files, it shall be an error if these files cannot be linked when the archive is extracted. For archive formats that do not store file contents with each name that causes a hard link, if the file that contains the data is not extracted during this *pax* session, either the data shall be restored from the original file, or a diagnostic message shall be displayed with the name of a file that can be used to extract the data. In traversing directories, *pax* shall detect infinite loops; that is, entering a previously visited directory that is an ancestor of the last file visited. When it detects an infinite loop, *pax* shall write a diagnostic message to standard error and shall terminate.

27175 OPTIONS

27154

27155 27156

2715727158

2715927160

27161

27162

27163

27164

27165

27166

27167

27168 27169

27170

27171

27172

27173 27174

27176

27177 27178

27179

The *pax* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines, except that the order of presentation of the $-\mathbf{o}$, $-\mathbf{p}$, and $-\mathbf{s}$ options is significant.

The following options shall be supported:

- 27180 Read an archive file from standard input.
- 27181 —w Write files to the standard output in the specified archive format.
- Append files to the end of the archive. It is implementation-defined which devices on the system support appending. Additional file formats unspecified by this volume of IEEE Std 1003.1-2001 may impose restrictions on appending.
- 27185 b blocksize Block the output at a positive decimal integer number of bytes per write to the archive file. Devices and archive formats may impose restrictions on blocking. Blocking shall be automatically determined on input. Conforming applications shall not specify a blocksize value larger than 32 256. Default blocking when creating archives depends on the archive format. (See the –x option below.)
- 27190 Match all file or archive members except those specified by the *pattern* or *file* operands.
- Cause files of type directory being copied or archived or archive members of type directory being extracted or listed to match only the file or archive member itself and not the file hierarchy rooted at the file.

27195 27196	-f archive	Specify the pathname of the input or output archive, overriding the default standard input (in list or read modes) or standard output (write mode).
27197 27198 27199 27200 27201 27202 27203	–Н	If a symbolic link referencing a file of type directory is specified on the command line, <i>pax</i> shall archive the file hierarchy rooted in the file referenced by the link, using the name of the link as the root of the file hierarchy. Otherwise, if a symbolic link referencing a file of any other file type which <i>pax</i> can normally archive is specified on the command line, then <i>pax</i> shall archive the file referenced by the link, using the name of the link. The default behavior shall be to archive the symbolic link itself.
27204 27205 27206 27207 27208 27209 27210 27211 27212 27213	−i	Interactively rename files or archive members. For each archive member matching a <i>pattern</i> operand or file matching a <i>file</i> operand, a prompt shall be written to the file /dev/tty. The prompt shall contain the name of the file or archive member, but the format is otherwise unspecified. A line shall then be read from /dev/tty. If this line is blank, the file or archive member shall be skipped. If this line consists of a single period, the file or archive member shall be processed with no modification to its name. Otherwise, its name shall be replaced with the contents of the line. The <i>pax</i> utility shall immediately exit with a non-zero exit status if end-of-file is encountered when reading a response or if /dev/tty cannot be opened for reading and writing.
27214 27215		The results of extracting a hard link to a file that has been renamed during extraction are unspecified.
27216	$-\mathbf{k}$	Prevent the overwriting of existing files.
27217 27218 27219 27220 27221 27222 27223	- l	(The letter ell.) In copy mode, hard links shall be made between the source and destination file hierarchies whenever possible. If specified in conjunction with $-\mathbf{H}$ or $-\mathbf{L}$, when a symbolic link is encountered, the hard link created in the destination file hierarchy shall be to the file referenced by the symbolic link. If specified when neither $-\mathbf{H}$ nor $-\mathbf{L}$ is specified, when a symbolic link is encountered, the implementation shall create a hard link to the symbolic link in the source file hierarchy or copy the symbolic link to the destination.
27224 27225 27226 27227 27228 27229 27230 27231	–L	If a symbolic link referencing a file of type directory is specified on the command line or encountered during the traversal of a file hierarchy, <i>pax</i> shall archive the file hierarchy rooted in the file referenced by the link, using the name of the link as the root of the file hierarchy. Otherwise, if a symbolic link referencing a file of any other file type which <i>pax</i> can normally archive is specified on the command line or encountered during the traversal of a file hierarchy, <i>pax</i> shall archive the file referenced by the link, using the name of the link. The default behavior shall be to archive the symbolic link itself.
27232 27233 27234	- n	Select the first archive member that matches each <i>pattern</i> operand. No more than one archive member shall be matched for each pattern (although members of type directory shall still match the file hierarchy rooted at that file).
27235 27236 27237	−o options	Provide information to the implementation to modify the algorithm for extracting or writing files. The value of <i>options</i> shall consist of one or more comma-separated keywords of the form:
27238		<pre>keyword[[:]=value][,keyword[[:]=value],]</pre>
27239 27240 27241		Some keywords apply only to certain file formats, as indicated with each description. Use of keywords that are inapplicable to the file format being processed produces undefined results.

Keywords in the *options* argument shall be a string that would be a valid portable filename as described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.276, Portable Filename Character Set.

Note: Keywords are not expected to be filenames, merely to follow the same character composition rules as portable filenames.

Keywords can be preceded with white space. The *value* field shall consist of zero or more characters; within *value*, the application shall precede any literal comma with a backslash, which shall be ignored, but preserves the comma as part of *value*. A comma as the final character, or a comma followed solely by white space as the final characters, in *options* shall be ignored. Multiple $-\mathbf{o}$ options can be specified; if keywords given to these multiple $-\mathbf{o}$ options conflict, the keywords and values appearing later in command line sequence shall take precedence and the earlier shall be silently ignored. The following keyword values of *options* shall be supported for the file formats as indicated:

delete=pattern

(Applicable only to the **–x pax** format.) When used in **write** or **copy** mode, *pax* shall omit from extended header records that it produces any keywords matching the string pattern. When used in **read** or **list** mode, *pax* shall ignore any keywords matching the string pattern in the extended header records. In both cases, matching shall be performed using the pattern matching notation described in Section 2.13.1 (on page 62) and Section 2.13.2 (on page 63). For example:

```
-o delete=security.*
```

would suppress security-related information. See **pax Extended Header** (on page 714) for extended header record keyword usage.

exthdr.name=string

(Applicable only to the **–x pax** format.) This keyword allows user control over the name that is written into the **ustar** header blocks for the extended header produced under the circumstances described in **pax Header Block** (on page 713). The name shall be the contents of *string*, after the following character substitutions have been made:

string Includes:	Replaced By:
%d	The directory name of the file, equivalent to the result of the <i>dirname</i> utility on the translated pathname.
%f	The filename of the file, equivalent to the result of the <i>basename</i> utility on the translated pathname.
%p	The process ID of the <i>pax</i> process.
왕왕	A'%' character.

Any other ' %' characters in *string* produce undefined results.

If no $-\mathbf{o}$ **exthdr.name**=*string* is specified, *pax* shall use the following default value:

```
%d/PaxHeaders.%p/%f
```

globexthdr.name=*string*

(Applicable only to the -x pax format.) When used in write or copy mode with the appropriate options, pax shall create global extended header records

27288	with ustar header blocks that will be treated as regular files by previous		
27289	versions of pax. This keyword allows user control over the name that is		
27290	written into the ustar header blocks for global extended header records. The		
27291	name shall be the contents of string, after the following character substitutions		
27292	have been ma	have been made:	
27293	string		
27294	Includes:	Replaced By:	
27295	%n	An integer that represents the sequence number of the global	
27296		extended header record in the archive, starting at 1.	
27297	%p	The process ID of the <i>pax</i> process.	
27298	% %	A '%' character.	
27299	Any other ' %	' characters in <i>string</i> produce undefined results.	
27300	If no -o glo	bexthdr.name=string is specified, pax shall use the following	
27301	default value		
27302	\$TMPDIR/Gl	obalHead.%p.%n	
27303	where \$TMP	DIR represents the value of the TMPDIR environment variable. If	
27304	TMPDIR is no	TMPDIR is not set, pax shall use /tmp.	
27305	invalid=action		
27306		only to the –x pax format.) This keyword allows user control over	
27307		x takes upon encountering values in an extended header record	
27308		or copy mode, are invalid in the destination hierarchy or, in list	
27309		ot be written in the codeset and current locale of the	
27310	•	ion. The following are invalid values that shall be recognized by	
27311	pax:		
27312		r copy mode, a filename or link name that contains character	
27313		encodings invalid in the destination hierarchy. (For example, the name	
27314	may contain embedded NULs.)		
27315		r copy mode, a filename or link name that is longer than the	
27316		n allowed in the destination hierarchy (for either a pathname	
27317	componer	component or the entire pathname).	
27318		ode, any character string value (filename, link name, user name,	
27319		n) that cannot be written in the codeset and current locale of the	
27320	implemer	ntation.	
27321	The following	ng mutually-exclusive values of the action argument are	
27322	supported:		
27323	bypass	In read or copy mode, <i>pax</i> shall bypass the file, causing no	

UTF-8 When used in **read**, **copy**, or **list** mode and a filename, link name, owner name, or any other field in an extended header

invalid values is unspecified.

change to the destination hierarchy. In **list** mode, pax shall write

all requested valid values for the file, but its method for writing

In **read** or **copy** mode, *pax* shall act as if the -i option were in

effect for each file with invalid filename or link name values,

allowing the user to provide a replacement name interactively.

rename

27324

27325

27326

27327

27328

2732927330

27331

record cannot be translated from the **pax** UTF-8 codeset format to the codeset and current locale of the implementation, *pax* shall use the actual UTF-8 encoding for the name.

write

In **read** or **copy** mode, *pax* shall write the file, translating or truncating the name, regardless of whether this may overwrite an existing file with a valid name. In **list** mode, *pax* shall behave identically to the **bypass** action.

If no **–o invalid**= option is specified, *pax* shall act as if **–oinvalid=bypass** were specified. Any overwriting of existing files that may be allowed by the **–oinvalid**= actions shall be subject to permission (**–p**) and modification time (**–u**) restrictions, and shall be suppressed if the **–k** option is also specified.

linkdata

(Applicable only to the -x pax format.) In write mode, pax shall write the contents of a file to the archive even when that file is merely a hard link to a file whose contents have already been written to the archive.

listopt=format

This keyword specifies the output format of the table of contents produced when the **-v** option is specified in **list** mode. See **List Mode Format Specifications** (on page 709). To avoid ambiguity, the **listopt**=*format* shall be the only or final **keyword**=*value* pair in a **-o** option-argument; all characters in the remainder of the option-argument shall be considered part of the format string. When multiple **-olistopt**=*format* options are specified, the format strings shall be considered a single, concatenated string, evaluated in command line order.

times

(Applicable only to the **–x** *pax* format.) When used in **write** or **copy** mode, *pax* | shall include **atime**, **ctime**, and **mtime** extended header records for each file. See **pax Extended Header File Times** (on page 717).

In addition to these keywords, if the **–x** *pax* format is specified, any of the keywords and values defined in **pax Extended Header** (on page 714), including implementation extensions, can be used in **–o** option-arguments, in either of two modes:

keyword=*value*

When used in **write** or **copy** mode, these keyword/value pairs shall be included at the beginning of the archive as **typeflag g** global extended header records. When used in **read** or **list** mode, these keyword/value pairs shall act as if they had been at the beginning of the archive as **typeflag g** global extended header records.

keyword:=value

When used in **write** or **copy** mode, these keyword/value pairs shall be included as records at the beginning of a **typeflag x** extended header for each file. (This shall be equivalent to the equal-sign form except that it creates no **typeflag g** global extended header records.) When used in **read** or **list** mode, these keyword/value pairs shall act as if they were included as records at the end of each extended header; thus, they shall override any global or file-specific extended header record keywords of the same names. For example, in the command:

27380 27381 27382		<pre>pax -r -o " gname:=mygroup, " <archive< pre=""></archive<></pre>
27383 27384		the group name will be forced to a new value for all files read from the archive.
27385 27386		The precedence of –o keywords over various fields in the archive is described in pax Extended Header Keyword Precedence (on page 717).
27387 27388 27389 27390 27391 27392	− p string	Specify one or more file characteristic options (privileges). The <i>string</i> optionargument shall be a string specifying file characteristics to be retained or discarded on extraction. The string shall consist of the specification characters a, e, m, o , and p . Other implementation-defined characters can be included. Multiple characteristics can be concatenated within the same string and multiple $-\mathbf{p}$ options can be specified. The meaning of the specification characters are as follows:
27393		a Do not preserve file access times.
27394 27395 27396		e Preserve the user ID, group ID, file mode bits (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 3.168, File Mode Bits), access time, modification time, and any other implementation-defined file characteristics.
27397		m Do not preserve file modification times.
27398		o Preserve the user ID and group ID.
27399 27400		p Preserve the file mode bits. Other implementation-defined file mode attributes may be preserved.
27401 27402 27403 27404 27405 27406		In the preceding list, "preserve" indicates that an attribute stored in the archive shall be given to the extracted file, subject to the permissions of the invoking process. The access and modification times of the file shall be preserved unless otherwise specified with the $-\mathbf{p}$ option or not stored in the archive. All attributes that are not preserved shall be determined as part of the normal file creation action (see Section 1.7.1.4 (on page 4)).
27407 27408 27409		If neither the \circ nor the \circ specification character is specified, or the user ID and group ID are not preserved for any reason, pax shall not set the S_ISUID and S_ISGID bits of the file mode.
27410 27411 27412		If the preservation of any of these items fails for any reason, <i>pax</i> shall write a diagnostic message to standard error. Failure to preserve these items shall affect the final exit status, but shall not cause the extracted file to be deleted.
27413 27414 27415		If file characteristic letters in any of the <i>string</i> option-arguments are duplicated or conflict with each other, the ones given last shall take precedence. For example, if $-\mathbf{p}$ eme is specified, file modification times are preserved.
27416 27417 27418 27419	− s replstr	Modify file or archive member names named by <i>pattern</i> or <i>file</i> operands according to the substitution expression <i>replstr</i> , using the syntax of the <i>ed</i> utility. The concepts of "address" and "line" are meaningless in the context of the <i>pax</i> utility, and shall not be supplied. The format shall be:
27420		-s /old/new/[gp]
27421 27422 27423		where as in ed , old is a basic regular expression and new can contain an ampersand, '\n' (where n is a digit) backreferences, or subexpression matching. The old string shall also be permitted to contain <newline>s.</newline>

27424 27425 27426 27427 27428 27429		expressions specified, ter 'g' is as de substitutions	ll character can be used as a delimiter ($'/'$ shown here). Multiple –s can be specified; the expressions shall be applied in the order rminating with the first successful substitution. The optional trailing fined in the ed utility. The optional trailing $'p'$ shall cause successful so to be written to standard error. File or archive member names that the empty string shall be ignored when reading and writing archives.
27430 27431 27432	-t	required by	ng files from the file system, and if the user has the permissions <i>utime</i> () to do so, set the access time of each file read to the access time efore being read by <i>pax</i> .
27433 27434 27435 27436 27437 27438 27439 27440 27441	−u	existing file member wit archive men the same na than the archiv replacement the destinati	that are older (having a less recent file modification time) than a pre- or archive member with the same name. In read mode, an archive h the same name as a file in the file system shall be extracted if the aber is newer than the file. In write mode, an archive file member with me as a file in the file system shall be superseded if the file is newer nive member. If $-\mathbf{a}$ is also specified, this is accomplished by appending we; otherwise, it is unspecified whether this is accomplished by actual in the archive or by appending to the archive. In copy mode, the file in on hierarchy shall be replaced by the file in the source hierarchy or by file in the source hierarchy if the file in the source hierarchy is newer.
27443 27444 27445	- v		e, produce a verbose table of contents (see the STDOUT section). write archive member pathnames to standard error (see the STDERR
27446 27447	-x format	Specify the formats:	output archive format. The pax utility shall support the following
27448 27449 27450 27451		cpio	The cpio interchange format; see the EXTENDED DESCRIPTION section. The default <i>blocksize</i> for this format for character special archive files shall be 5120. Implementations shall support all <i>blocksize</i> values less than or equal to 32 256 that are multiples of 512.
27452 27453 27454 27455		pax	The pax interchange format; see the EXTENDED DESCRIPTION section. The default <i>blocksize</i> for this format for character special archive files shall be 5120. Implementations shall support all <i>blocksize</i> values less than or equal to 32 256 that are multiples of 512.
27456 27457 27458 27459		ustar	The tar interchange format; see the EXTENDED DESCRIPTION section. The default <i>blocksize</i> for this format for character special archive files shall be 10 240. Implementations shall support all <i>blocksize</i> values less than or equal to 32 256 that are multiples of 512.
27460 27461		-	tion-defined formats shall specify a default block size as well as any sizes supported for character special archive files.
27462 27463			t to append to an archive file in a format different from the existing nat shall cause <i>pax</i> to exit immediately with a non-zero exit status.
27464 27465		In copy mospecified.	de, if no -x format is specified, pax shall behave as if -xpax were
27466 27467 27468	-X	into director	rsing the file hierarchy specified by a pathname, <i>pax</i> shall not descend ries that have a different device ID (<i>st_dev</i> ; see the System Interfaces EEE Std 1003.1-2001, <i>stat()</i>).

The options that operate on the names of files or archive members ($-\mathbf{c}$, $-\mathbf{i}$, $-\mathbf{n}$, $-\mathbf{s}$, $-\mathbf{u}$, and $-\mathbf{v}$) shall interact as follows. In **read** mode, the archive members shall be selected based on the user-specified *pattern* operands as modified by the $-\mathbf{c}$, $-\mathbf{n}$, and $-\mathbf{u}$ options. Then, any $-\mathbf{s}$ and $-\mathbf{i}$ options shall modify, in that order, the names of the selected files. The $-\mathbf{v}$ option shall write names resulting from these modifications.

In **write** mode, the files shall be selected based on the user-specified pathnames as modified by the $-\mathbf{n}$ and $-\mathbf{u}$ options. Then, any $-\mathbf{s}$ and $-\mathbf{i}$ options shall modify, in that order, the names of these selected files. The $-\mathbf{v}$ option shall write names resulting from these modifications.

If both the $-\mathbf{u}$ and $-\mathbf{n}$ options are specified, pax shall not consider a file selected unless it is newer than the file to which it is compared.

List Mode Format Specifications

In **list** mode with the **–o listopt**=*format* option, the *format* argument shall be applied for each selected file. The *pax* utility shall append a <newline> to the **listopt** output for each selected file. The *format* argument shall be used as the *format* string described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5, File Format Notation, with the exceptions 1. through 5. defined in the EXTENDED DESCRIPTION section of *printf*, plus the following exceptions:

- 6. The sequence (*keyword*) can occur before a format conversion specifier. The conversion argument is defined by the value of *keyword*. The implementation shall support the following keywords:
 - Any of the Field Name entries in Table 4-13 (on page 718) and Table 4-15 (on page 721). The implementation may support the *cpio* keywords without the leading **c**_ in addition to the form required by Table 4-16 (on page 722).
 - Any keyword defined for the extended header in **pax Extended Header** (on page 714).
 - Any keyword provided as an implementation-defined extension within the extended header defined in **pax Extended Header** (on page 714).

For example, the sequence "% (charset) s" is the string value of the name of the character set in the extended header.

The result of the keyword conversion argument shall be the value from the applicable header field or extended header, without any trailing NULs.

All keyword values used as conversion arguments shall be translated from the UTF-8 encoding to the character set appropriate for the local file system, user database, and so on, as applicable.

7. An additional conversion specifier character, T, shall be used to specify time formats. The T conversion specifier character can be preceded by the sequence (*keyword=subformat*), where *subformat* is a date format as defined by *date* operands. The default *keyword* shall be **mtime** and the default subformat shall be:

```
%b %e %H:%M %Y
```

- 8. An additional conversion specifier character, M, shall be used to specify the file mode string as defined in *ls* Standard Output. If (*keyword*) is omitted, the **mode** keyword shall be used. For example, % . 1M writes the single character corresponding to the *<entry type>* field of the *ls* –1 command.
- 9. An additional conversion specifier character, D, shall be used to specify the device for block or special files, if applicable, in an implementation-defined format. If not applicable, and (keyword) is specified, then this conversion shall be equivalent to % (keyword) u. If not

27513 applicable, and (keyword) is omitted, then this conversion shall be equivalent to <space>. 10. An additional conversion specifier character, F, shall be used to specify a pathname. The F 27514 27515 conversion character can be preceded by a sequence of comma-separated keywords: 27516 (keyword[,keyword] ...) The values for all the keywords that are non-null shall be concatenated together, each 27517 27518 separated by a '/'. The default shall be (path) if the keyword path is defined; otherwise, 27519 the default shall be (**prefix**,**name**). 11. An additional conversion specifier character, L, shall be used to specify a symbolic line 27520 27521 expansion. If the current file is a symbolic link, then %L shall expand to: "%s -> %s", <value of keyword>, <contents of link> 27522 27523 Otherwise, the %L conversion specification shall be the equivalent of %F. 27524 OPERANDS 27525 The following operands shall be supported: directory The destination directory pathname for **copy** mode. 27526 file 27527 A pathname of a file to be copied or archived. A pattern matching one or more pathnames of archive members. A pattern must 27528 pattern 27529 be given in the name-generating notation of the pattern matching notation in Section 2.13 (on page 62), including the filename expansion rules in Section 2.13.3 27530 27531 (on page 63). The default, if no pattern is specified, is to select all members in the 27532 archive. 27533 **STDIN** In write mode, the standard input shall be used only if no file operands are specified. It shall be a 27534 27535 text file containing a list of pathnames, one per line, without leading or trailing

 llank>s. 27536 In **list** and **read** modes, if **-f** is not specified, the standard input shall be an archive file. 27537 Otherwise, the standard input shall not be used. 27538 INPUT FILES The input file named by the archive option-argument, or standard input when the archive is read 27539 27540 from there, shall be a file formatted according to one of the specifications in the EXTENDED 27541 DESCRIPTION section or some other implementation-defined format. 27542 The file **/dev/tty** shall be used to write prompts and read responses. 27543 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *pax*: 27544 LANG Provide a default value for the internationalization variables that are unset or null. 27545 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 27546 27547 Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 27548 27549 LC_ALL If set to a non-empty string value, override the values of all the other 27550 internationalization variables. LC_COLLATE 27551 27552 Determine the locale for the behavior of ranges, equivalence classes, and multi-27553 character collating elements used in the pattern matching expressions for the 27554 pattern operand, the basic regular expression for the -s option, and the extended 27555 regular expression defined for the **yesexpr** locale keyword in the *LC_MESSAGES*

27556		category.
27557 27558 27559 27560 27561	LC_CTYPE	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), the behavior of character classes used in the extended regular expression defined for the yesexpr locale keyword in the <i>LC_MESSAGES</i> category, and pattern matching.
27562 27563 27564 27565	LC_MESSAG	Determine the locale for the processing of affirmative responses that should be used to affect the format and contents of diagnostic messages written to standard error.
27566 27567	LC_TIME	Determine the format and contents of date and time strings when the $-\mathbf{v}$ option is specified.
27568 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
27569 27570	TMPDIR	Determine the pathname that provides part of the default global extended header record file, as described for the $-\mathbf{o}$ globexthdr= keyword in the OPTIONS section.
27571 27572	TZ	Determine the timezone used to calculate date and time strings when the $-\mathbf{v}$ option is specified. If TZ is unset or null, an unspecified default timezone shall be used.

27573 ASYNCHRONOUS EVENTS

27574 Default.

27575 STDOUT

In **write** mode, if **-f** is not specified, the standard output shall be the archive formatted according to one of the specifications in the EXTENDED DESCRIPTION section, or some other implementation-defined format (see **-x** format).

In **list** mode, when the **-olistopt**=*format* has been specified, the selected archive members shall be written to standard output using the format described under **List Mode Format**Specifications (on page 709). In **list** mode without the **-olistopt**=*format* option, the table of contents of the selected archive members shall be written to standard output using the following format:

27584 "%s\n", <pathname>

27585 If the **-v** option is specified in **list** mode, the table of contents of the selected archive members shall be written to standard output using the following formats.

For pathnames representing hard links to previous members of the archive:

27588 " $s\Delta = \Delta s n$ ", $< ls -1 \ listing>$, < linkname>

For all other pathnames:

27590 "%s\n", <ls -l listing>

where *<ls* – *listing>* shall be the format specified by the *ls* utility with the –*l* option. When writing pathnames in this format, it is unspecified what is written for fields for which the underlying archive format does not have the correct information, although the correct number of *<*blank>-separated fields shall be written.

27595 In **list** mode, standard output shall not be buffered more than a line at a time.

27596 STDERR

27597 If **-v** is specified in **read**, **write**, or **copy** modes, *pax* shall write the pathnames it processes to the standard error output using the following format:

27599 "%s\n", <pathname>

These pathnames shall be written as soon as processing is begun on the file or archive member, and shall be flushed to standard error. The trailing <newline>, which shall not be buffered, is written when the file has been read or written.

27603 If the -s option is specified, and the replacement string has a trailing 'p', substitutions shall be written to standard error in the following format:

27605 " $s\Delta > \Delta s n$ ", <original pathname>, <new pathname>

In all operating modes of *pax*, optional messages of unspecified format concerning the input archive format and volume number, the number of files, blocks, volumes, and media parts as well as other diagnostic messages may be written to standard error.

In all formats, for both standard output and standard error, it is unspecified how non-printable characters in pathnames or link names are written.

When pax is in **read** mode or **list** mode, using the **-xpax** archive format, and a filename, link name, owner name, or any other field in an extended header record cannot be translated from the **pax** UTF-8 codeset format to the codeset and current locale of the implementation, pax shall write a diagnostic message to standard error, shall process the file as described for the **-o** invalid=option, and then shall process the next file in the archive.

27616 OUTPUT FILES

27617

27618

27619

2762027621

27622

27623

27625

27626 27627

27628 27629

27630

2763127632

27633

27634

27635

27636

In **read** mode, the extracted output files shall be of the archived file type. In **copy** mode, the copied output files shall be the type of the file being copied. In either mode, existing files in the destination hierarchy shall be overwritten only when all permission $(-\mathbf{p})$, modification time $(-\mathbf{u})$, and invalid-value $(-\mathbf{oinvalid}=)$ tests allow it.

In write mode, the output file named by the **-f** option-argument shall be a file formatted according to one of the specifications in the EXTENDED DESCRIPTION section, or some other implementation-defined format.

27624 EXTENDED DESCRIPTION

pax Interchange Format

A *pax* archive tape or file produced in the **–xpax** format shall contain a series of blocks. The physical layout of the archive shall be identical to the **ustar** format described in **ustar Interchange Format** (on page 717). Each file archived shall be represented by the following sequence:

- An optional header block with extended header records. This header block is of the form described in **pax Header Block** (on page 713), with a *typeflag* value of **x** or **g**. The extended header records, described in **pax Extended Header** (on page 714), shall be included as the data for this header block.
- A header block that describes the file. Any fields in the preceding optional extended header shall override the associated fields in this header block for this file.
- Zero or more blocks that contain the contents of the file.

At the end of the archive file there shall be two 512-byte blocks filled with binary zeros, interpreted as an end-of-archive indicator.

 A schematic of an example archive with global extended header records and two actual files is shown in Figure 4-1. In the example, the second file in the archive has no extended header preceding it, presumably because it has no need for extended attributes.

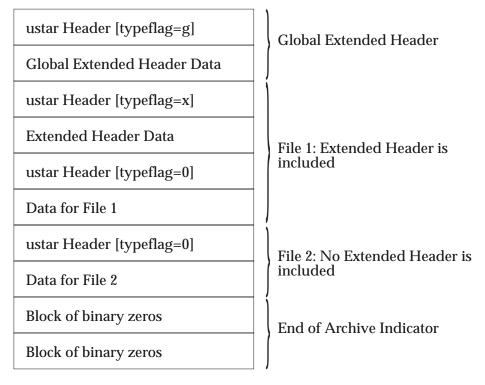


Figure 4-1 pax Format Archive Example

pax Header Block

The **pax** header block shall be identical to the **ustar** header block described in **ustar Interchange Format** (on page 717), except that two additional *typeflag* values are defined:

- x Represents extended header records for the following file in the archive (which shall have its own **ustar** header block). The format of these extended header records shall be as described in **pax Extended Header** (on page 714).
- g Represents global extended header records for the following files in the archive. The format of these extended header records shall be as described in **pax Extended Header** (on page 714). Each value shall affect all subsequent files that do not override that value in their own extended header record and until another global extended header record is reached that provides another value for the same field. The *typeflag g* global headers should not be used with interchange media that could suffer partial data loss in transporting the archive.

For both of these types, the *size* field shall be the size of the extended header records in octets. The other fields in the header block are not meaningful to this version of the *pax* utility. However, if this archive is read by a *pax* utility conforming to the ISO POSIX-2: 1993 standard, the header block fields are used to create a regular file that contains the extended header records as data. Therefore, header block field values should be selected to provide reasonable file access to this regular file.

A further difference from the **ustar** header block is that data blocks for files of *typeflag* 1 (the digit one) (hard link) may be included, which means that the size field may be greater than zero. Archives created by *pax* –**o** linkdata shall include these data blocks with the hard links.

pax Extended Header

A pax extended header contains values that are inappropriate for the ustar header block because of limitations in that format: fields requiring a character encoding other than that described in the ISO/IEC 646: 1991 standard, fields representing file attributes not described in the ustar header, and fields whose format or length do not fit the requirements of the ustar header. The values in an extended header add attributes to the following file (or files; see the description of the typeflag g header block) or override values in the following header block(s), as indicated in

An extended header shall consist of one or more records, each constructed as follows:

27674 "%d %s=%s\n", <length>, <keyword>, <value>

the following list of keywords.

The extended header records shall be encoded according to the ISO/IEC 10646-1:2000 standard (UTF-8). The <*length*> field, <*blank*>, equals sign, and <*newline*> shown shall be limited to the portable character set, as encoded in UTF-8. The <*keyword*> and <*value*> fields can be any UTF-8 characters. The <*length*> field shall be the decimal length of the extended header record in octets, including the trailing <*newline*>.

The < keyword> field shall be one of the entries from the following list or a keyword provided as an implementation extension. Keywords consisting entirely of lowercase letters, digits, and periods are reserved for future standardization. A keyword shall not include an equals sign. (In the following list, the notations "file(s)" or "block(s)" is used to acknowledge that a keyword affects the following single file after a *typeflag* x extended header, but possibly multiple files after *typeflag* y. Any requirements in the list for y to include a record when in y to y mode shall apply only when such a record has not already been provided through the use of the y option. When used in y mode, y shall behave as if an archive had been created with applicable extended header records and then extracted.)

atime

The file access time for the following file(s), equivalent to the value of the *st_atime* member of the **stat** structure for a file, as described by the *stat*() function. The access time shall be restored if the process has the appropriate privilege required to do so. The format of the *<value>* shall be as described in **pax Extended Header File Times** (on page 717).

charset

The name of the character set used to encode the data in the following file(s). The entries in the following table are defined to refer to known standards; additional names may be agreed on between the originator and recipient.

27697				
27698		<value></value>	Formal Standard	
27699		ISO-IR Δ 646 Δ 1990	ISO/IEC 646: 1990	
27700		ISO-IR Δ 8859 Δ 1 Δ 1998	ISO/IEC 8859-1:1998	
27701		ISO-IR Δ 8859 Δ 2 Δ 1999	ISO/IEC 8859-2: 1999	
27702		ISO-IR Δ 8859 Δ 3 Δ 1999	ISO/IEC 8859-3: 1999	
27703		ISO-IR Δ 8859 Δ 4 Δ 1998	ISO/IEC 8859-4: 1998	
27704		ISO-IR Δ 8859 Δ 5 Δ 1999	ISO/IEC 8859-5: 1999	
27705		ISO-IR Δ 8859 Δ 6 Δ 1999	ISO/IEC 8859-6: 1999	
27706		ISO-IR Δ 8859 Δ 7 Δ 1987	ISO/IEC 8859-7: 1987	
27707		ISO-IRΔ8859Δ8Δ1999	ISO/IEC 8859-8: 1999	
27708		ISO-IR Δ 8859 Δ 9 Δ 1999	ISO/IEC 8859-9: 1999	
27709		ISO-IR Δ 8859 Δ 10 Δ 1998	ISO/IEC 8859-10: 1998	
27710		ISO-IR Δ 8859 Δ 13 Δ 1998	ISO/IEC 8859-13: 1998	
27711		ISO-IR Δ 8859 Δ 14 Δ 1998	ISO/IEC 8859-14: 1998	
27712		ISO-IR Δ 8859 Δ 15 Δ 1999	ISO/IEC 8859-15: 1999	
27713		ISO-IR Δ 10646 Δ 2000	ISO/IEC 10646: 2000	
27714		ISO-IR Δ 10646 Δ 2000 Δ UTF-8	ISO/IEC 10646, UTF-8 encoding	
27715		BINARY	None.	
27716		The encoding is included in an extended	header for information only, when now	
27717		used as described in IEEE Std 1003.1-20		
27717		any other encoding. The BINARY entry		
21110		· ·	·	
27719		When used in write or copy mode, it		
27720		includes a charset extended header recor	d for a file.	
27721	comment	A series of characters used as a comment. All characters in the <value> field shall</value>		
27722		be ignored by pax.		
27723	ctime	The file creation time for the following	or file(s) equivalent to the value of the	
27724	CHIIC	st_ctime member of the stat structure for		
27725		The creation time shall be restored if t		
27726		required to do so. The format of the <i><va< i=""></va<></i>		
27727		Header File Times (on page 717).	ide> shan be as described in pax Extende	
21121		1 0		
27728	gid	The group ID of the group that owns the		
27729		digits from the ISO/IEC 646: 1991 standa		
27730		in the following header block(s). When used in write or copy mode, pax shall		
27731		include a <i>gid</i> extended header record for each file whose group ID is greater than		
27732		2 097 151 (octal 7 777 777).		
27733	gname	The group of the file(s), formatted as a	group name in the group database. Th	
27734	8	record shall override the gid and gname f		
27734		any gid extended header record. When the		
27736		translate the name from the UTF-8 enco	20 1	
27737		set appropriate for the group database of	•	
		characters cannot be translated, and if th		
27738		the results are implementation-defined.		
27739		<u>-</u>		
27740		shall include a gname extended header		
27741		cannot be represented entirely with the	ietters and digits of the portable charact	
27742		set.		
27743	linkpath	The pathname of a link being created	to another file, of any type, previous	
27744	_	archived. This record shall override the		
27745		block(s). The following ustar header block	•	

27746 If typeflag of the following header block is 1, it shall be a hard link. If typeflag is 2, it 27747 shall be a symbolic link and the linkpath value shall be the contents of the symbolic link. The pax utility shall translate the name of the link (contents of the 27748 symbolic link) from the UTF-8 encoding to the character set appropriate for the 27749 27750 local file system. When used in write or copy mode, pax shall include a linkpath 27751 extended header record for each link whose pathname cannot be represented entirely with the members of the portable character set other than NUL. 27752 27753 mtime The file modification time of the following file(s), equivalent to the value of the st_mtime member of the stat structure for a file, as described in the stat() function. 27754 This record shall override the *mtime* field in the following header block(s). The 27755 modification time shall be restored if the process has the appropriate privilege 27756 required to do so. The format of the *<value>* shall be as described in **pax Extended** 27757 27758 **Header File Times** (on page 717). path The pathname of the following file(s). This record shall override the *name* and 27759 prefix fields in the following header block(s). The pax utility shall translate the 27760 pathname of the file from the UTF-8 encoding to the character set appropriate for 27761 the local file system. 27762 27763 When used in **write** or **copy** mode, *pax* shall include a *path* extended header record for each file whose pathname cannot be represented entirely with the members of 27764 the portable character set other than NUL. 27765 **realtime.** The keywords prefixed by "realtime." are reserved for future standardization. 27766 security.any The keywords prefixed by "security." are reserved for future standardization. 27767 size The size of the file in octets, expressed as a decimal number using digits from the 27768 ISO/IEC 646:1991 standard. This record shall override the *size* field in the 27769 following header block(s). When used in write or copy mode, pax shall include a 27770 size extended header record for each file with a size value greater than 8 589 934 591 27771 (octal 77 777 777 777). 27772 uid The user ID of the file owner, expressed as a decimal number using digits from the 27773 ISO/IEC 646:1991 standard. This record shall override the uid field in the 27774 27775 following header block(s). When used in write or copy mode, pax shall include a uid extended header record for each file whose owner ID is greater than 2 097 151 27776 (octal 7777777). 27777 uname The owner of the following file(s), formatted as a user name in the user database. 27778 This record shall override the *uid* and *uname* fields in the following header block(s), 27779 and any *uid* extended header record. When used in **read**, **copy**, or **list** mode, *pax* 27780 shall translate the name from the UTF-8 encoding in the header record to the 27781 character set appropriate for the user database on the receiving system. If any of 27782 the UTF-8 characters cannot be translated, and if the -oinvalid= UTF-8 option is 27783 not specified, the results are implementation-defined. When used in write or copy 27784 mode, pax shall include a uname extended header record for each file whose user 27785 name cannot be represented entirely with the letters and digits of the portable 27786 27787 character set.

If the *<value>* field is zero length, it shall delete any header block field, previously entered extended header value, or global extended header value of the same name.

If a keyword in an extended header record (or in a $-\mathbf{o}$ option-argument) overrides or deletes a corresponding field in the **ustar** header block, *pax* shall ignore the contents of that header block field.

27788

27789

27791

Unlike the **ustar** header block fields, NULs shall not delimit *<value>s*; all characters within the *<value>* field shall be considered data for the field. None of the length limitations of the **ustar** header block fields in Table 4-13 (on page 718) shall apply to the extended header records.

pax Extended Header Keyword Precedence

This section describes the precedence in which the various header records and fields and command line options are selected to apply to a file in the archive. When *pax* is used in **read** or **list** modes, it shall determine a file attribute in the following sequence:

- 1. If **-odelete**=*keyword-prefix* is used, the affected attributes shall be determined from step 7., if applicable, or ignored otherwise.
- 2. If **–o***keyword*:= is used, the affected attributes shall be ignored.
- 3. If **-o**keyword:=value is used, the affected attribute shall be assigned the value.
- 4. If there is a *typeflag* **x** extended header record, the affected attribute shall be assigned the <*value>*. When extended header records conflict, the last one given in the header shall take precedence.
- 5. If **-o**keyword=value is used, the affected attribute shall be assigned the value.
- 6. If there is a *typeflag* **g** global extended header record, the affected attribute shall be assigned the *<value>*. When global extended header records conflict, the last one given in the global header shall take precedence.
- 7. Otherwise, the attribute shall be determined from the **ustar** header block.

pax Extended Header File Times

The *pax* utility shall write an **mtime** record for each file in **write** or **copy** modes if the file's modification time cannot be represented exactly in the **ustar** header logical record described in **ustar Interchange Format**. This can occur if the time is out of **ustar** range, or if the file system of the underlying implementation supports non-integer time granularities and the time is not an integer. All of these time records shall be formatted as a decimal representation of the time in seconds since the Epoch. If a period ('.') decimal point character is present, the digits to the right of the point shall represent the units of a subsecond timing granularity, where the first digit is tenths of a second and each subsequent digit is a tenth of the previous digit. In **read** or **copy** mode, the *pax* utility shall truncate the time of a file to the greatest value that is not greater than the input header file time. In **write** or **copy** mode, the *pax* utility shall output a time exactly if it can be represented exactly as a decimal number, and otherwise shall generate only enough digits so that the same time shall be recovered if the file is extracted on a system whose underlying implementation supports the same time granularity.

ustar Interchange Format

A ustar archive tape or file shall contain a series of logical records. Each logical record shall be a fixed-size logical record of 512 octets (see below). Although this format may be thought of as being stored on 9-track industry-standard 12.7 mm (0.5 in) magnetic tape, other types of transportable media are not excluded. Each file archived shall be represented by a header logical record that describes the file, followed by zero or more logical records that give the contents of the file. At the end of the archive file there shall be two 512-octet logical records filled with binary zeros, interpreted as an end-of-archive indicator.

The logical records may be grouped for physical I/O operations, as described under the -bblocksize and -x ustar options. Each group of logical records may be written with a single operation equivalent to the write() function. On magnetic tape, the result of this write shall be a

single tape physical block. The last physical block shall always be the full size, so logical records after the two zero logical records may contain undefined data.

The header logical record shall be structured as shown in the following table. All lengths and offsets are in decimal.

Table 4-13 ustar Header Block

Field Name	Octet Offset	Length (in Octets)
name	0	100
mode	100	8
uid	108	8
gid	116	8
size	124	12
mtime	136	12
chksum	148	8
typeflag	156	1
linkname	157	100
magic	257	6
version	263	2
uname	265	32
gname	297	32
devmajor	329	8
devminor	337	8
prefix	345	155

All characters in the header logical record shall be represented in the coded character set of the ISO/IEC 646: 1991 standard. For maximum portability between implementations, names should be selected from characters represented by the portable filename character set as octets with the most significant bit zero. If an implementation supports the use of characters outside of slash and the portable filename character set in names for files, users, and groups, one or more implementation-defined encodings of these characters shall be provided for interchange purposes.

However, the *pax* utility shall never create filenames on the local system that cannot be accessed via the procedures described in IEEE Std 1003.1-2001. If a filename is found on the medium that would create an invalid filename, it is implementation-defined whether the data from the file is stored on the file hierarchy and under what name it is stored. The *pax* utility may choose to ignore these files as long as it produces an error indicating that the file is being ignored.

Each field within the header logical record is contiguous; that is, there is no padding used. Each character on the archive medium shall be stored contiguously.

The fields *magic*, *uname*, and *gname* are character strings each terminated by a NUL character. The fields *name*, *linkname*, and *prefix* are NUL-terminated character strings except when all characters in the array contain non-NUL characters including the last character. The *version* field is two octets containing the characters "00" (zero-zero). The *typeflag* contains a single character. All other fields are leading zero-filled octal numbers using digits from the ISO/IEC 646: 1991 standard IRV. Each numeric field is terminated by one or more <space> or NUL characters.

The *name* and the *prefix* fields shall produce the pathname of the file. A new pathname shall be formed, if *prefix* is not an empty string (its first character is not NUL), by concatenating *prefix* (up to the first NUL character), a slash character, and *name*; otherwise, *name* is used alone. In either case, *name* is terminated at the first NUL character. If *prefix* begins with a NUL character, it shall be ignored. In this manner, pathnames of at most 256 characters can be supported. If a pathname

does not fit in the space provided, *pax* shall notify the user of the error, and shall not store any part of the file—header or data—on the medium.

The *linkname* field, described below, shall not use the *prefix* to produce a pathname. As such, a *linkname* is limited to 100 characters. If the name does not fit in the space provided, *pax* shall notify the user of the error, and shall not attempt to store the link on the medium.

The *mode* field provides 12 bits encoded in the ISO/IEC 646:1991 standard octal digit representation. The encoded bits shall represent the following values:

Table 4-14 ustar mode Field

Bit Value	IEEE Std 1003.1-2001 Bit	Description
04 000	S_ISUID	Set UID on execution.
02 000	S_ISGID	Set GID on execution.
01 000	<reserved></reserved>	Reserved for future standardization.
00 400	S_IRUSR	Read permission for file owner class.
00 200	S_IWUSR	Write permission for file owner class.
00 100	S_IXUSR	Execute/search permission for file owner class.
00 040	S_IRGRP	Read permission for file group class.
00 020	S_IWGRP	Write permission for file group class.
00 010	S_IXGRP	Execute/search permission for file group class.
00 004	S_IROTH	Read permission for file other class.
00 002	S_IWOTH	Write permission for file other class.
00 001	S_IXOTH	Execute/search permission for file other class.

When appropriate privilege is required to set one of these mode bits, and the user restoring the files from the archive does not have the appropriate privilege, the mode bits for which the user does not have appropriate privilege shall be ignored. Some of the mode bits in the archive format are not mentioned elsewhere in this volume of IEEE Std 1003.1-2001. If the implementation does not support those bits, they may be ignored.

The *uid* and *gid* fields are the user and group ID of the owner and group of the file, respectively.

The *size* field is the size of the file in octets. If the *typeflag* field is set to specify a file to be of type 1 (a link) or 2 (a symbolic link), the *size* field shall be specified as zero. If the *typeflag* field is set to specify a file of type 5 (directory), the *size* field shall be interpreted as described under the definition of that record type. No data logical records are stored for types 1, 2, or 5. If the *typeflag* field is set to 3 (character special file), 4 (block special file), or 6 (FIFO), the meaning of the *size* field is unspecified by this volume of IEEE Std 1003.1-2001, and no data logical records shall be stored on the medium. Additionally, for type 6, the *size* field shall be ignored when reading. If the *typeflag* field is set to any other value, the number of logical records written following the header shall be (*size*+511)/512, ignoring any fraction in the result of the division.

The *mtime* field shall be the modification time of the file at the time it was archived. It is the ISO/IEC 646: 1991 standard representation of the octal value of the modification time obtained from the *stat()* function.

The *chksum* field shall be the ISO/IEC 646: 1991 standard IRV representation of the octal value of the simple sum of all octets in the header logical record. Each octet in the header shall be treated as an unsigned value. These values shall be added to an unsigned integer, initialized to zero, the precision of which is not less than 17 bits. When calculating the checksum, the *chksum* field is treated as if it were all spaces.

The typeflag field specifies the type of file archived. If a particular implementation does not recognize the type, or the user does not have appropriate privilege to create that type, the file

shall be extracted as if it were a regular file if the file type is defined to have a meaning for the *size* field that could cause data logical records to be written on the medium (see the previous description for *size*). If conversion to a regular file occurs, the *pax* utility shall produce an error indicating that the conversion took place. All of the *typeflag* fields shall be coded in the ISO/IEC 646:1991 standard IRV:

- Represents a regular file. For backwards-compatibility, a *typeflag* value of binary zero ('\0') should be recognized as meaning a regular file when extracting files from the archive. Archives written with this version of the archive file format create regular files with a *typeflag* value of the ISO/IEC 646: 1991 standard IRV '0'.
- Represents a file linked to another file, of any type, previously archived. Such files are identified by each file having the same device and file serial number. The linked-to name is specified in the *linkname* field with a NUL-character terminator if it is less than 100 octets in length.
- 2 Represents a symbolic link. The contents of the symbolic link shall be stored in the *linkname* field.
- Represent character special files and block special files respectively. In this case the *devmajor* and *devminor* fields shall contain information defining the device, the format of which is unspecified by this volume of IEEE Std 1003.1-2001. Implementations may map the device specifications to their own local specification or may ignore the entry.
- Specifies a directory or subdirectory. On systems where disk allocation is performed on a directory basis, the *size* field shall contain the maximum number of octets (which may be rounded to the nearest disk block allocation unit) that the directory may hold. A *size* field of zero indicates no such limiting. Systems that do not support limiting in this manner should ignore the *size* field.
- Specifies a FIFO special file. Note that the archiving of a FIFO file archives the existence of this file and not its contents.
- Reserved to represent a file to which an implementation has associated some highperformance attribute. Implementations without such extensions should treat this file as a regular file (type 0).
- A-Z The letters 'A' to 'Z', inclusive, are reserved for custom implementations. All other values are reserved for future versions of IEEE Std 1003.1-2001.

Attempts to archive a socket using **ustar** interchange format shall produce a diagnostic message. Handling of other file types is implementation-defined.

The *magic* field is the specification that this archive was output in this archive format. If this field contains **ustar** (the five characters from the ISO/IEC 646: 1991 standard IRV shown followed by NUL), the *uname* and *gname* fields shall contain the ISO/IEC 646: 1991 standard IRV representation of the owner and group of the file, respectively (truncated to fit, if necessary). When the file is restored by a privileged, protection-preserving version of the utility, the user and group databases shall be scanned for these names. If found, the user and group IDs contained within these files shall be used rather than the values contained within the *uid* and *gid* fields.

cpio Interchange Format

The octet-oriented **cpio** archive format shall be a series of entries, each comprising a header that describes the file, the name of the file, and then the contents of the file.

An archive may be recorded as a series of fixed-size blocks of octets. This blocking shall be used only to make physical I/O more efficient. The last group of blocks shall always be at the full size.

For the octet-oriented **cpio** archive format, the individual entry information shall be in the order indicated and described by the following table; see also the **<cpio.h>** header.

Table 4-15 Octet-Oriented cpio Archive Entry

Header Field Name	Length (in Octets)	Interpreted as
c_magic	6	Octal number
c_dev	6	Octal number
c_ino	6	Octal number
c_mode	6	Octal number
c_uid	6	Octal number
c_gid	6	Octal number
c_nlink	6	Octal number
c_rdev	6	Octal number
c_mtime	11	Octal number
c_namesize	6	Octal number
c_filesize	11	Octal number
Filename Field Name	Length	Interpreted as
c_name	c_namesize	Pathname string
File Data Field Name	Length	Interpreted as
c_filedata	c_filesize	Data

cpio Header

For each file in the archive, a header as defined previously shall be written. The information in the header fields is written as streams of the ISO/IEC 646: 1991 standard characters interpreted as octal numbers. The octal numbers shall be extended to the necessary length by appending the ISO/IEC 646: 1991 standard IRV zeros at the most-significant-digit end of the number; the result is written to the most-significant digit of the stream of octets first. The fields shall be interpreted as follows:

- *c_magic* Identify the archive as being a transportable archive by containing the identifying value "070707".
- c_{-} dev, c_{-} ino Contains values that uniquely identify the file within the archive (that is, no files contain the same pair of c_{-} dev and c_{-} ino values unless they are links to the same file). The values shall be determined in an unspecified manner.
- c_{-} mode Contains the file type and access permissions as defined in the following table.

Table 4-16 Values for cpio c_mode Field

File Permissions Name	Value	Indicates
C_IRUSR	000 400	Read by owner
C_IWUSR	000 200	Write by owner
C_IXUSR	000 100	Execute by owner
C_IRGRP	000 040	Read by group
C_IWGRP	000 020	Write by group
C_IXGRP	000 010	Execute by group
C_IROTH	000 004	Read by others
C_IWOTH	000 002	Write by others
C_IXOTH	000 001	Execute by others
C_ISUID	004 000	Set <i>uid</i>
C_ISGID	002 000	Set <i>gid</i>
C_ISVTX	001 000	Reserved
File Type Name	Value	Indicates
C_ISDIR	040 000	Directory
C_ISFIFO	010 000	FIFO
C_ISREG	0100 000	Regular file
C_ISLNK	0120 000	Symbolic link
C_ISBLK	060 000	Block special file
C_ISCHR	020 000	Character special file
C_ISSOCK	0140 000	Socket
C_ISCTG	0110 000	Reserved

Directories, FIFOs, symbolic links, and regular files shall be supported on a system conforming to this volume of IEEE Std 1003.1-2001; additional values defined previously are reserved for compatibility with existing systems. Additional file types may be supported; however, such files should not be written to archives intended to be transported to other systems.

28035 28036		types may be supported; however, such files should not be written to archives intended to be transported to other systems.
28037	c_uid	Contains the user ID of the owner.
28038	c_gid	Contains the group ID of the group.
28039 28040	c_nlink	Contains the number of links referencing the file at the time the archive was created.
28041	c_rdev	Contains implementation-defined information for character or block special files.
28042 28043	c_mtime	Contains the latest time of modification of the file at the time the archive was created.
28044	c_namesize	Contains the length of the pathname, including the terminating NUL character.
28045	c filesize	Contains the length of the file in octets. This shall be the length of the data section

Contains the length of the file in octets. This shall be the length of the data section following the header structure.

cpio Filename

The c_n ame field shall contain the pathname of the file. The length of this field in octets is the value of c_n amesize.

If a filename is found on the medium that would create an invalid pathname, it is implementation-defined whether the data from the file is stored on the file hierarchy and under what name it is stored.

All characters shall be represented in the ISO/IEC 646:1991 standard IRV. For maximum portability between implementations, names should be selected from characters represented by the portable filename character set as octets with the most significant bit zero. If an implementation supports the use of characters outside the portable filename character set in names for files, users, and groups, one or more implementation-defined encodings of these characters shall be provided for interchange purposes. However, the *pax* utility shall never create filenames on the local system that cannot be accessed via the procedures described previously in this volume of IEEE Std 1003.1-2001. If a filename is found on the medium that would create an invalid filename, it is implementation-defined whether the data from the file is stored on the local file system and under what name it is stored. The *pax* utility may choose to ignore these files as long as it produces an error indicating that the file is being ignored.

cpio File Data

Following c_name , there shall be $c_filesize$ octets of data. Interpretation of such data occurs in a manner dependent on the file. If $c_filesize$ is zero, no data shall be contained in $c_filedata$.

When restoring from an archive:

- If the user does not have the appropriate privilege to create a file of the specified type, *pax* shall ignore the entry and write an error message to standard error.
- Only regular files have data to be restored. Presuming a regular file meets any selection criteria that might be imposed on the format-reading utility by the user, such data shall be restored.
- If a user does not have appropriate privilege to set a particular mode flag, the flag shall be ignored. Some of the mode flags in the archive format are not mentioned elsewhere in this volume of IEEE Std 1003.1-2001. If the implementation does not support those flags, they may be ignored.

cpio Special Entries

FIFO special files, directories, and the trailer shall be recorded with $c_filesize$ equal to zero. For other special files, $c_filesize$ is unspecified by this volume of IEEE Std 1003.1-2001. The header for the next file entry in the archive shall be written directly after the last octet of the file entry preceding it. A header denoting the filename **TRAILER!!!** shall indicate the end of the archive; the contents of octets in the last block of the archive following such a header are undefined.

28083 EXIT STATUS

The following exit values shall be returned:

28085 0 All files were processed successfully.

28086 >0 An error occurred.

28087 CONSEQUENCES OF ERRORS

If *pax* cannot create a file or a link when reading an archive or cannot find a file when writing an archive, or cannot preserve the user ID, group ID, or file mode when the **–p** option is specified, a diagnostic message shall be written to standard error and a non-zero exit status shall be returned, but processing shall continue. In the case where *pax* cannot create a link to a file, *pax* shall not, by default, create a second copy of the file.

If the extraction of a file from an archive is prematurely terminated by a signal or error, pax may have only partially extracted the file or (if the $-\mathbf{n}$ option was not specified) may have extracted a file of the same name as that specified by the user, but which is not the file the user wanted. Additionally, the file modes of extracted directories may have additional bits from the S_IRWXU mask set as well as incorrect modification and access times.

28098 APPLICATION USAGE

The $-\mathbf{p}$ (privileges) option was invented to reconcile differences between historical *tar* and *cpio* implementations. In particular, the two utilities use $-\mathbf{m}$ in diametrically opposed ways. The $-\mathbf{p}$ option also provides a consistent means of extending the ways in which future file attributes can be addressed, such as for enhanced security systems or high-performance files. Although it may seem complex, there are really two modes that are most commonly used:

- -p e "Preserve everything". This would be used by the historical superuser, someone with all the appropriate privileges, to preserve all aspects of the files as they are recorded in the archive. The e flag is the sum of o and p, and other implementation-defined attributes.
- -p p "Preserve" the file mode bits. This would be used by the user with regular privileges who wished to preserve aspects of the file other than the ownership. The file times are preserved by default, but two other flags are offered to disable these and use the time of extraction.

The one pathname per line format of standard input precludes pathnames containing <newline>s. Although such pathnames violate the portable filename guidelines, they may exist and their presence may inhibit usage of *pax* within shell scripts. This problem is inherited from historical archive programs. The problem can be avoided by listing filename arguments on the command line instead of on standard input.

It is almost certain that appropriate privileges are required for *pax* to accomplish parts of this volume of IEEE Std 1003.1-2001. Specifically, creating files of type block special or character special, restoring file access times unless the files are owned by the user (the -t option), or preserving file owner, group, and mode (the -**p** option) all probably require appropriate privileges.

In **read** mode, implementations are permitted to overwrite files when the archive has multiple members with the same name. This may fail if permissions on the first version of the file do not permit it to be overwritten.

The **cpio** and **ustar** formats can only support files up to 8589934592 bytes ($8*2^30$) in size.

28126 EXAMPLES

28127 The following command:

28128 pax -w - f / dev/rmt/1m.

copies the contents of the current directory to tape drive 1, medium density (assuming historical System V device naming procedures—the historical BSD device name would be /dev/rmt9).

28131 The following commands:

```
28132
             mkdir newdir
28133
             pax -rw olddir newdir
28134
             copy the olddir directory hierarchy to newdir.
             pax -r -s ',^//*usr//*,,' -f a.pax
28135
             reads the archive a.pax, with all files rooted in /usr in the archive extracted relative to the current
28136
28137
             directory.
28138
             Using the option:
             -o listopt="%M %(atime)T %(size)D %(name)s"
28139
             overrides the default output description in Standard Output and instead writes:
28140
             -rw-rw--- Jan 12 15:53 1492 /usr/foo/bar
28141
             Using the options:
28142
28143
             -o listopt='%L\t%(size)D\n%.7' \
28144
             -o listopt='(name)s\n%(ctime)T\n%T'
             overrides the default output description in Standard Output and instead writes:
28145
             /usr/foo/bar -> /tmp
28146
                                          1492
             /usr/fo
28147
             Jan 12 1991
28148
28149
             Jan 31 15:53
```

28150 RATIONALE

The *pax* utility was new for the ISO POSIX-2: 1993 standard. It represents a peaceful compromise between advocates of the historical *tar* and *cpio* utilities.

A fundamental difference between *cpio* and *tar* was in the way directories were treated. The *cpio* utility did not treat directories differently from other files, and to select a directory and its contents required that each file in the hierarchy be explicitly specified. For *tar*, a directory matched every file in the file hierarchy it rooted.

The pax utility offers both interfaces; by default, directories map into the file hierarchy they root. The $-\mathbf{d}$ option causes pax to skip any file not explicitly referenced, as cpio historically did. The tar -style behavior was chosen as the default because it was believed that this was the more common usage and because tar is the more commonly available interface, as it was historically provided on both System V and BSD implementations.

The data interchange format specification in this volume of IEEE Std 1003.1-2001 requires that processes with "appropriate privileges" shall always restore the ownership and permissions of extracted files exactly as archived. If viewed from the historic equivalence between superuser and "appropriate privileges", there are two problems with this requirement. First, users running as superusers may unknowingly set dangerous permissions on extracted files. Second, it is needlessly limiting, in that superusers cannot extract files and own them as superuser unless the archive was created by the superuser. (It should be noted that restoration of ownerships and permissions for the superuser, by default, is historical practice in *cpio*, but not in *tar*.) In order to avoid these two problems, the *pax* specification has an additional "privilege" mechanism, the $-\mathbf{p}$ option. Only a *pax* invocation with the privileges needed, and which has the $-\mathbf{p}$ option set using the e specification character, has the "appropriate privilege" to restore full ownership and permission information.

Note also that this volume of IEEE Std 1003.1-2001 requires that the file ownership and access permissions shall be set, on extraction, in the same fashion as the *creat()* function when provided

with the mode stored in the archive. This means that the file creation mask of the user is applied to the file permissions.

Users should note that directories may be created by *pax* while extracting files with permissions that are different from those that existed at the time the archive was created. When extracting sensitive information into a directory hierarchy that no longer exists, users are encouraged to set their file creation mask appropriately to protect these files during extraction.

The table of contents output is written to standard output to facilitate pipeline processing.

An early proposal had hard links displaying for all pathnames. This was removed because it complicates the output of the case where $-\mathbf{v}$ is not specified and does not match historical *cpio* usage. The hard-link information is available in the $-\mathbf{v}$ display.

The description of the $-\mathbf{l}$ option allows implementations to make hard links to symbolic links. IEEE Std 1003.1-2001 does not specify any way to create a hard link to a symbolic link, but many implementations provide this capability as an extension. If there are hard links to symbolic links when an archive is created, the implementation is required to archive the hard link in the archive (unless $-\mathbf{H}$ or $-\mathbf{L}$ is specified). When in **read** mode and in **copy** mode, implementations supporting hard links to symbolic links should use them when appropriate.

The archive formats inherited from the POSIX.1-1990 standard have certain restrictions that have been brought along from historical usage. For example, there are restrictions on the length of pathnames stored in the archive. When *pax* is used in **copy**(**–rw**) mode (copying directory hierarchies), the ability to use extensions from the **–xpax** format overcomes these restrictions.

The default *blocksize* value of 5 120 bytes for *cpio* was selected because it is one of the standard block-size values for *cpio*, set when the $-\mathbf{B}$ option is specified. (The other default block-size value for *cpio* is 512 bytes, and this was considered to be too small.) The default block value of 10 240 bytes for *tar* was selected because that is the standard block-size value for BSD *tar*. The maximum block size of 32 256 bytes (2^{15} –512 bytes) is the largest multiple of 512 bytes that fits into a signed 16-bit tape controller transfer register. There are known limitations in some historical systems that would prevent larger blocks from being accepted. Historical values were chosen to improve compatibility with historical scripts using *dd* or similar utilities to manipulate archives. Also, default block sizes for any file type other than character special file has been deleted from this volume of IEEE Std 1003.1-2001 as unimportant and not likely to affect the structure of the resulting archive.

Implementations are permitted to modify the block-size value based on the archive format or the device to which the archive is being written. This is to provide implementations with the opportunity to take advantage of special types of devices, and it should not be used without a great deal of consideration as it almost certainly decreases archive portability.

The intended use of the $-\mathbf{n}$ option was to permit extraction of one or more files from the archive without processing the entire archive. This was viewed by the standard developers as offering significant performance advantages over historical implementations. The $-\mathbf{n}$ option in early proposals had three effects; the first was to cause special characters in patterns to not be treated specially. The second was to cause only the first file that matched a pattern to be extracted. The third was to cause pax to write a diagnostic message to standard error when no file was found matching a specified pattern. Only the second behavior is retained by this volume of IEEE Std 1003.1-2001, for many reasons. First, it is in general not acceptable for a single option to have multiple effects. Second, the ability to make pattern matching characters act as normal characters is useful for parts of pax other than file extraction. Third, a finer degree of control over the special characters is useful because users may wish to normalize only a single special character in a single filename. Fourth, given a more general escape mechanism, the previous behavior of the $-\mathbf{n}$ option can be easily obtained using the $-\mathbf{s}$ option or a sed script. Finally,

writing a diagnostic message when a pattern specified by the user is unmatched by any file is useful behavior in all cases.

In this version, the $-\mathbf{n}$ was removed from the **copy** mode synopsis of *pax*; it is inapplicable because there are no pattern operands specified in this mode.

There is another method than *pax* for copying subtrees in IEEE Std 1003.1-2001 described as part of the *cp* utility. Both methods are historical practice: *cp* provides a simpler, more intuitive interface, while *pax* offers a finer granularity of control. Each provides additional functionality to the other; in particular, *pax* maintains the hard-link structure of the hierarchy while *cp* does not. It is the intention of the standard developers that the results be similar (using appropriate option combinations in both utilities). The results are not required to be identical; there seemed insufficient gain to applications to balance the difficulty of implementations having to guarantee that the results would be exactly identical.

A single archive may span more than one file. It is suggested that implementations provide informative messages to the user on standard error whenever the archive file is changed.

The $-\mathbf{d}$ option (do not create intermediate directories not listed in the archive) found in early proposals was originally provided as a complement to the historic $-\mathbf{d}$ option of *cpio*. It has been deleted.

The -s option in early proposals specified a subset of the substitution command from the ed utility. As there was no reason for only a subset to be supported, the -s option is now compatible with the current ed specification. Since the delimiter can be any non-null character, the following usage with single spaces is valid:

```
pax -s " foo bar " ...
```

 The **-t** description is worded so as to note that this may cause the access time update caused by some other activity (which occurs while the file is being read) to be overwritten.

The default behavior of *pax* with regard to file modification times is the same as historical implementations of *tar*. It is not the historical behavior of *cpio*.

Because the -i option uses /dev/tty, utilities without a controlling terminal are not able to use this option.

The -y option, found in early proposals, has been deleted because a line containing a single period for the -i option has equivalent functionality. The special lines for the -i option (a single period and the empty line) are historical practice in *cpio*.

In early drafts, a —echarmap option was included to increase portability of files between systems using different coded character sets. This option was omitted because it was apparent that consensus could not be formed for it. In this version, the use of UTF-8 should be an adequate substitute.

The $-\mathbf{k}$ option was added to address international concerns about the dangers involved in the character set transformations of $-\mathbf{e}$ (if the target character set were different from the source, the filenames might be transformed into names matching existing files) and also was made more general to protect files transferred between file systems with different {NAME_MAX} values (truncating a filename on a smaller system might also inadvertently overwrite existing files). As stated, it prevents any overwriting, even if the target file is older than the source. This version adds more granularity of options to solve this problem by introducing the $-\mathbf{oinvalid} = \text{option}$ —specifically the UTF-8 action. (Note that an existing file that is named with a UTF-8 encoding is still subject to overwriting in this case. The $-\mathbf{k}$ option closes that loophole.)

Some of the file characteristics referenced in this volume of IEEE Std 1003.1-2001 might not be supported by some archive formats. For example, neither the **tar** nor **cpio** formats contain the

file access time. For this reason, the e specification character has been provided, intended to cause all file characteristics specified in the archive to be retained.

It is required that extracted directories, by default, have their access and modification times and permissions set to the values specified in the archive. This has obvious problems in that the directories are almost certainly modified after being extracted and that directory permissions may not permit file creation. One possible solution is to create directories with the mode specified in the archive, as modified by the *umask* of the user, with sufficient permissions to allow file creation. After all files have been extracted, *pax* would then reset the access and modification times and permissions as necessary.

The list-mode formatting description borrows heavily from the one defined by the *printf* utility. However, since there is no separate operand list to get conversion arguments, the format was extended to allow specifying the name of the conversion argument as part of the conversion specification.

The T conversion specifier allows time fields to be displayed in any of the date formats. Unlike the ls utility, pax does not adjust the format when the date is less than six months in the past. This makes parsing the output more predictable.

The D conversion specifier handles the ability to display the major/minor or file size, as with ls, by using -8 (size) D.

The \bot conversion specifier handles the *ls* display for symbolic links.

Conversion specifiers were added to generate existing known types used for *ls*.

pax Interchange Format

The new POSIX data interchange format was developed primarily to satisfy international concerns that the **ustar** and **cpio** formats did not provide for file, user, and group names encoded in characters outside a subset of the ISO/IEC 646: 1991 standard. The standard developers realized that this new POSIX data interchange format should be very extensible because there were other requirements they foresaw in the near future:

- Support international character encodings and locale information
- Support security information (ACLs, and so on)
- Support future file types, such as realtime or contiguous files
- Include data areas for implementation use
- Support systems with words larger than 32 bits and timers with subsecond granularity

The following were not goals for this format because these are better handled by separate utilities or are inappropriate for a portable format:

- Encryption
- Compression
- Data translation between locales and codesets
- *inode* storage

The format chosen to support the goals is an extension of the **ustar** format. Of the two formats previously available, only the **ustar** format was selected for extensions because:

 It was easier to extend in an upwards-compatible way. It offered version flags and header block type fields with room for future standardization. The cpio format, while possessing a more flexible file naming methodology, could not be extended without breaking some

 theoretical implementation or using a dummy filename that could be a legitimate filename.

• Industry experience since the original "tar wars" fought in developing the ISO POSIX-1 standard has clearly been in favor of the **ustar** format, which is generally the default output format selected for *pax* implementations on new systems.

The new format was designed with one additional goal in mind: reasonable behavior when an older *tar* or *pax* utility happened to read an archive. Since the POSIX.1-1990 standard mandated that a "format-reading utility" had to treat unrecognized *typeflag* values as regular files, this allowed the format to include all the extended information in a pseudo-regular file that preceded each real file. An option is given that allows the archive creator to set up reasonable names for these files on the older systems. Also, the normative text suggests that reasonable file access values be used for this **ustar** header block. Making these header files inaccessible for convenient reading and deleting would not be reasonable. File permissions of 600 or 700 are suggested.

The **ustar** *typeflag* field was used to accommodate the additional functionality of the new format rather than magic or version because the POSIX.1-1990 standard (and, by reference, the previous version of *pax*), mandated the behavior of the format-reading utility when it encountered an unknown *typeflag*, but was silent about the other two fields.

Early proposals of the first revision to IEEE Std 1003.1-2001 contained a proposed archive format that was based on compatibility with the standard for tape files (ISO 1001, similar to the format used historically on many mainframes and minicomputers). This format was overly complex and required considerable overhead in volume and header records. Furthermore, the standard developers felt that it would not be acceptable to the community of POSIX developers, so it was later changed to be a format more closely related to historical practice on POSIX systems.

The prefix and name split of pathnames in **ustar** was replaced by the single path extended header record for simplicity.

The concept of a global extended header (*typeflagg*) was controversial. If this were applied to an archive being recorded on magnetic tape, a few unreadable blocks at the beginning of the tape could be a serious problem; a utility attempting to extract as many files as possible from a damaged archive could lose a large percentage of file header information in this case. However, if the archive were on a reliable medium, such as a CD-ROM, the global extended header offers considerable potential size reductions by eliminating redundant information. Thus, the text warns against using the global method for unreliable media and provides a method for implanting global information in the extended header for each file, rather than in the *typeflag g* records.

No facility for data translation or filtering on a per-file basis is included because the standard developers could not invent an interface that would allow this in an efficient manner. If a filter, such as encryption or compression, is to be applied to all the files, it is more efficient to apply the filter to the entire archive as a single file. The standard developers considered interfaces that would invoke a shell script for each file going into or out of the archive, but the system overhead in this approach was considered to be too high.

One such approach would be to have **filter=** records that give a pathname for an executable. When the program is invoked, the file and archive would be open for standard input/output and all the header fields would be available as environment variables or command-line arguments. The standard developers did discuss such schemes, but they were omitted from IEEE Std 1003.1-2001 due to concerns about excessive overhead. Also, the program itself would need to be in the archive if it were to be used portably.

There is currently no portable means of identifying the character set(s) used for a file in the file system. Therefore, *pax* has not been given a mechanism to generate charset records automatically. The only portable means of doing this is for the user to write the archive using the

28360 — ocharset=string command line option. This assumes that all of the files in the archive use the same encoding. The "implementation-defined" text is included to allow for a system that can identify the encodings used for each of its files.

The table of standards that accompanies the charset record description is acknowledged to be very limited. Only a limited number of character set standards is reasonable for maximal interchange. Any character set is, of course, possible by prior agreement. It was suggested that EBCDIC be listed, but it was omitted because it is not defined by a formal standard. Formal standards, and then only those with reasonably large followings, can be included here, simply as a matter of practicality. The *<value>*s represent names of officially registered character sets in the format required by the ISO 2375: 1985 standard.

The normal comma or <blank>-separated list rules are not followed in the case of keyword options to allow ease of argument parsing for *getopts*.

Further information on character encodings is in **pax Archive Character Set Encoding/Decoding** (on page 732).

The standard developers have reserved keyword name space for vendor extensions. It is suggested that the format to be used is:

VENDOR.keyword

where *VENDOR* is the name of the vendor or organization in all uppercase letters. It is further suggested that the keyword following the period be named differently than any of the standard keywords so that it could be used for future standardization, if appropriate, by omitting the *VENDOR* prefix.

The *<length>* field in the extended header record was included to make it simpler to step through the records, even if a record contains an unknown format (to a particular *pax*) with complex interactions of special characters. It also provides a minor integrity checkpoint within the records to aid a program attempting to recover files from a damaged archive.

There are no extended header versions of the *devmajor* and *devminor* fields because the unspecified format **ustar** header field should be sufficient. If they are not, vendor-specific extended keywords (such as *VENDOR.devmajor*) should be used.

Device and *i*-number labeling of files was not adopted from *cpio*; files are interchanged strictly on a symbolic name basis, as in **ustar**.

Just as with the **ustar** format descriptions, the new format makes no special arrangements for multi-volume archives. Each of the *pax* archive types is assumed to be inside a single POSIX file and splitting that file over multiple volumes (diskettes, tape cartridges, and so on), processing their labels, and mounting each in the proper sequence are considered to be implementation details that cannot be described portably.

The **pax** format is intended for interchange, not only for backup on a single (family of) systems. It is not as densely packed as might be possible for backup:

- It contains information as coded characters that could be coded in binary.
- It identifies extended records with name fields that could be omitted in favor of a fixed-field layout.
- It translates names into a portable character set and identifies locale-related information, both of which are probably unnecessary for backup.

The requirements on restoring from an archive are slightly different from the historical wording, allowing for non-monolithic privilege to bring forward as much as possible. In particular, attributes such as "high performance file" might be broadly but not universally granted while

set-user-ID or *chown*() might be much more restricted. There is no implication in IEEE Std 1003.1-2001 that the security information be honored after it is restored to the file hierarchy, in spite of what might be improperly inferred by the silence on that topic. That is a topic for another standard.

Links are recorded in the fashion described here because a link can be to any file type. It is desirable in general to be able to restore part of an archive selectively and restore all of those files completely. If the data is not associated with each link, it is not possible to do this. However, the data associated with a file can be large, and when selective restoration is not needed, this can be a significant burden. The archive is structured so that files that have no associated data can always be restored by the name of any link name of any link, and the user may choose whether data is recorded with each instance of a file that contains data. The format permits mixing of both types of links in a single archive; this can be done for special needs, and pax is expected to interpret such archives on input properly, despite the fact that there is no pax option that would force this mixed case on output. (When **–o linkdata** is used, the output must contain the duplicate data, but the implementation is free to include it or omit it when **–o linkdata** is not used.)

The time values are included as extended header records for those implementations needing more than the eleven octal digits allowed by the **ustar** format. Portable file timestamps cannot be negative. If *pax* encounters a file with a negative timestamp in **copy** or **write** mode, it can reject the file, substitute a non-negative timestamp, or generate a non-portable timestamp with a leading '-'. Even though some implementations can support finer file-time granularities than seconds, the normative text requires support only for seconds since the Epoch because the ISO POSIX-1 standard states them that way. The **ustar** format includes only *mtime*; the new format adds *atime* and *ctime* for symmetry. The *atime* access time restored to the file system will be affected by the -**p a** and -**p e** options. The *ctime* creation time (actually *inode* modification time) is described with "appropriate privilege" so that it can be ignored when writing to the file system. POSIX does not provide a portable means to change file creation time. Nothing is intended to prevent a non-portable implementation of *pax* from restoring the value.

The *gid*, *size*, and *uid* extended header records were included to allow expansion beyond the sizes specified in the regular *tar* header. New file system architectures are emerging that will exhaust the 12-digit size field. There are probably not many systems requiring more than 8 digits for user and group IDs, but the extended header values were included for completeness, allowing overrides for all of the decimal values in the *tar* header.

The standard developers intended to describe the effective results of *pax* with regard to file ownerships and permissions; implementations are not restricted in timing or sequencing the restoration of such, provided the results are as specified.

Much of the text describing the extended headers refers to use in "write or copy modes". The copy mode references are due to the normative text: "The effect of the copy shall be as if the copied files were written to an archive file and then subsequently extracted ...". There is certainly no way to test whether *pax* is actually generating the extended headers in copy mode, but the effects must be as if it had.

pax Archive Character Set Encoding/Decoding

There is a need to exchange archives of files between systems of different native codesets. Filenames, group names, and user names must be preserved to the fullest extent possible when an archive is read on the receiving platform. Translation of the contents of files is not within the scope of the *pax* utility.

There will also be the need to represent characters that are not available on the receiving platform. These unsupported characters cannot be automatically folded to the local set of characters due to the chance of collisions. This could result in overwriting previous extracted files from the archive or pre-existing files on the system.

For these reasons, the codeset used to represent characters within the extended header records of the *pax* archive must be sufficiently rich to handle all commonly used character sets. The fields requiring translation include, at a minimum, filenames, user names, group names, and link pathnames. Implementations may wish to have localized extended keywords that use non-portable characters.

The standard developers considered the following options:

- The archive creator specifies the well-defined name of the source codeset. The receiver must then recognize the codeset name and perform the appropriate translations to the destination codeset.
- The archive creator includes within the archive the character mapping table for the source codeset used to encode extended header records. The receiver must then read the character mapping table and perform the appropriate translations to the destination codeset.
- The archive creator translates the extended header records in the source codeset into a canonical form. The receiver must then perform the appropriate translations to the destination codeset.

The approach that incorporates the name of the source codeset poses the problem of codeset name registration, and makes the archive useless to *pax* archive decoders that do not recognize that codeset.

Because parts of an archive may be corrupted, the standard developers felt that including the character map of the source codeset was too fragile. The loss of this one key component could result in making the entire archive useless. (The difference between this and the global extended header decision was that the latter has a workaround—duplicating extended header records on unreliable media—but this would be too burdensome for large character set maps.)

Both of the above approaches also put an undue burden on the *pax* archive receiver to handle the cross-product of all source and destination codesets.

To simplify the translation from the source codeset to the canonical form and from the canonical form to the destination codeset, the standard developers decided that the internal representation should be a stateless encoding. A stateless encoding is one where each codepoint has the same meaning, without regard to the decoder being in a specific state. An example of a stateful encoding would be the Japanese Shift-JIS; an example of a stateless encoding would be the ISO/IEC 646: 1991 standard (equivalent to 7-bit ASCII).

For these reasons, the standard developers decided to adopt a canonical format for the representation of file information strings. The obvious, well-endorsed candidate is the ISO/IEC 10646-1:2000 standard (based in part on Unicode), which can be used to represent the characters of virtually all standardized character sets. The standard developers initially agreed upon using UCS2 (16-bit Unicode) as the internal representation. This repertoire of characters provides a sufficiently rich set to represent all commonly-used codesets.

However, the standard developers found that the 16-bit Unicode representation had some problems. It forced the issue of standardizing byte ordering. The 2-byte length of each character made the extended header records twice as long for the case of strings coded entirely from historical 7-bit ASCII. For these reasons, the standard developers chose the UTF-8 defined in the ISO/IEC 10646-1:2000 standard. This multi-byte representation encodes UCS2 or UCS4 characters reliably and deterministically, eliminating the need for a canonical byte ordering. In addition, NUL octets and other characters possibly confusing to POSIX file systems do not appear, except to represent themselves. It was realized that certain national codesets take up more space after the encoding, due to their placement within the UCS range; it was felt that the usefulness of the encoding of the names outweighs the disadvantage of size increase for file, user, and group names.

The encoding of UTF-8 is as follows:

```
UCS4 Hex Encoding
                              UTF-8 Binary Encoding
28504
28505
           00000000-0000007F
                               0xxxxxxx
28506
           00000080-000007FF
                              110xxxxx 10xxxxxx
28507
           00000800-0000FFFF
                               1110xxxx 10xxxxxx 10xxxxxx
           00010000-001FFFFF
                               11110xxx 10xxxxxx 10xxxxxx 10xxxxxx
28508
           00200000-03FFFFFF
                               111110xx 10xxxxxx 10xxxxxx 10xxxxxx 10xxxxxx
28509
           0400000-7FFFFFF
                               1111110x 10xxxxxx 10xxxxxx 10xxxxxx 10xxxxxx
28510
```

where each 'x' represents a bit value from the character being translated.

ustar Interchange Format

The description of the **ustar** format reflects numerous enhancements over pre-1988 versions of the historical *tar* utility. The goal of these changes was not only to provide the functional enhancements desired, but also to retain compatibility between new and old versions. This compatibility has been retained. Archives written using the old archive format are compatible with the new format.

Implementors should be aware that the previous file format did not include a mechanism to archive directory type files. For this reason, the convention of using a filename ending with slash was adopted to specify a directory on the archive.

The total size of the *name* and *prefix* fields have been set to meet the minimum requirements for {PATH_MAX}. If a pathname will fit within the *name* field, it is recommended that the pathname be stored there without the use of the *prefix* field. Although the name field is known to be too small to contain {PATH_MAX} characters, the value was not changed in this version of the archive file format to retain backwards-compatibility, and instead the prefix was introduced. Also, because of the earlier version of the format, there is no way to remove the restriction on the *linkname* field being limited in size to just that of the *name* field.

The *size* field is required to be meaningful in all implementation extensions, although it could be zero. This is required so that the data blocks can always be properly counted.

It is suggested that if device special files need to be represented that cannot be represented in the standard format, that one of the extension types (**A-Z**) be used, and that the additional information for the special file be represented as data and be reflected in the *size* field.

Attempting to restore a special file type, where it is converted to ordinary data and conflicts with an existing filename, need not be specially detected by the utility. If run as an ordinary user, *pax* should not be able to overwrite the entries in, for example, /dev in any case (whether the file is converted to another type or not). If run as a privileged user, it should be able to do so, and it would be considered a bug if it did not. The same is true of ordinary data files and similarly

named special files; it is impossible to anticipate the needs of the user (who could really intend to overwrite the file), so the behavior should be predictable (and thus regular) and rely on the protection system as required.

The value 7 in the *typeflag* field is intended to define how contiguous files can be stored in a **ustar** archive. IEEE Std 1003.1-2001 does not require the contiguous file extension, but does define a standard way of archiving such files so that all conforming systems can interpret these file types in a meaningful and consistent manner. On a system that does not support extended file types, the *pax* utility should do the best it can with the file and go on to the next.

The file protection modes are those conventionally used by the *ls* utility. This is extended beyond the usage in the ISO POSIX-2 standard to support the "shared text" or "sticky" bit. It is intended that the conformance document should not document anything beyond the existence of and support of such a mode. Further extensions are expected to these bits, particularly with overloading the set-user-ID and set-group-ID flags.

cpio Interchange Format

The reference to appropriate privilege in the **cpio** format refers to an error on standard output; the **ustar** format does not make comparable statements.

The model for this format was the historical System V *cpio*–c data interchange format. This model documents the portable version of the **cpio** format and not the binary version. It has the flexibility to transfer data of any type described within IEEE Std 1003.1-2001, yet is extensible to transfer data types specific to extensions beyond IEEE Std 1003.1-2001 (for example, contiguous files). Because it describes existing practice, there is no question of maintaining upwards-compatibility.

cpio Header

There has been some concern that the size of the c_ino field of the header is too small to handle those systems that have very large inode numbers. However, the c_ino field in the header is used strictly as a hard-link resolution mechanism for archives. It is not necessarily the same value as the inode number of the file in the location from which that file is extracted.

The name *c_magic* is based on historical usage.

cpio Filename

For most historical implementations of the *cpio* utility, {PATH_MAX} octets can be used to describe the pathname without the addition of any other header fields (the NUL character would be included in this count). {PATH_MAX} is the minimum value for pathname size, documented as 256 bytes. However, an implementation may use $c_namesize$ to determine the exact length of the pathname. With the current description of the <cpio.h> header, this pathname size can be as large as a number that is described in six octal digits.

Two values are documented under the c_mode field values to provide for extensibility for known file types:

0110 000 Reserved for contiguous files. The implementation may treat the rest of the information for this archive like a regular file. If this file type is undefined, the implementation may create the file as a regular file.

This provides for extensibility of the **cpio** format while allowing for the ability to read old archives. Files of an unknown type may be read as "regular files" on some implementations. On a system that does not support extended file types, the *pax* utility should do the best it can with the file and go on to the next.

28582 FUTURE DIRECTIONS

28583 None.

28584 SEE ALSO

Chapter 2 (on page 29), *cp*, *ed*, *getopts*, *ls*, *printf*, the Base Definitions volume of IEEE Std 1003.1-2001, <**cpio.h**>, the System Interfaces volume of IEEE Std 1003.1-2001, *chown*(), *creat*(), *mkdir*(), *mkfifo*(), *stat*(), *utime*(), *write*()

28588 CHANGE HISTORY

First released in Issue 4.

28590 Issue 5

A note is added to the APPLICATION USAGE indicating that the **cpio** and **tar** formats can only support files up to 8 gigabytes in size.

28593 **Issue 6**

28594

28595

28597

28598

28599

28600

The *pax* utility is aligned with the IEEE P1003.2b draft standard:

- Support has been added for symbolic links in the options and interchange formats.
- A new format has been devised, based on extensions to ustar.
 - References to the "extended" tar and cpio formats derived from the POSIX.1-1990 standard have been changed to remove the "extended" adjective because this could cause confusion with the extended *tar* header added in this revision. (All references to *tar* are actually to ustar.)
- The *TZ* entry is added to the ENVIRONMENT VARIABLES section.
- IEEE PASC Interpretation 1003.2 #168 is applied, clarifying that *mkdir()* and *mkfifo()* calls can ignore an [EEXIST] error when extracting an archive.
- IEEE PASC Interpretation 1003.2 #180 is applied, clarifying how extracted files are created when in **read** mode.
- 28606 IEEE PASC Interpretation 1003.2 #181 is applied, clarifying the description of the -t option.
- 28607 IEEE PASC Interpretation 1003.2 #195 is applied.
- IEEE PASC Interpretation 1003.2 #206 is applied, clarifying the handling of links for the –H, –L, and –l options.
- IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/35 is applied, adding the process ID of the *pax* process into certain fields. This change provides a method for the implementation to ensure that different instances of *pax* extracting a file named /a/b/foo will not collide when processing the extended header information associated with foo.
- IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/36 is applied, changing -x **B** to -x pax in the OPTIONS section.

pr Utilities

```
28616 NAME
28617 pr — print files
28618 SYNOPSIS
28619 pr [+page] [-co]u
```

DESCRIPTION

 The *pr* utility is a printing and pagination filter. If multiple input files are specified, each shall be read, formatted, and written to standard output. By default, the input shall be separated into 66-line pages, each with:

- A 5-line header that includes the page number, date, time, and the pathname of the file
- A 5-line trailer consisting of blank lines

If standard output is associated with a terminal, diagnostic messages shall be deferred until the *pr* utility has completed processing.

When options specifying multi-column output are specified, output text columns shall be of equal width; input lines that do not fit into a text column shall be truncated. By default, text columns shall be separated with at least one

 lank>.

28633 OPTIONS

The pr utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines, except that: the page option has a '+' delimiter; page and column can be multi-digit numbers; some of the option-arguments are optional; and some of the option-arguments cannot be specified as separate arguments from the preceding option letter. In particular, the -s option does not allow the option letter to be separated from its argument, and the options -e, -i, and -n require that both arguments, if present, not be separated from the option letter.

The following options shall be supported. In the following option descriptions, *column*, *lines*, *offset*, *page*, and *width* are positive decimal integers; *gap* is a non-negative decimal integer.

28643 +page Begin output at page number page of the formatted input.

-column Produce multi-column output that is arranged in column columns (the default shall be 1) and is written down each column in the order in which the text is received from the input file. This option should not be used with -m. The options -e and -i shall be assumed for multiple text-column output. Whether or not text columns are produced with identical vertical lengths is unspecified, but a text column shall never exceed the length of the page (see the -l option). When used with -t, use the minimum number of lines to write the output.

Modify the effect of the *-column* option so that the columns are filled across the page in a round-robin order (for example, when *column* is 2, the first input line heads column 1, the second heads column 2, the third is the second line in column 1, and so on).

Produce output that is double-spaced; append an extra <newline> following every <newline> found in the input.

-e[char][gap]

-a

 $-\mathbf{d}$

Expand each input <tab> to the next greater column position specified by the formula n^*gap+1 , where n is an integer > 0. If gap is zero or is omitted, it shall default to 8. All <tab>s in the input shall be expanded into the appropriate number of <space>s. If any non-digit character, char, is specified, it shall be used as the

Utilities **pr**

28662		input <tab>.</tab>
28663 XSI 28664 28665	−f	Use a <form-feed> for new pages, instead of the default behavior that uses a sequence of <newline>s. Pause before beginning the first page if the standard output is associated with a terminal.</newline></form-feed>
28666 28667	−F	Use a <form-feed> for new pages, instead of the default behavior that uses a sequence of <newline>s.</newline></form-feed>
28668	-h header	Use the string <i>header</i> to replace the contents of the <i>file</i> operand in the page header.
28669 28670 28671 28672 28673	-i[char][gap]	In output, replace multiple \langle space \rangle s with \langle tab \rangle s wherever two or more adjacent \langle space \rangle s reach column positions $gap+1$, 2^* $gap+1$, 3^* $gap+1$, and so on. If gap is zero or is omitted, default tab settings at every eighth column position shall be assumed. If any non-digit character, $char$, is specified, it shall be used as the output \langle tab \rangle .
28674 28675 28676	–l lines	Override the 66-line default and reset the page length to <i>lines</i> . If <i>lines</i> is not greater than the sum of both the header and trailer depths (in lines), the pr utility shall suppress both the header and trailer, as if the $-\mathbf{t}$ option were in effect.
28677 28678 28679 28680	-m	Merge files. Standard output shall be formatted so the <i>pr</i> utility writes one line from each file specified by a <i>file</i> operand, side by side into text columns of equal fixed widths, in terms of the number of column positions. Implementations shall support merging of at least nine <i>file</i> operands.
28681 28682 28683 28684 28685 28686	-n[char][wide	Provide <i>width</i> -digit line numbering (default for <i>width</i> shall be 5). The number shall occupy the first <i>width</i> column positions of each text column of default output or each line of - m output. If <i>char</i> (any non-digit character) is given, it shall be appended to the line number to separate it from whatever follows (default for <i>char</i> is a <tab>).</tab>
28687 28688 28689	− o offset	Each line of output shall be preceded by offset <space>s. If the $-\mathbf{o}$ option is not specified, the default offset shall be zero. The space taken is in addition to the output line width (see the $-\mathbf{w}$ option below).</space>
28690 28691 28692	-p	Pause before beginning each page if the standard output is directed to a terminal (<i>pr</i> shall write an <alert> to standard error and wait for a <carriage-return> to be read on /dev/tty).</carriage-return></alert>
28693	-r	Write no diagnostic reports on failure to open files.
28694 28695	-s[char]	Separate text columns by the single character <i>char</i> instead of by the appropriate number of <space>s (default for <i>char</i> shall be <tab>).</tab></space>
28696 28697 28698	-t	Write neither the five-line identifying header nor the five-line trailer usually supplied for each page. Quit writing after the last line of each file without spacing to the end of the page.
28699 28700 28701 28702	−w width	Set the width of the line to <i>width</i> column positions for multiple text-column output only. If the $-\mathbf{w}$ option is not specified and the $-\mathbf{s}$ option is not specified, the default width shall be 72. If the $-\mathbf{w}$ option is not specified and the $-\mathbf{s}$ option is specified, the default width shall be 512.
28703		For single column output, input lines shall not be truncated.

pr Utilities

28704 **OPERANDS**

The following operand shall be supported:

28706 *file* A pathname of a file to be written. If no *file* operands are specified, or if a *file* 28707 operand is '-', the standard input shall be used.

28708 **STDIN**

The standard input shall be used only if no *file* operands are specified, or if a *file* operand is '-'. See the INPUT FILES section.

28711 INPUT FILES

The input files shall be text files.

The file $\langle \text{dev/tty} \text{ shall be used to read responses required by the } -\mathbf{p} \text{ option.}$

28714 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *pr*:

28716 LANG Provide a default value for the internationalization variables that are unset or null.
28717 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
28718 Internationalization Variables for the precedence of internationalization variables
28719 used to determine the values of locale categories.)

28720 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

28722 LC_CTYPE 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) and which characters are defined as printable (character class **print**). Non-printable characters are still written to standard output, but are not counted for the purpose for column-width and line-length calculations.

28727 LC_MESSAGES

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

28730 *LC_TIME* Determine the format of the date and time for use in writing header lines.

28731 XSI NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

Determine the timezone used to calculate date and time strings written in header lines. If *TZ* is unset or null, an unspecified default timezone shall be used.

28734 ASYNCHRONOUS EVENTS

28735 If *pr* receives an interrupt while writing to a terminal, it shall flush all accumulated error messages to the screen before terminating.

28737 **STDOUT**

The *pr* utility output shall be a paginated version of the original file (or files). This pagination shall be accomplished using either <form-feed>s or a sequence of <newline>s, as controlled by the **-F** or **-f** option. Page headers shall be generated unless the **-t** option is specified. The page headers shall be of the form:

28742 "\n\n%s %s Page %d\n\n\n", <output of date>, <file>, <page number>

In the POSIX locale, the *<output of date>* field, representing the date and time of last modification of the input file (or the current date and time if the input file is standard input), shall be equivalent to the output of the following command as it would appear if executed at the given time:

Utilities pr

```
28747
              date "+%b %e %H:%M %Y"
              without the trailing <newline>, if the page being written is from standard input. If the page
28748
28749
              being written is not from standard input, in the POSIX locale, the same format shall be used, but
              the time used shall be the modification time of the file corresponding to file instead of the current
28750
              time. When the LC_TIME locale category is not set to the POSIX locale, a different format and
28751
              order of presentation of this field may be used.
28752
              If the standard input is used instead of a file operand, the <file> field shall be replaced by a null
28753
28754
28755
              If the -\mathbf{h} option is specified, the <file> field shall be replaced by the header argument.
28756 STDERR
28757
              The standard error shall be used for diagnostic messages and for alerting the terminal when -\mathbf{p}
28758
              is specified.
28759 OUTPUT FILES
28760
              None.
28761 EXTENDED DESCRIPTION
              None.
28762
28763 EXIT STATUS
              The following exit values shall be returned:
28764
28765
                  Successful completion.
              >0 An error occurred.
28766
28767 CONSEQUENCES OF ERRORS
28768
              Default.
28769 APPLICATION USAGE
              None.
28771 EXAMPLES
28772
                1. Print a numbered list of all files in the current directory:
                    ls -a | pr -n -h "Files in $(pwd)."
28773
                2. Print file1 and file2 as a double-spaced, three-column listing headed by "file list":
28774
                    pr -3d -h "file list" file1 file2
28775
                3. Write file1 on file2, expanding tabs to columns 10, 19, 28, ...:
28776
                    pr -e9 -t <file1 >file2
28777
28778 RATIONALE
              This utility is one of those that does not follow the Utility Syntax Guidelines because of its
28779
28780
              historical origins. The standard developers could have added new options that obeyed the
              guidelines (and marked the old options obsolescent) or devised an entirely new utility; there are
28781
28782
```

examples of both actions in this volume of IEEE Std 1003.1-2001. Because of its widespread use by historical applications, the standard developers decided to exempt this version of pr from many of the guidelines.

Implementations are required to accept option-arguments to the -h, -l, -o, and -w options whether presented as part of the same argument or as a separate argument to pr, as suggested by the Utility Syntax Guidelines. The -n and -s options, however, are specified as in historical practice because they are frequently specified without their optional arguments. If a

 dlank>

28783

28784 28785

28786 28787

pr Utilities

were allowed before the option-argument in these cases, a *file* operand could mistakenly be interpreted as an option-argument in historical applications.

The text about the minimum number of lines in multi-column output was included to ensure that a best effort is made in balancing the length of the columns. There are known historical implementations in which, for example, 60-line files are listed by pr-2 as one column of 56 lines and a second of 4. Although this is not a problem when a full page with headers and trailers is produced, it would be relatively useless when used with -t.

Historical implementations of the *pr* utility have differed in the action taken for the **–f** option. BSD uses it as described here for the **–F** option; System V uses it to change trailing <newline>s on each page to a <form-feed> and, if standard output is a TTY device, sends an <alert> to standard error and reads a line from **/dev/tty** before the first page. There were strong arguments from both sides of this issue concerning historical practice and as a result the **–F** option was added. XSI-conformant systems support the System V historical actions for the **–f** option.

The *<output of date>* field in the *-*l format is specified only for the POSIX locale. As noted, the format can be different in other locales. No mechanism for defining this is present in this volume of IEEE Std 1003.1-2001, as the appropriate vehicle is a message catalog; that is, the format should be specified as a "message".

28806 FUTURE DIRECTIONS

28807 None.

28808 SEE ALSO

28791

28792

28793

28794

28795

28796

28797 28798

28799

28800

28801

28802

28803

28804 28805

28809 expand, lp

28810 CHANGE HISTORY

First released in Issue 2.

28812 Issue 6

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

• The -**p** option is added.

The normative text is reworded to avoid use of the term "must" for application requirements.

Utilities printf

28817 **NAME** 28818 printf — write formatted output 28819 SYNOPSIS 28820 printf format[argument...] 28821 **DESCRIPTION** The printf utility shall write formatted operands to the standard output. The argument operands 28822 shall be formatted under control of the *format* operand. 28823 28824 OPTIONS 28825 None. 28826 OPERANDS 28827 The following operands shall be supported: format A string describing the format to use to write the remaining operands. See the 28828 EXTENDED DESCRIPTION section. 28829 The strings to be written to standard output, under the control of *format*. See the 28830 argument EXTENDED DESCRIPTION section. 28831 28832 **STDIN** 28833 Not used. 28834 INPUT FILES None. 28835 28836 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *printf*: 28837 LANG 28838 Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 28839 Internationalization Variables for the precedence of internationalization variables 28840 used to determine the values of locale categories.) 28841 LC_ALL If set to a non-empty string value, override the values of all the other 28842 internationalization variables. 28843 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 28844 28845 characters (for example, single-byte as opposed to multi-byte characters in 28846 arguments). LC MESSAGES 28847 28848 Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. 28849 LC_NUMERIC 28850 Determine the locale for numeric formatting. It shall affect the format of numbers 28851 written using the e, E, f, g, and G conversion specifier characters (if supported). 28852 **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 28853 XSI 28854 ASYNCHRONOUS EVENTS Default. 28855 28856 **STDOUT**

See the EXTENDED DESCRIPTION section.

printf Utilities

28858 STDERR

The standard error shall be used only for diagnostic messages.

28860 OUTPUT FILES

28861 None.

28862 EXTENDED DESCRIPTION

The *format* operand shall be used as the *format* string described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5, File Format Notation with the following exceptions:

- 1. A <space> in the format string, in any context other than a flag of a conversion specification, shall be treated as an ordinary character that is copied to the output.
- 2. A ' Δ ' character in the format string shall be treated as a ' Δ ' character, not as a <space>.
- 3. In addition to the escape sequences shown in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5, File Format Notation ('\\', '\a', '\b', '\f', '\n', '\r', '\t', '\t
- 4. The implementation shall not precede or follow output from the d or u conversion specifiers with
blank>s not specified by the *format* operand.
- 5. The implementation shall not precede output from the o conversion specifier with zeros not specified by the *format* operand.
- 6. The e, E, f, g, and G conversion specifiers need not be supported.
- 7. An additional conversion specifier character, b, shall be supported as follows. The argument shall be taken to be a string that may contain backslash-escape sequences. The following backslash-escape sequences shall be supported:
 - The escape sequences listed in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5, File Format Notation ('\\', '\a', '\b', '\f', '\n', '\r', '\t', '\v'), which shall be converted to the characters they represent
 - "\oddd", where *ddd* is a zero, one, two, or three-digit octal number that shall be converted to a byte with the numeric value specified by the octal number
 - '\c', which shall not be written and shall cause *printf* to ignore any remaining characters in the string operand containing it, any remaining string operands, and any additional characters in the *format* operand

The interpretation of a backslash followed by any other sequence of characters is unspecified.

Bytes from the converted string shall be written until the end of the string or the number of bytes indicated by the precision specification is reached. If the precision is omitted, it shall be taken to be infinite, so all bytes up to the end of the converted string shall be written.

- 8. For each conversion specification that consumes an argument, the next argument operand shall be evaluated and converted to the appropriate type for the conversion as specified below.
- 9. The *format* operand shall be reused as often as necessary to satisfy the argument operands. Any extra c or s conversion specifiers shall be evaluated as if a null string argument were supplied; other extra conversion specifications shall be evaluated as if a zero argument were supplied. If the *format* operand contains no conversion specifications and *argument* operands are present, the results are unspecified.

Utilities printf

10. If a character sequence in the *format* operand begins with a '%' character, but does not form a valid conversion specification, the behavior is unspecified.

The *argument* operands shall be treated as strings if the corresponding conversion specifier is b, c, or s; otherwise, it shall be evaluated as a C constant, as described by the ISO C standard, with the following extensions:

- A leading plus or minus sign shall be allowed.
- If the leading character is a single-quote or double-quote, the value shall be the numeric value in the underlying codeset of the character following the single-quote or double-quote.

If an argument operand cannot be completely converted into an internal value appropriate to the corresponding conversion specification, a diagnostic message shall be written to standard error and the utility shall not exit with a zero exit status, but shall continue processing any remaining operands and shall write the value accumulated at the time the error was detected to standard output.

It is not considered an error if an argument operand is not completely used for a c or s conversion or if a string operand's first or second character is used to get the numeric value of a character.

28917 EXIT STATUS

The following exit values shall be returned:

28919 0 Successful completion.

28920 >0 An error occurred.

28921 CONSEQUENCES OF ERRORS

28922 Default.

28923 APPLICATION USAGE

The floating-point formatting conversion specifications of *printf*() are not required because all arithmetic in the shell is integer arithmetic. The *awk* utility performs floating-point calculations and provides its own **printf** function. The *bc* utility can perform arbitrary-precision floating-point arithmetic, but does not provide extensive formatting capabilities. (This *printf* utility cannot really be used to format *bc* output; it does not support arbitrary precision.) Implementations are encouraged to support the floating-point conversions as an extension.

Note that this *printf* utility, like the *printf*() function defined in the System Interfaces volume of IEEE Std 1003.1-2001 on which it is based, makes no special provision for dealing with multibyte characters when using the %c conversion specification or when a precision is specified in a %b or %s conversion specification. Applications should be extremely cautious using either of these features when there are multi-byte characters in the character set.

No provision is made in this volume of IEEE Std 1003.1-2001 which allows field widths and precisions to be specified as '*' since the '*' can be replaced directly in the *format* operand using shell variable substitution. Implementations can also provide this feature as an extension if they so choose.

Hexadecimal character constants as defined in the ISO C standard are not recognized in the *format* operand because there is no consistent way to detect the end of the constant. Octal character constants are limited to, at most, three octal digits, but hexadecimal character constants are only terminated by a non-hex-digit character. In the ISO C standard, the "##" concatenation operator can be used to terminate a constant and follow it with a hexadecimal character to be written. In the shell, concatenation occurs before the *printf* utility has a chance to parse the end of the hexadecimal constant.

printf Utilities

The %b conversion specification is not part of the ISO C standard; it has been added here as a portable way to process backslash escapes expanded in string operands as provided by the *echo* utility. See also the APPLICATION USAGE section of *echo* (on page 333) for ways to use *printf* as a replacement for all of the traditional versions of the *echo* utility.

If an argument cannot be parsed correctly for the corresponding conversion specification, the *printf* utility is required to report an error. Thus, overflow and extraneous characters at the end of an argument being used for a numeric conversion shall be reported as errors.

28953 EXAMPLES

To alert the user and then print and read a series of prompts:

```
28955 printf "\aPlease fill in the following: \nName: "
28956 read name
28957 printf "Phone number: "
28958 read phone
```

To read out a list of right and wrong answers from a file, calculate the percentage correctly, and print them out. The numbers are right-justified and separated by a single <tab>. The percentage is written to one decimal place of accuracy:

```
28962
            while read right wrong ; do
                percent=$(echo "scale=1;($right*100)/($right+$wrong)" | bc)
28963
                printf "%2d right\t%2d wrong\t(%s%%)\n" \
28964
                     $right $wrong $percent
28965
28966
            done < database file
            The command:
28967
            printf "%5d%4d\n" 1 21 321 4321 54321
28968
            produces:
28969
28970
                1 21
28971
              3214321
```

Note that the *format* operand is used three times to print all of the given strings and that a '0' was supplied by *printf* to satisfy the last %4d conversion specification.

The *printf* utility is required to notify the user when conversion errors are detected while producing numeric output; thus, the following results would be expected on an implementation with 32-bit twos-complement integers when %d is specified as the *format* operand:

Argument	Standard Output	Diagnostic Output
5a	5	printf: "5a" not completely converted
999999999	2147483647	printf: "9999999999" arithmetic overflow
-9999999999	-2147483648	printf: "-9999999999" arithmetic overflow
ABC	0	printf: "ABC" expected numeric value

The diagnostic message format is not specified, but these examples convey the type of information that should be reported. Note that the value shown on standard output is what would be expected as the return value from the *strtol()* function as defined in the System Interfaces volume of IEEE Std 1003.1-2001. A similar correspondence exists between %u and *strtoul()* and %e, %f, and %g (if the implementation supports floating-point conversions) and *strtod()*.

Utilities printf

28990	In a locale using the ISO/IEC 646: 1991 standard as the underlying codeset, the command:				
28991	printf "%d\n" 3 +3 -3 \'3 \"+3 "'-3"				
28992	produces:				
28993	3 Numeric value of constant 3				
28994	3 Numeric value of constant 3				
28995	-3 Numeric value of constant -3				
28996	51 Numeric value of the character '3' in the ISO/IEC 646: 1991 standard codeset				
28997	43 Numeric value of the character '+' in the ISO/IEC 646: 1991 standard codeset				
28998	45 Numeric value of the character '-' in the ISO/IEC 646: 1991 standard codeset				
28999 29000 29001	Note that in a locale with multi-byte characters, the value of a character is intended to be the value of the equivalent of the wchar_t representation of the character as described in the System Interfaces volume of IEEE Std 1003.1-2001.				
29002 RATIO 29003 29004 29005 29006	TIONALE The <i>printf</i> utility was added to provide functionality that has historically been provided by <i>echo</i> . However, due to irreconcilable differences in the various versions of <i>echo</i> extant, the version has few special features, leaving those to this new <i>printf</i> utility, which is based on one in the Ninth Edition system.				
29007 29008 29009	The EXTENDED DESCRIPTION section almost exactly matches the <i>printf()</i> function in the ISO C standard, although it is described in terms of the file format notation in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 5, File Format Notation.				
29010 FUTURE DIRECTIONS					
29011	None.				
29012 SEE AL 29013	29012 SEE ALSO 29013 <i>awk</i> , <i>bc</i> , <i>echo</i> , the System Interfaces volume of IEEE Std 1003.1-2001, <i>printf</i> ()				
29014 CHANGE HISTORY					

First released in Issue 4.

prs Utilities

	NAME		CCCC CL. (DEVELORMENT)		
29017		prs — print an SCCS file (DEVELOPMENT)			
	SYNOP				
29019 X	KSI	prs [-a][-	-d dataspec][-r[SID]] file		
29020 X	KSI	prs [-e	-1] -c cutoff [-d dataspec] file		
29021 X	KSI	prs [-e	-l] -r[SID][-d dataspec]file		
29022 I	DESCR	IPTION			
29023			ty shall write to standard output parts or all of an SCCS file in a user-supplied		
29024		format.			
29025 (OPTIO				
29026			y shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2,		
29027			ax Guidelines, except that the $-\mathbf{r}$ option has an optional option-argument. This ion-argument cannot be presented as a separate argument. The following options		
29028 29029		shall be supp			
29030		-d dataspec	Specify the output data specification. The <i>dataspec</i> shall be a string consisting of		
29031		u unuspec	SCCS file <i>data keywords</i> (see Data Keywords (on page 747)) interspersed with		
29032			optional user-supplied text.		
29033		-r[SID]	Specify the SCCS identification string (SID) of a delta for which information is		
29034			desired. If no SID option-argument is specified, the SID of the most recently		
29035			created delta shall be assumed.		
29036		-е	Request information for all deltas created earlier than and including the delta		
29037			designated via the $-\mathbf{r}$ option or the date-time given by the $-\mathbf{c}$ option.		
29038		-l	Request information for all deltas created later than and including the delta		
29039			designated via the $-\mathbf{r}$ option or the date-time given by the $-\mathbf{c}$ option.		
29040		−c cutoff	Indicate the <i>cutoff</i> date-time, in the form:		
29041			YY[MM[DD[HH[MM[SS]]]]]		
29042			For the YY component, values in the range [69,99] shall refer to years 1969 to 1999		
29043			inclusive, and values in the range [00,68] shall refer to years 2000 to 2068 inclusive.		
29044 29045			Note: It is expected that in a future version of IEEE Std 1003.1-2001 the default century inferred from a 2-digit year will change. (This would apply to all		
29046			commands accepting a 2-digit year as input.)		
29047			No changes (deltas) to the SCCS file that were created after the specified cutoff		
29048			date-time shall be included in the output. Units omitted from the date-time default		
29049			to their maximum possible values; for example, -c 7502 is equivalent to		
29050			-c 750228235959.		
29051		- a	Request writing of information for both removed—that is, $delta$ $type=R$ (see		
29052			rmdel)—and existing—that is, delta type=D,—deltas. If the -a option is not		
29053			specified, information for existing deltas only shall be provided.		
	OPERA		or an area of all the areas and all		
29055			ng operand shall be supported:		
29056		file	A pathname of an existing SCCS file or a directory. If <i>file</i> is a directory, the <i>prs</i>		
29057			utility shall behave as though each file in the directory were specified as a named		
29058 29059			file, except that non-SCCS files (last component of the pathname does not begin with s .) and unreadable files shall be silently ignored.		
23033			with 5.7 and diffeatable files shall be shelly ignored.		

Utilities prs

29060 If exactly one *file* operand appears, and it is '-', the standard input shall be read; each line of the standard input shall be taken to be the name of an SCCS file to be processed. Non-SCCS files and unreadable files shall be silently ignored.

29063 STDIN

The standard input shall be a text file used only when the *file* operand is specified as '-'. Each line of the text file shall be interpreted as an SCCS pathname.

29066 INPUT FILES

29067 Any SCCS files displayed are files of an unspecified format.

29068 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *prs*:

29070 LANG Provide a default value for the internationalization variables that are unset or null.
29071 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
29072 Internationalization Variables for the precedence of internationalization variables
29073 used to determine the values of locale categories.)

29074 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

29076 *LC_CTYPE* 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).

29079 *LC_MESSAGES*

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

29082 *NLSPATH* Determine the location of message catalogs for the processing of *LC_MESSAGES*.

29083 ASYNCHRONOUS EVENTS

29084 Default.

29085 **STDOUT**

The standard output shall be a text file whose format is dependent on the data keywords specified with the $-\mathbf{d}$ option.

29088 Data Keywords

Data keywords specify which parts of an SCCS file shall be retrieved and output. All parts of an SCCS file have an associated data keyword. A data keyword may appear in a *dataspec* multiple times.

The information written by *prs* shall consist of:

- 29093 1. The user-supplied text
- 29094 2. Appropriate values (extracted from the SCCS file) substituted for the recognized data keywords in the order of appearance in the *dataspec*

The format of a data keyword value shall either be simple ('S'), in which keyword substitution is direct, or multi-line ('M').

User-supplied text shall be any text other than recognized data keywords. A <tab> shall be specified by '\t' and <newline> by '\n'. When the $-\mathbf{r}$ option is not specified, the default dataspec shall be:

29101 :PN::\n\n

prs Utilities

and the following *dataspec* shall be used for each selected delta:

:Dt:\t:DL:\nMRs:\n:MR:COMMENTS:\n:C:

29103
29104
29105
29106
29107
29108
29109
29110
29111
29112
29113
29114
29115
29116
29117
29118
29119
29120
29121
29122
29123
29124
29125
29126 29127
29127
29129
29130
29131
29132
29133
29134
29135
29136
29137
29138
29139
29140
29141
29142
29143
29144

29102

	SCCS File Data Keywords				
Keyword	Data Item	File Section	Value	Format	
:Dt:	Delta information	Delta Table	See below*	S	
:DL:	Delta line statistics	"	:Li:/:Ld:/:Lu:	S	
:Li:	Lines inserted by Delta	"	nnnnn***	S	
:Ld:	Lines deleted by Delta	"	nnnnn***	S	
:Lu:	Lines unchanged by Delta	"	nnnnn***	S	
:DT:	Delta type	"	D or R	S	
:I:	SCCS ID string (SID)	"	See below**	S	
:R:	Release number	"	nnnn	S	
:L:	Level number	"	nnnn	S	
:B:	Branch number	"	nnnn	S	
:S:	Sequence number	"	nnnn	S	
:D:	Date delta created	"	:Dy:/:Dm:/:Dd:	S	
:Dy:	Year delta created	"	nn	S	
:Dm:	Month delta created	"	nn	S	
:Dd:	Day delta created	"	nn	S	
:T:	Time delta created	"	:Th:::Tm:::Ts:	S	
:Th:	Hour delta created	"	nn	S	
:Tm:	Minutes delta created	"	nn	S	
:Ts:	Seconds delta created	"	nn	S	
:P:	Programmer who created Delta	"	logname	S	
:DS:	Delta sequence number	"	nnnn	S	
:DP:	Predecessor Delta sequence	"	nnnn	S	
	number			_	
:DI:	Sequence number of deltas	"	:Dn:/:Dx:/:Dg:	S	
	included, excluded, or ignored		ibini ibini ibgi		
:Dn:	Deltas included (sequence #)	"	:DS::DS:	S	
:Dx:	Deltas included (sequence #)	"	:DS::DS:	S	
:Dg:	Deltas ignored (sequence #)	"	:DS::DS:	S	
:MR:	MR numbers for delta	"	text	M	
:C:	Comments for delta	"	text	M	
:UN:	User names	User Names	text	M	
:FL:	Flag list	Flags	text	M	
.г.с. :Y:	Module type flag	riags	text	S	
:MF:	MR validation flag	"		S	
:MP:	MR validation program name	,,	yes or no	S	
		"	text	S	
:KF:	Keyword error, warning flag	"	yes or no		
:KV:	Keyword validation string	"	text	S	
:BF:	Branch flag	,,	yes or no	S	
:J:	Joint edit flag	,,	yes or no	S	
:LK:	Locked releases	,,	:R:	S	
:Q:	User-defined keyword	"	text	S	
:M:	Module name		text	S	

Utilities prs

150		SCCS File	Data Keywords		
151	Keyword	Data Item	File Section	Value	Forma
152	:FB:	Floor boundary	"	:R:	S
153	:CB:	Ceiling boundary	"	:R:	S
154	:Ds:	Default SID	"	:I:	S
155	:ND:	Null delta flag	"	yes or no	S
156	:FD:	File descriptive text	Comments	text	M
157	:BD:	Body	Body	text	M
158	:GB:	Gotten body	"	text	M
159	:W:	A form of <i>what</i> string	N/A	:Z::M:\t:I:	S
160	:A:	A form of what string	N/A	:Z::Y: :M: :I::Z:	S
161	:Z:	what string delimiter	N/A	@(#)	S
162	:F:	SCCS filename	N/A	text	S
163	:PN:	SCCS file pathname	N/A	text	S
160 161 162 163	:A: :Z: :F: :PN:	A form of <i>what</i> string <i>what</i> string delimiter SCCS filename	N/A N/A N/A		:Z::Y: :M: :I::Z: @ (#) text
			(DE)		
IJ		S: if the delta is a branch delta			
166	:R:.:L: if th	ne delta is not a branch delta (:	BF:= =no)		
167	*** The line s	tatistics are capped at 99 999.	For example, if 100	000 lines were un	changed
168		rision, :Lu: shall produce the v			

29169 **STDERR**

29170 The standard error shall be used only for diagnostic messages.

```
29171 OUTPUT FILES
```

29172 None.

29173 EXTENDED DESCRIPTION

29174 None.

29175 EXIT STATUS

The following exit values shall be returned:

29177 0 Successful completion.

29178 >0 An error occurred.

29179 CONSEQUENCES OF ERRORS

29180 Default.

29181 APPLICATION USAGE

29182 None.

29183 EXAMPLES

29185

29184 1. The following example:

prs -d "User Names for :F: are:\n:UN:" s.file

29186 might write to standard output:

29187 User Names for s.file are:

29188 xyz 29189 131 29190 abc

29191 2. The following example:

prs Utilities

```
29192
                   prs -d "Delta for pgm :M:: :I: - :D: By :P:" -r s.file
                   might write to standard output:
29193
                   Delta for pgm main.c: 3.7 - 77/12/01 By cas
29194
29195
               3. As a special case:
                   prs s.file
29196
29197
                   might write to standard output:
                   s.file:
29198
29199
                   <black line>
                   D 1.1 77/12/01 00:00:00 cas 1 000000/00000/00000
29200
29201
                   MRs:
29202
                   bl78-12345
                   bl79-54321
29203
29204
                   COMMENTS:
                   this is the comment line for s.file initial delta
29205
                   <black line>
29206
                   for each delta table entry of the D type. The only option allowed to be used with this
29207
29208
                   special case is the -a option.
29209 RATIONALE
             None.
29210
29211 FUTURE DIRECTIONS
29212
             None.
29213 SEE ALSO
29214
             admin, delta, get, what
29215 CHANGE HISTORY
             First released in Issue 2.
29216
29217 Issue 5
             The phrase "in which keyword substitution is followed by a <newline>" is deleted from the end
29218
             of the second paragraph of Data Keywords (on page 747).
29219
             The interpretation of the YY component of the –c cutoff argument is noted.
29220
29221 Issue 6
29222
             The normative text is reworded to emphasize the term "shall" for implementation requirements.
             The Open Group Base Resolution bwg2001-007 is applied, updating the table in STDOUT with a
29223
29224
             note that line statistics are capped at 99 999 for the :Li:, :Ld:, :Lu:, and :DL: keywords.
29225
             The Open Group Interpretation PIN4C.00009 is applied.
```

Utilities ps

```
29226 NAME
29227
                                           ps — report process status
29228 SYNOPSIS
                                           ps [-aA] [-defl] [-G grouplist] [-o format]...[-p proclist] [-t termlist]
29229 UP XSI
29230
                                            [-U userlist] [-g grouplist] [-n namelist] [-u userlist]
29231
29232 DESCRIPTION
                                           The ps utility shall write information about processes, subject to having the appropriate
29233
29234
                                           privileges to obtain information about those processes.
                                           By default, ps shall select all processes with the same effective user ID as the current user and the
29235
29236
                                           same controlling terminal as the invoker.
29237 OPTIONS
                                           The ps utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2,
29238
29239
                                           Utility Syntax Guidelines.
                                           The following options shall be supported:
29240
                                                                                     Write information for all processes associated with terminals. Implementations
29241
29242
                                                                                     may omit session leaders from this list.
                                           -\mathbf{A}
29243
                                                                                     Write information for all processes.
29244 XSI
                                           -\mathbf{d}
                                                                                     Write information for all processes, except session leaders.
                                                                                     Write information for all processes. (Equivalent to −A.)
29245 XSI
                                            -е
                                                                                     Generate a full listing. (See the STDOUT section for the contents of a full listing.)
                                           -\mathbf{f}
29246 XSI
                                            –g grouplist
                                                                                     Write information for processes whose session leaders are given in grouplist. The
29247 XSI
                                                                                     application shall ensure that the grouplist is a single argument in the form of a
29248
29249
                                                                                     <br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>

                                           -G grouplist Write information for processes whose real group ID numbers are given in
29250
                                                                                     grouplist. The application shall ensure that the grouplist is a single argument in the
29251
29252
                                                                                     form of a <blank> or comma-separated list.
                                           -\mathbf{l}
                                                                                     Generate a long listing. (See STDOUT for the contents of a long listing.)
29253 XSI
                                           -n namelist
                                                                                    Specify the name of an alternative system namelist file in place of the default. The
29254 XSI
                                                                                     name of the default file and the format of a namelist file are unspecified.
29255
                                           -o format
                                                                                     Write information according to the format specification given in format. This is
29256
                                                                                     fully described in the STDOUT section. Multiple -o options can be specified; the
29257
                                                                                     format specification shall be interpreted as the <space>-separated concatenation of
29258
29259
                                                                                     all the format option-arguments.
                                           -p proclist
                                                                                     Write information for processes whose process ID numbers are given in proclist.
29260
                                                                                     The application shall ensure that the proclist is a single argument in the form of a
29261
                                                                                     <br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>

29262
29263
                                           –t termlist
                                                                                     Write information for processes associated with terminals given in termlist. The
                                                                                     application shall ensure that the termlist is a single argument in the form of a
29264
                                                                                     <br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>
<br/>

29265
                                                                                     implementation-defined format. On XSI-conformant systems, they shall be given
29266 XSI
29267
                                                                                     in one of two forms: the device's filename (for example, tty04) or, if the device's
                                                                                     filename starts with tty, just the identifier following the characters tty (for
29268
```

ps Utilities

29269		example, "04").	
29270 XSI 29271 29272 29273 29274	–u userlist	Write information for processes whose user ID numbers or login names are given in <i>userlist</i> . The application shall ensure that the <i>userlist</i> is a single argument in the form of a <blank> or comma-separated list. In the listing, the numerical user ID shall be written unless the –f option is used, in which case the login name shall be written.</blank>	
29275 29276 29277	–U userlist	Write information for processes whose real user ID numbers or login names are given in <i>userlist</i> . The application shall ensure that the <i>userlist</i> is a single argument in the form of a <black> or comma-separated list.</black>	
29278 29279 29280	specified, th	ception of $-\mathbf{o}$ <i>format</i> , all of the options shown are used to select processes. If any are the default list shall be ignored and ps shall select the processes represented by the R of all the selection-criteria options.	
29281 OPERA 29282	N DS None.		
29283 STDIN 29284	Not used.		
29285 INPUT 29286	FILES None.		
29287 ENVIR 29288	The following environment variables shall affect the execution of <i>ps</i> :		
29289 29290 29291 29292	COLUMNS	Override the system-selected horizontal display line size, used to determine the number of text columns to display. See the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables for valid values and results when it is unset or null.	
29293 29294 29295 29296	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)	
29297 29298	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.	
29299 29300 29301	LC_CTYPE	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).	
29302 29303 29304 29305	LC_MESSA	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.	
29306	LC_TIME	Determine the format and contents of the date and time strings displayed.	
29307 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.	
29308 29309	TZ	Determine the timezone used to calculate date and time strings displayed. If TZ is unset or null, an unspecified default timezone shall be used.	

Utilities ps

29310 ASYNCHRONOUS EVENTS

29311 Default.

STDOUT

29314 XSI

When the $-\mathbf{o}$ option is not specified, the standard output format is unspecified.

On XSI-conformant systems, the output format shall be as follows. The column headings and descriptions of the columns in a *ps* listing are given below. The precise meanings of these fields are implementation-defined. The letters 'f' and 'l' (below) indicate the option (**full** or **long**) that shall cause the corresponding heading to appear; **all** means that the heading always appears. Note that these two options determine only what information is provided for a process; they do not determine which processes are listed.

29320	F	(l)	Flags (octal and additive) associated with the process.
29321	S	(l)	The state of the process.
29322	UID	(f,l)	The user ID number of the process owner; the login name is printed
29323			under the -f option.
29324	PID	(all)	The process ID of the process; it is possible to kill a process if this
29325			datum is known.
29326	PPID	(f,l)	The process ID of the parent process.
29327	C	(f,l)	Processor utilization for scheduling.
29328	PRI	(l)	The priority of the process; higher numbers mean lower priority.
29329	NI	(l)	Nice value; used in priority computation.
29330	ADDR	(l)	The address of the process.
29331	SZ	(l)	The size in blocks of the core image of the process.
29332	WCHAN	(l)	The event for which the process is waiting or sleeping; if blank, the
29333			process is running.
29334	STIME	(f)	Starting time of the process.
29335	TTY	(all)	The controlling terminal for the process.
29336	TIME	(all)	The cumulative execution time for the process.
29337	CMD	(all)	The command name; the full command name and its arguments are
29338			written under the –f option.

A process that has exited and has a parent, but has not yet been waited for by the parent, shall be marked **defunct**.

Under the option $-\mathbf{f}$, ps tries to determine the command name and arguments given when the process was created by examining memory or the swap area. Failing this, the command name, as it would appear without the option $-\mathbf{f}$, is written in square brackets.

The $-\mathbf{o}$ option allows the output format to be specified under user control.

The application shall ensure that the format specification is a list of names presented as a single argument,
 blank> or comma-separated. Each variable has a default header. The default header can be overridden by appending an equals sign and the new text of the header. The rest of the characters in the argument shall be used as the header text. The fields specified shall be written in the order specified on the command line, and should be arranged in columns in the output. The field widths shall be selected by the system to be at least as wide as the header text (default or overridden value). If the header text is null, such as $-\mathbf{o}$ user=, the field width shall be at least as wide as the default header text. If all header text fields are null, no header line shall be written.

The following names are recognized in the POSIX locale:

ps Utilities

29355 29356	ruser	The real user ID of the process. This shall be the textual user ID, if it can be obtained and the field width permits, or a decimal representation otherwise.		
29357 29358	user	The effective user ID of the process. This shall be the textual user ID, if it can be obtained and the field width permits, or a decimal representation otherwise.		
29359 29360	rgroup	The real group ID of the process. This shall be the textual group ID, if it can be obtand the field width permits, or a decimal representation otherwise.		
29361 29362	group	The effective group ID of the process. This shall be the textual group ID, if it can obtained and the field width permits, or a decimal representation otherwise.		
29363	pid	The decimal value of the process ID.		
29364	ppid	The decimal value of the parent process ID.		
29365	pgid	The decimal value of the process group ID.		
29366 29367 29368	pcpu	The ratio of CPU time used recently to CPU time available in the same period, expressed as a percentage. The meaning of "recently" in this context is unspecified. The CPU time available is determined in an unspecified manner.		
29369	VSZ	The size of the process in (virtual) memory in 1 024 byte units as a decimal integer.		
29370	nice	The decimal value of the nice value of the process; see <i>nice</i> .		
29371	etime	In the POSIX locale, the elapsed time since the process was started, in the form:		
29372		[[dd-]hh:]mm:ss		
29373 29374 29375		where <i>dd</i> shall represent the number of days, <i>hh</i> the number of hours, <i>mm</i> the number of minutes, and <i>ss</i> the number of seconds. The <i>dd</i> field shall be a decimal integer. The <i>hh</i> , <i>mm</i> , and <i>ss</i> fields shall be two-digit decimal integers padded on the left with zeros.		
29376	time	In the POSIX locale, the cumulative CPU time of the process in the form:		
29377		[dd-]hh:mm:ss		
29378		The <i>dd</i> , <i>hh</i> , <i>mm</i> , and <i>ss</i> fields shall be as described in the etime specifier.		
29379 29380	tty	The name of the controlling terminal of the process (if any) in the same format used by the <i>who</i> utility.		
29381	comm	The name of the command being executed (argv[0] value) as a string.		
29382 29383 29384 29385 29386 29387 29388	args	The command with all its arguments as a string. The implementation may truncate this value to the field width; it is implementation-defined whether any further truncation occurs. It is unspecified whether the string represented is a version of the argument list as it was passed to the command when it started, or is a version of the arguments as they may have been modified by the application. Applications cannot depend on being able to modify their argument list and having that modification be reflected in the output of <i>ps</i> .		
29389 29390		Any field need not be meaningful in all implementations. In such a case a hyphen $('-')$ should be output in place of the field value.		
29391 29392 29393	Only comm and args shall be allowed to contain blank>s; all others shall not. Any implementation-defined variables shall be specified in the system documentation along with the default header and indicating whether the field may contain blank>s.			
29394 29395	The following table specifies the default header to be used in the POSIX locale corresponding to each format specifier.			

Utilities ps

20	2	n	o
7.3	١.٦	м	n

Table 4-17 Variable Names and Default Headers in *ps*

Format	Specifier	Default Header	Format Specifier	Default Header
args		COMMAND	ppid	PPID
comm		COMMAND	rgroup	RGROUP
etime		ELAPSED	ruser	RUSER
group		GROUP	time	TIME
nice		NI	tty	TT
pcpu		%CPU	user	USER
pgid		PGID	VSZ	VSZ
pid		PID		

29406 STDERR

29407 The standard error shall be used only for diagnostic messages.

29408 OUTPUT FILES

29409 None.

29410 EXTENDED DESCRIPTION

29411 None.

29412 EXIT STATUS

The following exit values shall be returned:

29414 0 Successful completion.

29415 >0 An error occurred.

29416 CONSEQUENCES OF ERRORS

29417 Default.

29421 29422

29423

29424

29425

29426

29427

29418 APPLICATION USAGE

Things can change while *ps* is running; the snapshot it gives is only true for an instant, and might not be accurate by the time it is displayed.

The **args** format specifier is allowed to produce a truncated version of the command arguments. In some implementations, this information is no longer available when the *ps* utility is executed.

If the field width is too narrow to display a textual ID, the system may use a numeric version. Normally, the system would be expected to choose large enough field widths, but if a large number of fields were selected to write, it might squeeze fields to their minimum sizes to fit on one line. One way to ensure adequate width for the textual IDs is to override the default header for a field to make it larger than most or all user or group names.

There is no special quoting mechanism for header text. The header text is the rest of the argument. If multiple header changes are needed, multiple **–o** options can be used, such as:

```
29430 ps -o "user=User Name" -o pid=Process\ ID
```

On some implementations, especially multi-level secure systems, *ps* may be severely restricted and produce information only about child processes owned by the user.

29433 EXAMPLES

29434 The command:

29435 ps -o user, pid, ppid=MOM -o args

29436 writes at least the following in the POSIX locale:

29437 USER PID MOM COMMAND

29438 helene 34 12 ps -o uid,pid,ppid=MOM -o args

ps Utilities

The contents of the **COMMAND** field need not be the same in all implementations, due to possible truncation.

29441 RATIONALE

There is very little commonality between BSD and System V implementations of *ps.* Many options conflict or have subtly different usages. The standard developers attempted to select a set of options for the base standard that were useful on a wide range of systems and selected options that either can be implemented on both BSD and System V-based systems without breaking the current implementations or where the options are sufficiently similar that any changes would not be unduly problematic for users or implementors.

It is recognized that on some implementations, especially multi-level secure systems, *ps* may be nearly useless. The default output has therefore been chosen such that it does not break historical implementations and also is likely to provide at least some useful information on most systems.

The major change is the addition of the format specification capability. The motivation for this invention is to provide a mechanism for users to access a wider range of system information, if the system permits it, in a portable manner. The fields chosen to appear in this volume of IEEE Std 1003.1-2001 were arrived at after considering what concepts were likely to be both reasonably useful to the "average" user and had a reasonable chance of being implemented on a wide range of systems. Again it is recognized that not all systems are able to provide all the information and, conversely, some may wish to provide more. It is hoped that the approach adopted will be sufficiently flexible and extensible to accommodate most systems. Implementations may be expected to introduce new format specifiers.

The default output should consist of a short listing containing the process ID, terminal name, cumulative execution time, and command name of each process.

The preference of the standard developers would have been to make the format specification an operand of the *ps* command. Unfortunately, BSD usage precluded this.

At one time a format was included to display the environment array of the process. This was deleted because there is no portable way to display it.

The $-\mathbf{A}$ option is equivalent to the BSD $-\mathbf{g}$ and the SVID $-\mathbf{e}$. Because the two systems differed, a mnemonic compromise was selected.

The **–a** option is described with some optional behavior because the SVID omits session leaders, but BSD does not.

In an early proposal, format specifiers appeared for priority and start time. The former was not defined adequately in this volume of IEEE Std 1003.1-2001 and was removed in deference to the defined nice value; the latter because elapsed time was considered to be more useful.

In a new BSD version of ps, a $-\mathbf{O}$ option can be used to write all of the default information, followed by additional format specifiers. This was not adopted because the default output is implementation-defined. Nevertheless, this is a useful option that should be reserved for that purpose. In the $-\mathbf{o}$ option for the POSIX Shell and Utilities ps, the format is the concatenation of each $-\mathbf{o}$. Therefore, the user can have an alias or function that defines the beginning of their desired format and add more fields to the end of the output in certain cases where that would be useful.

The format of the terminal name is unspecified, but the descriptions of *ps, talk, who,* and *write* require that they all use the same format.

The **pcpu** field indicates that the CPU time available is determined in an unspecified manner. This is because it is difficult to express an algorithm that is useful across all possible machine

Utilities ps

29485 architectures. Historical counterparts to this value have attempted to show percentage of use in 29486 the recent past, such as the preceding minute. Frequently, these values for all processes did not add up to 100%. Implementations are encouraged to provide data in this field to users that will 29487 help them identify processes currently affecting the performance of the system. 29488 29489 FUTURE DIRECTIONS None. 29490 29491 **SEE ALSO** 29492 kill, nice, renice 29493 CHANGE HISTORY First released in Issue 2. 29494 29495 Issue 6 This utility is marked as part of the User Portability Utilities option. 29496 The normative text is reworded to avoid use of the term "must" for application requirements. 29497

The TZ entry is added to the ENVIRONMENT VARIABLES section.

29498

pwd **Utilities**

29499 NAME 29500	pwd — retur	n working directory name		
29501 SYNOP 29502	SIS pwd [-L	-P]		
29503 DESCR 29504 29505	The pwd util	lity shall write to standard output an absolute pathname of the current working hich does not contain the filenames dot or dot-dot.		
29506 OPTIO I 29507 29508	The <i>pwd</i> util	lity shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section Syntax Guidelines.		
29509	The followin	g options shall be supported by the implementation:		
29510 29511 29512 29513	-L	If the PWD environment variable contains an absolute pathname of the current directory that does not contain the filenames dot or dot-dot, pwd shall write this pathname to standard output. Otherwise, the $-\mathbf{L}$ option shall behave as the $-\mathbf{P}$ option.		
29514 29515	- P	The absolute pathname written shall not contain filenames that, in the context of the pathname, refer to files of type symbolic link.		
29516 29517	If both $-L$ and $-P$ are specified, the last one shall apply. If neither $-L$ nor $-P$ is specified, the pwd utility shall behave as if $-L$ had been specified.			
29518 OPERA 29519	NDS None.			
29520 STDIN 29521	Not used.			
29522 INPUT 29523	FILES None.			
29524 ENVIR 0 29525		ARIABLES ag environment variables shall affect the execution of <i>pwd</i> :		
29526 29527 29528 29529	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)		
29530 29531	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.		
29532 29533 29534	LC_MESSAC	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.		
29535 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .		
29536 29537 29538 29539 29540	PWD	If the –P option is in effect, this variable shall be set to an absolute pathname of the current working directory that does not contain any components that specify symbolic links, does not contain any components that are dot, and does not contain any components that are dot-dot. If an application sets or unsets the value of <i>PWD</i> , the behavior of <i>pwd</i> is unspecified.		

Utilities pwd

29541 ASYNCHRONOUS EVENTS

29542 Default.

29543 **STDOUT**

The *pwd* utility output is an absolute pathname of the current working directory: 29544

29545 "%s\n", <directory pathname>

29546 STDERR

29547 The standard error shall be used only for diagnostic messages.

29548 OUTPUT FILES

29549 None.

29550 EXTENDED DESCRIPTION

29551 None.

29552 EXIT STATUS

29553 The following exit values shall be returned:

Successful completion. 29554

>0 An error occurred. 29555

29556 CONSEQUENCES OF ERRORS

If an error is detected, output shall not be written to standard output, a diagnostic message shall 29557 29558 be written to standard error, and the exit status is not zero.

29559 APPLICATION USAGE

None. 29560

29561 EXAMPLES

None. 29562

29563 RATIONALE

Some implementations have historically provided *pwd* as a shell special built-in command. 29564

In most utilities, if an error occurs, partial output may be written to standard output. This does 29565 not happen in historical implementations of pwd. Because pwd is frequently used in historical 29566 29567 shell scripts without checking the exit status, it is important that the historical behavior is required here; therefore, the CONSEQUENCES OF ERRORS section specifically disallows any 29568

partial output being written to standard output. 29569

29570 FUTURE DIRECTIONS

None. 29571

29572 **SEE ALSO**

29573 cd, the System Interfaces volume of IEEE Std 1003.1-2001, getcwd()

29574 CHANGE HISTORY

First released in Issue 2. 29575

29576 Issue 6

29577 The -P and -L options are added to describe actions relating to symbolic links as specified in the

IEEE P1003.2b draft standard. 29578

qalter Utilities

```
      29579 NAME

      29580 qalter — alter batch job

      29581 SYNOPSIS

      29582 BE qalter [-a date_time] [-A account_string] [-c interval] [-e path_name]

      29583 [-h hold_list] [-j join_list] [-k keep_list] [-l resource_list]

      29584 [-m mail_options] [-M mail_list] [-N name] [-o path_name]

      29585 [-p priority] [-r y|n] [-S path_name_list] [-u user_list]

      29586 job_identifier ...
```

29588 **DESCRIPTION**

29589

29590

29591

29602

29603

29604

29605 29606

29607

29609

29610 29611

2961229613

29614 29615

29616 29617

29618

29619 29620

29621

29622

The attributes of a batch job are altered by a request to the batch server that manages the batch job. The *qalter* utility is a user-accessible batch client that requests the alteration of the attributes of one or more batch jobs.

The *qalter* utility shall alter the attributes of those batch jobs, and only those batch jobs, for which a batch *job_identifier* is presented to the utility.

The *qalter* utility shall alter the attributes of batch jobs in the order in which the batch *job_identifiers* are presented to the utility.

29596 If the *qalter* utility fails to process a batch *job_identifier* successfully, the utility shall proceed to process the remaining batch *job_identifiers*, if any.

For each batch *job_identifier* for which the *qalter* utility succeeds, each attribute of the identified batch job shall be altered as indicated by all the options presented to the utility.

For each identified batch job for which the *qalter* utility fails, the utility shall not alter any attribute of the batch job.

For each batch job that the *qalter* utility processes, the utility shall not modify any attribute other than those required by the options and option-arguments presented to the utility.

The *qalter* utility shall alter batch jobs by sending a *Modify Job Request* to the batch server that manages each batch job. At the time the *qalter* utility exits, it shall have modified the batch job corresponding to each successfully processed batch *job_identifier*. An attempt to alter the attributes of a batch job in the RUNNING state is implementation-defined.

29608 OPTIONS

The *qalter* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported by the implementation:

-a date_time Redefine the time at which the batch job becomes eligible for execution.

The *date_time* argument shall be in the same form and represent the same time as for the *touch* utility. The time so represented shall be set into the *Execution_Time* attribute of the batch job. If the time specified is earlier than the current time, the —a option shall have no effect.

-A account_string

Redefine the account to which the resource consumption of the batch job should be charged.

The syntax of the *account_string* option-argument is unspecified.

The *qalter* utility shall set the *Account_Name* attribute of the batch job to the value of the *account_string* option-argument.

Utilities qalter

29623	-c interval	Redefine who	ether the batch job should be checkpointed, and if so, how often.
29624 29625		The <i>qalter</i> uti	lity shall accept a value for the interval option-argument that is one of g:
29626 29627		n	No checkpointing is to be performed on the batch job (NO_CHECKPOINT).
29628 29629		S	Checkpointing is to be performed only when the batch server is shut down (CHECKPOINT_AT_SHUTDOWN).
29630 29631 29632		С	Automatic periodic checkpointing is to be performed at the <i>Minimum_Cpu_Interval</i> attribute of the batch queue, in units of CPU minutes (CHECKPOINT_AT_MIN_CPU_INTERVAL).
29633 29634 29635 29636		c=minutes	Automatic periodic checkpointing is to be performed every <i>minutes</i> of CPU time, or every <i>Minimum_Cpu_Interval</i> minutes, whichever is greater. The <i>minutes</i> argument shall conform to the syntax for unsigned integers and shall be greater than zero.
29637 29638 29639 29640		document for	entation may define other checkpoint intervals. The conformance or an implementation shall describe any alternative checkpoint w they are specified, their internal behavior, and how they affect the he utility.
29641 29642		The <i>qalter</i> uti <i>interval</i> option	lity shall set the <i>Checkpoint</i> attribute of the batch job to the value of the on-argument.
29643 29644	−e path_name		path to be used for the standard error stream of the batch job.
29645 29646 29647 29648		syntax of th	ility shall accept a <i>path_name</i> option-argument that conforms to the ne <i>path_name</i> element defined in the System Interfaces volume of 3.1-2001, which can be preceded by a host name element of the form
29649 29650 29651		utility shall	name option-argument constitutes an absolute pathname, the <i>qalter</i> set the <i>Error_Path</i> attribute of the batch job to the value of the otion-argument, including the host name element, if present.
29652 29653 29654 29655 29656		name element batch job to	name option-argument constitutes a relative pathname and no host it is specified, the <i>qalter</i> utility shall set the <i>Error_Path</i> attribute of the other value of the absolute pathname derived by expanding the ption-argument relative to the current directory of the process that <i>qalter</i> utility.
29657 29658 29659		•	ame option-argument constitutes a relative pathname and a host name
		-	pecified, the <i>qalter</i> utility shall set the <i>Error_Path</i> attribute of the batch use of the option-argument without expansion.
29660 29661 29662		job to the val If the <i>path_na</i> utility shall j	- · · · · · · · · · · · · · · · · · · ·
29661	− h hold_list	job to the val If the path_na utility shall p hostname is th Redefine the accept a value	ue of the option-argument without expansion. ame option-argument does not include a host name element, the qalter prefix the pathname in the Error_Path attribute with hostname:, where

qalter Utilities

29668 29669 29670		'n'. For each unique character in the <i>hold_list</i> option-argument, the <i>qalter</i> utility shall add a value to the <i>Hold_Types</i> attribute of the batch job as follows, each representing a different hold type:
29671		u USER
29672		s SYSTEM
29673		o OPERATOR
29674 29675 29676		If any of these characters are duplicated in the <i>hold_list</i> option-argument, the duplicates shall be ignored. An existing <i>Hold_Types</i> attribute can be cleared by the hold type:
29677		n NO_HOLD
29678 29679 29680 29681 29682 29683 29684 29685		The <i>qalter</i> utility shall consider it an error if any hold type other than 'n' is combined with hold type 'n'. Strictly conforming applications shall not repeat any of the characters 'u', 's', 'o', or 'n' within the <i>hold_list</i> option-argument. The <i>qalter</i> utility shall permit the repetition of characters, but shall not assign additional meaning to the repeated characters. An implementation may define other hold types. The conformance document for an implementation shall describe any additional hold types, how they are specified, their internal behavior, and how they affect the behavior of the utility.
29686 29687 29688	– j join_list	Redefine which streams of the batch job are to be merged. The <i>qalter</i> – j option shall accept a value for the <i>join_list</i> option-argument that is a string of alphanumeric characters in the portable character set.
29689 29690		The <i>qalter</i> utility shall accept a <i>join_list</i> option-argument that consists of one or more of the characters $'e'$ and $'o'$, or the single character $'n'$.
29691 29692		All of the other batch job output streams specified shall be merged into the output stream represented by the character listed first in the <i>join_list</i> option-argument.
29693 29694 29695		For each unique character in the <i>join_list</i> option-argument, the <i>qalter</i> utility shall add a value to the <i>Join_Path</i> attribute of the batch job as follows, each representing a different batch job stream to join:
29696		e The standard error of the batch job (JOIN_STD_ERROR).
29697		 The standard output of the batch job (JOIN_STD_OUTPUT).
29698		An existing Join_Path attribute can be cleared by the join type:
29699		n NO_JOIN
29700 29701		If 'n' is specified, then no files are joined. The <i>qalter</i> utility shall consider it an error if any join type other than 'n' is combined with join type 'n'.
29702 29703 29704 29705		Strictly conforming applications shall not repeat any of the characters 'e', 'o', or 'n' within the <i>join_list</i> option-argument. The <i>qalter</i> utility shall permit the repetition of characters, but shall not assign additional meaning to the repeated characters.
29706 29707 29708		An implementation may define other join types. The conformance document for an implementation shall describe any additional batch job streams, how they are specified, their internal behavior, and how they affect the behavior of the utility.
29709	− k keep_list	Redefine which output of the batch job to retain on the execution host.

Utilities qalter

29710 The qalter -k option shall accept a value for the keep_list option-argument that is a 29711 string of alphanumeric characters in the portable character set. 29712 The *qalter* utility shall accept a *keep_list* option-argument that consists of one or more of the characters 'e' and 'o', or the single character 'n'. 29713 29714 For each unique character in the *keep_list* option-argument, the *qalter* utility shall add a value to the *Keep_Files* attribute of the batch job as follows, each representing 29715 a different batch job stream to keep: 29716 The standard error of the batch job (KEEP_STD_ERROR). 29717 е 29718 The standard output of the batch job (KEEP_STD_OUTPUT). 0 If both 'e' and 'o' are specified, then both files are retained. An existing 29719 29720 *Keep_Files* attribute can be cleared by the keep type: NO_KEEP 29721 If 'n' is specified, then no files are retained. The *qalter* utility shall consider it an 29722 error if any keep type other than 'n' is combined with keep type 'n'. 29723 Strictly conforming applications shall not repeat any of the characters 'e', 'o', or 29724 'n' within the keep list option-argument. The qalter utility shall permit the 29725 repetition of characters, but shall not assign additional meaning to the repeated 29726 29727 characters. An implementation may define other keep types. The conformance document for an implementation shall describe any additional keep types, how 29728 they are specified, their internal behavior, and how they affect the behavior of the 29729 29730 utility. –l resource_list 29731 29732 Redefine the resources that are allowed or required by the batch job. The qalter utility shall accept a resource_list option-argument that conforms to the 29733 29734 following syntax: resource=value[,,resource=value,,...] 29735 The qalter utility shall set one entry in the value of the Resource_List attribute of the 29736 batch job for each resource listed in the resource_list option-argument. 29737 Because the list of supported resource names might vary by batch server, the *qalter* 29738 utility shall rely on the batch server to validate the resource names and associated 29739 29740 values. See Section 3.3.3 (on page 123) for a means of removing keyword=value (and 29741 *value@keyword*) pairs and other general rules for list-oriented batch job attributes. -m mail_options 29742 29743 Redefine the points in the execution of the batch job at which the batch server is to send mail about a change in the state of the batch job. 29744 The *qalter* –**m** option shall accept a value for the *mail options* option-argument that 29745 is a string of alphanumeric characters in the portable character set. 29746 29747 The qalter utility shall accept a value for the mail_options option-argument that is a 29748 string of one or more of the characters 'e', 'b', and 'a', or the single character 'n'. For each unique character in the mail_options option-argument, the qalter 29749 utility shall add a value to the *Mail Users* attribute of the batch job as follows, each 29750 representing a different time during the life of a batch job at which to send mail: 29751 29752 MAIL_AT_EXIT

qalter Utilities

29753		b MAIL_AT_BEGINNING
29754		a MAIL_AT_ABORT
29755		If any of these characters are duplicated in the <i>mail_options</i> option-argument, the
29756		duplicates shall be ignored.
29757		An existing Mail_Points attribute can be cleared by the mail type:
29758		n NO_MAIL
29759 29760 29761 29762 29763 29764		If 'n' is specified, then mail is not sent. The <i>qalter</i> utility shall consider it an error if any mail type other than 'n' is combined with mail type 'n'. Strictly conforming applications shall not repeat any of the characters 'e', 'b', 'a', or 'n' within the <i>mail_options</i> option-argument. The <i>qalter</i> utility shall permit the repetition of characters but shall not assign additional meaning to the repeated characters.
29765 29766 29767		An implementation may define other mail types. The conformance document for an implementation shall describe any additional mail types, how they are specified, their internal behavior, and how they affect the behavior of the utility.
29768 29769	− M mail_list	Redefine the list of users to which the batch server that executes the batch job is to send mail, if the batch server sends mail about the batch job.
29770 29771 29772		The syntax of the <i>mail_list</i> option-argument is unspecified. If the implementation of the <i>qalter</i> utility uses a name service to locate users, the utility shall accept the syntax used by the name service.
29773 29774		If the implementation of the <i>qalter</i> utility does not use a name service to locate users, the implementation shall accept the following syntax for user names:
29775		<pre>mail_address[,,mail_address,,]</pre>
29776		The interpretation of <i>mail_address</i> is implementation-defined.
29777 29778		The <i>qalter</i> utility shall set the <i>Mail_Users</i> attribute of the batch job to the value of the <i>mail_list</i> option-argument.
29779	-N name	Redefine the name of the batch job.
29780 29781 29782		The $qalter$ $-N$ option shall accept a value for the $name$ option-argument that is a string of up to 15 alphanumeric characters in the portable character set where the first character is alphabetic.
29783		The syntax of the <i>name</i> option-argument is unspecified.
29784 29785		The <i>qalter</i> utility shall set the <i>Job_Name</i> attribute of the batch job to the value of the <i>name</i> option-argument.
29786 29787	-o path_name	Redefine the path for the standard output of the batch job.
29788 29789 29790 29791		The <i>qalter</i> utility shall accept a <i>path_name</i> option-argument that conforms to the syntax of the <i>path_name</i> element defined in the System Interfaces volume of IEEE Std 1003.1-2001, which can be preceded by a host name element of the form <i>hostname</i> :.
29792 29793 29794		If the <code>path_name</code> option-argument constitutes an absolute pathname, the <code>qalter</code> utility shall set the <code>Output_Path</code> attribute of the batch job to the value of the <code>path_name</code> option-argument.

Utilities qalter

29795 29796 29797 29798 29799		If the <i>path_name</i> option-argument constitutes a relative pathname and no host name element is specified, the <i>qalter</i> utility shall set the <i>Output_Path</i> attribute of the batch job to the absolute pathname derived by expanding the <i>path_name</i> option-argument relative to the current directory of the process that executes the <i>qalter</i> utility.
29800 29801 29802 29803		If the <i>path_name</i> option-argument constitutes a relative pathname and a host name element is specified, the <i>qalter</i> utility shall set the <i>Output_Path</i> attribute of the batch job to the value of the <i>path_name</i> option-argument without any expansion of the pathname.
29804 29805 29806		If the <code>path_name</code> option-argument does not include a host name element, the <code>qalter</code> utility shall prefix the pathname in the <code>Output_Path</code> attribute with <code>hostname</code> :, where <code>hostname</code> is the name of the host upon which the <code>qalter</code> utility is being executed.
29807	- p priority	Redefine the priority of the batch job.
29808 29809 29810		The <i>qalter</i> utility shall accept a value for the priority option-argument that conforms to the syntax for signed decimal integers, and which is not less than -1024 and not greater than 1023 .
29811 29812		The <i>qalter</i> utility shall set the <i>Priority</i> attribute of the batch job to the value of the <i>priority</i> option-argument.
29813	-r y n	Redefine whether the batch job is rerunnable.
29814 29815		If the value of the option-argument is $'y'$, the <i>qalter</i> utility shall set the <i>Rerunable</i> attribute of the batch job to TRUE.
29816 29817		If the value of the option-argument is 'n', the <i>qalter</i> utility shall set the <i>Rerunable</i> attribute of the batch job to FALSE.
29818 29819		The <i>qalter</i> utility shall consider it an error if any character other than $'y'$ or $'n'$ is specified in the option-argument.
29820 29821	-S path_nam	ne_list Redefine the shell that interprets the script at the destination system.
29822 29823		The <i>qalter</i> utility shall accept a <i>path_name_list</i> option-argument that conforms to the following syntax:
29824		<pre>pathname[@host][,pathname[@host],]</pre>
29825 29826		The <i>qalter</i> utility shall accept only one pathname that is missing a corresponding host name. The <i>qalter</i> utility shall allow only one pathname per named host.
29827 29828 29829 29830		The <i>qalter</i> utility shall add a value to the <i>Shell_Path_List</i> attribute of the batch job for each entry in the <i>path_name_list</i> option-argument. See Section 3.3.3 (on page 123) for a means of removing <i>keyword=value</i> (and <i>value@keyword</i>) pairs and other general rules for list-oriented batch job attributes.
29831 29832	-u user_list	Redefine the user name under which the batch job is to run at the destination system.
29833 29834		The <i>qalter</i> utility shall accept a <i>user_list</i> option-argument that conforms to the following syntax:
29835		username[@host][,,username[@host],,]
29836 29837		The <i>qalter</i> utility shall accept only one user name that is missing a corresponding host name. The <i>qalter</i> utility shall accept only one user name per named host.

qalter Utilities

29838 The *qalter* utility shall add a value to the *User_List* attribute of the batch job for each 29839 entry in the user_list option-argument. See Section 3.3.3 (on page 123) for a means of removing keyword=value (and value@keyword) pairs and other general rules for 29840 list-oriented batch job attributes. 29841 29842 OPERANDS The qalter utility shall accept one or more operands that conform to the syntax for a batch 29843 29844 *job_identifier* (see Section 3.3.1 (on page 122)). 29845 **STDIN** Not used. 29846 29847 INPUT FILES None. 29849 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *qalter*: 29850 LANG Provide a default value for the internationalization variables that are unset or null. 29851 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 29852 29853 Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 29854 LC ALL If set to a non-empty string value, override the values of all the other 29855 internationalization variables. 29856 29857 LC_CTYPE 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 29858 arguments). 29859 LC_MESSAGES 29860 Determine the locale that should be used to affect the format and contents of 29861 diagnostic messages written to standard error. 29862 *LOGNAME* Determine the login name of the user. 29863 TZDetermine the timezone used to interpret the *date-time* option-argument. If TZ is 29864 29865 unset or null, an unspecified default timezone shall be used. 29866 ASYNCHRONOUS EVENTS Default. 29867 29868 STDOUT None. 29869 **29870 STDERR** 29871 The standard error shall be used only for diagnostic messages. 29872 OUTPUT FILES None. 29873 29874 EXTENDED DESCRIPTION

29875 None.

29876 EXIT STATUS

The following exit values shall be returned:

29878 0 Successful completion.

29879 >0 An error occurred.

Utilities qalter

29880 CONSEQUENCES OF ERRORS

In addition to the default behavior, the *qalter* utility shall not be required to write a diagnostic message to standard error when the error reply received from a batch server indicates that the batch *job_identifier* does not exist on the server. Whether or not the *qalter* utility attempts to locate the batch job on other batch servers is implementation-defined.

29885 APPLICATION USAGE

29886 None.29887 **EXAMPLES**

29888 None.29889 **RATIONALE**

The *qalter* utility allows users to change the attributes of a batch job.

As a means of altering a queued job, the *qalter* utility is superior to deleting and requeuing the batch job insofar as an altered job retains its place in the queue with some traditional selection algorithms. In addition, the *qalter* utility is both shorter and simpler than a sequence of *qdel* and *qsub* utilities.

The result of an attempt on the part of a user to alter a batch job in a RUNNING state is implementation-defined because a batch job in the RUNNING state will already have opened its output files and otherwise performed any actions indicated by the options in effect at the time the batch job began execution.

The options processed by the *qalter* utility are identical to those of the *qsub* utility, with a few exceptions: $-\mathbf{V}$, $-\mathbf{v}$, and $-\mathbf{q}$. The $-\mathbf{V}$ and $-\mathbf{v}$ are inappropriate for the *qalter* utility, since they capture potentially transient environment information from the submitting process. The $-\mathbf{q}$ option would specify a new queue, which would largely negate the previously stated advantage of using *qalter*; furthermore, the *qmove* utility provides a superior means of moving jobs.

Each of the following paragraphs provides the rationale for a *qalter* option.

Additional rationale concerning these options can be found in the rationale for the *qsub* utility.

The -a option allows users to alter the date and time at which a batch job becomes eligible to run.

The **–A** option allows users to change the account that will be charged for the resources consumed by the batch job. Support for the **–A** option is mandatory for conforming implementations of *qalter*, even though support of accounting is optional for servers. Whether or not to support accounting is left to the implementor of the server, but mandatory support of the **–A** option assures users of a consistent interface and allows them to control accounting on servers that support accounting.

The –c option allows users to alter the checkpointing interval of a batch job. A checkpointing system, which is not defined by IEEE Std 1003.1-2001, allows recovery of a batch job at the most recent checkpoint in the event of a crash. Checkpointing is typically used for jobs that consume expensive computing time or must meet a critical schedule. Users should be allowed to make the tradeoff between the overhead of checkpointing and the risk to the timely completion of the batch job; therefore, this volume of IEEE Std 1003.1-2001 provides the checkpointing interval option. Support for checkpointing is optional for servers.

The **–e** option allows users to alter the name and location of the standard error stream written by a batch job. However, the path of the standard error stream is meaningless if the value of the *Join_Path* attribute of the batch job is TRUE.

The $-\mathbf{h}$ option allows users to set the hold type in the $Hold_Types$ attribute of a batch job. The qhold and qrls utilities add or remove hold types to the $Hold_Types$ attribute, respectively. The $-\mathbf{h}$

qalter **Utilities**

29926	option has been modified to allow for implementation-defined hold types.
29927 29928	The $-\mathbf{j}$ option allows users to alter the decision to join (merge) the standard error stream of the batch job with the standard output stream of the batch job.
29929	The $-\mathbf{l}$ option allows users to change the resource limits imposed on a batch job.
29930 29931	The $-m$ option allows users to modify the list of points in the life of a batch job at which the designated users will receive mail notification.
29932 29933	The $-M$ option allows users to alter the list of users who will receive notification about events in the life of a batch job.
29934	The $-N$ option allows users to change the name of a batch job.
29935 29936	The $-\mathbf{o}$ option allows users to alter the name and path to which the standard output stream of the batch job will be written.
29937 29938	The $-\mathbf{P}$ option allows users to modify the priority of a batch job. Support for priority is optional for batch servers.
29939	The $-\mathbf{r}$ option allows users to alter the rerunability status of a batch job.
29940 29941 29942	The $-S$ option allows users to change the name and location of the shell image that will be invoked to interpret the script of the batch job. This option has been modified to allow a list of shell name and locations associated with different hosts.
29943	The $-\mathbf{u}$ option allows users to change the user identifier under which the batch job will execute.
29944 29945 29946 29947	The <i>job_identifier</i> operand syntax is provided so that the user can differentiate between the originating and destination (or executing) batch server. These may or may not be the same. The <i>.server_name</i> portion identifies the originating batch server, while the <i>@server</i> portion identifies the destination batch server.
29948 29949	Historically, the <i>qalter</i> utility has been a component of the Network Queuing System (NQS), the existing practice from which this utility has been derived.
29950 FUTUR 29951	None.
29952 SEE AL	
29953	Chapter 3 (on page 101), qdel, qhold, qmove, qrls, qsub, touch
29954 CHAN 0 29955	GE HISTORY Derived from IEEE Std 1003.2d-1994.
29956 Issue 6 29957	The TZ entry is added to the ENVIRONMENT VARIABLES section.
29958	IEEE PASC Interpretation 1003.2 $\#182$ is applied, clarifying the description of the $-\mathbf{a}$ option.

Utilities qdel

29959 **NAME**

29960 qdel — delete batch jobs

29961 SYNOPSIS

29962 BE qdel job_identifier ...

29963

29964 **DESCRIPTION**

A batch job is deleted by sending a request to the batch server that manages the batch job. A batch job that has been deleted is no longer subject to management by batch services.

The *qdel* utility is a user-accessible client of batch services that requests the deletion of one or more batch jobs.

The *qdel* utility shall request a batch server to delete those batch jobs for which a batch *job_identifier* is presented to the utility.

The *qdel* utility shall delete batch jobs in the order in which their batch *job_identifiers* are presented to the utility.

If the *qdel* utility fails to process any batch *job_identifier* successfully, the utility shall proceed to process the remaining batch *job_identifiers*, if any.

The *qdel* utility shall delete each batch job by sending a *Delete Job Request* to the batch server that manages the batch job.

The *qdel* utility shall not exit until the batch job corresponding to each successfully processed batch *job_identifier* has been deleted.

29979 OPTIONS

29980 None.

29981 OPERANDS

The *qdel* utility shall accept one or more operands that conform to the syntax for a batch *job_identifier* (see Section 3.3.1 (on page 122)).

29984 STDIN

29989

29985 Not used.

29986 INPUT FILES

29987 None.

29988 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *qdel*:

29990 LANG Provide a default value for the internationalization variables that are unset or null.
29991 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
29992 Internationalization Variables for the precedence of internationalization variables
29993 used to determine the values of locale categories.)

29994 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

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).

29999 LC_MESSAGES

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

qdel Utilities

30002 *LOGNAME* Determine the login name of the user.

30003 ASYNCHRONOUS EVENTS

30004 Default.

30005 STDOUT

30006 An implementation of the *qdel* utility may write informative messages to standard output.

30007 STDERR

30008 The standard error shall be used only for diagnostic messages.

30009 OUTPUT FILES

30010 None.

30011 EXTENDED DESCRIPTION

30012 None.

30013 EXIT STATUS

The following exit values shall be returned:

30015 0 Successful completion.

30016 >0 An error occurred.

30017 CONSEQUENCES OF ERRORS

In addition to the default behavior, the *qdel* utility shall not be required to write a diagnostic message to standard error when the error reply received from a batch server indicates that the batch *job_identifier* does not exist on the server. Whether or not the *qdel* utility waits to output the diagnostic message while attempting to locate the job on other servers is implementation-defined.

30023 APPLICATION USAGE

30024 None.

30025 EXAMPLES

30026 None.

30027 RATIONALE

30028

30029

30030

30031

30032 30033

30034

30035

30036

30037

30038

30039

30040

30041

30042

30044

30045

The *qdel* utility allows users and administrators to delete jobs.

The *qdel* utility provides functionality that is not otherwise available. For example, the *kill* utility of the operating system does not suffice. First, to use the *kill* utility, the user might have to log in on a remote node, because the *kill* utility does not operate across the network. Second, unlike *qdel*, *kill* cannot remove jobs from queues. Lastly, the arguments of the *qdel* utility are job identifiers rather than process identifiers, and so this utility can be passed the output of the *qselect* utility, thus providing users with a means of deleting a list of jobs.

Because a set of jobs can be selected using the *qselect* utility, the *qdel* utility has not been complicated with options that provide for selection of jobs. Instead, the batch jobs to be deleted are identified individually by their job identifiers.

Historically, the *qdel* utility has been a component of NQS, the existing practice on which it is based. However, the *qdel* utility defined in this volume of IEEE Std 1003.1-2001 does not provide an option for specifying a signal number to send to the batch job prior to the killing of the process; that capability has been subsumed by the *qsig* utility.

A discussion was held about the delays of networking and the possibility that the batch server may never respond, due to a down router, down batch server, or other network mishap. The DESCRIPTION records this under the words "fails to process any job identifier". In the broad sense, the network problem is also an error, which causes the failure to process the batch job

Utilities qdel

30046 identifier.

30047 FUTURE DIRECTIONS

30048 None.

30049 SEE ALSO

30050 Chapter 3 (on page 101), kill, qselect, qsig

30051 CHANGE HISTORY

30052 Derived from IEEE Std 1003.2d-1994.

30053 **Issue 6**

30054 The LC_TIME and TZ entries are removed from the ENVIRONMENT VARIABLES section.

qhold **Utilities**

30055 **NAME**

30056 qhold — hold batch jobs

30057 SYNOPSIS

qhold [-h hold list] job identifier ... 30058 BE

30059

30060 DESCRIPTION

A hold is placed on a batch job by a request to the batch server that manages the batch job. A 30061 batch job that has one or more holds is not eligible for execution. The qhold utility is a user-30062 30063 accessible client of batch services that requests one or more types of hold to be placed on one or 30064 more batch jobs.

The *qhold* utility shall place holds on those batch jobs for which a batch *job_identifier* is presented 30065 to the utility. 30066

The *qhold* utility shall place holds on batch jobs in the order in which their batch *job_identifiers* 30067 are presented to the utility. If the *qhold* utility fails to process any batch *job identifier* successfully, 30068 the utility shall proceed to process the remaining batch *job_identifiers*, if any. 30069

30070 The qhold utility shall place holds on each batch job by sending a Hold Job Request to the batch 30071 server that manages the batch job.

The *qhold* utility shall not exit until holds have been placed on the batch job corresponding to 30072 30073 each successfully processed batch *job_identifier*.

30074 OPTIONS

The *qhold* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section

The following option shall be supported by the implementation:

The *qhold* -**h** option shall accept a value for the *hold_list* option-argument that is a string of alphanumeric characters in the portable character set (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.1, Portable Character Set).

The *qhold* utility shall accept a value for the *hold_list* option-argument that is a string of one or more of the characters 'u', 's', or 'o', or the single character 'n'.

For each unique character in the *hold_list* option-argument, the *qhold* utility shall add a value to the *Hold Types* attribute of the batch job as follows, each representing a different hold type:

USER 11

SYSTEM S

OPERATOR

If any of these characters are duplicated in the *hold_list* option-argument, the duplicates shall be ignored.

An existing *Hold_Types* attribute can be cleared by the following hold type:

NO HOLD

The *qhold* utility shall consider it an error if any hold type other than 'n' is combined with hold type 'n'.

30075 12.2, Utility Syntax Guidelines. 30076

30077 30078

30079 30080

30081 30082

30083

30084

30085

30086

30087

30088

30089

30090

30091 30092

30093 30094

30095 30096

-h *hold_list* Define the types of holds to be placed on the batch job.

772

Utilities **qhold**

An implementation may define other hold types. The conformance document for an implementation shall describe any additional hold types, how they are specified, their internal behavior, and how they affect the behavior of the utility. 30104 If the –h option is not presented to the qhold utility, the implementation shall set the Hold_Types attribute to USER. 30106 OPERANDS 30107 The qhold utility shall accept one or more operands that conform to the syntax for a batch job_identifier (see Section 3.3.1 (on page 122)). 30108 STDIN 30110 Not used. 30111 INPUT FILES 30112 None. 30113 ENVIRONMENT VARIABLES 30114 The following environment variables shall affect the execution of qhold: 30115 LANG Provide a default value for the internationalization variables that are unset or null. See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 30119 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 3012 LC_CTYPE 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). 30124 LC_MESSAGES 30125 Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. 30126 ASYNCHRONOUS EVENTS 30127 LOGNAME Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. 30128 STDERR 30130 The standard error shall be used only for diagnostic messages. 3014 OUTPUT FILES 3015 None. 3015 None. 3016 STILES None. 3017 None STDOUT	30097 30098 30099 30100		Strictly conforming applications shall not repeat any of the characters 'u', 's', 'o', or 'n' within the <i>hold_list</i> option-argument. The <i>qhold</i> utility shall permit the repetition of characters, but shall not assign additional meaning to the repeated characters.
the Hold_Types attribute to USER. OPERANDS The qhold utility shall accept one or more operands that conform to the syntax for a batch job_identifier (see Section 3.3.1 (on page 122)). OPERANDS The qhold utility shall accept one or more operands that conform to the syntax for a batch job_identifier (see Section 3.3.1 (on page 122)). OPERANDS The qhold utility shall accept one or more operands that conform to the syntax for a batch job_identifier (see Section 3.3.1 (on page 122)). OPERANDS STDIN	30102		an implementation shall describe any additional hold types, how they are
The qhold utility shall accept one or more operands that conform to the syntax for a batch job_identifier (see Section 3.3.1 (on page 122)). The following environment variables shall affect the execution of qhold: The following environment variables shall affect the execution of qhold: LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. LC_CTYPE 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). LC_MESSAGES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. LOGNAME Determine the login name of the user. LOGNAME Determine the login name of the user. LOGNAME STDEUR None. The standard error shall be used only for diagnostic messages. OUTPUT FILES None.			
The qhold utility shall accept one or more operands that conform to the syntax for a batch job_identifier (see Section 3.3.1 (on page 122)). The following environment variables shall affect the execution of qhold: The following environment variables shall affect the execution of qhold: LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. LC_CTYPE 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). LC_MESSAGES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. LOGNAME Determine the login name of the user. LOGNAME Determine the login name of the user. LOGNAME STDEUR None. The standard error shall be used only for diagnostic messages. OUTPUT FILES None.	30106 OPERA	NDS	
Not used. None.	30107	The <i>qhold</i> u	
None.	30109 STDIN		
Solita None.	30110	Not used.	
The following environment variables shall affect the execution of qhold: LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. LC_CTYPE 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). LC_MESSAGES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. LOGNAME Determine the login name of the user. COGNAME Determine the login name of the user. COGNAME STDOUT None. The standard error shall be used only for diagnostic messages. The standard error shall be used only for diagnostic messages.			
30115 LANG Provide a default value for the internationalization variables that are unset or null. 30116 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 30117 Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 30119 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 30121 LC_CTYPE 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). 30124 LC_MESSAGES 30125 Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. 30127 LOGNAME Determine the login name of the user. 30128 ASYNCHRONOUS EVENTS 30129 Default. 30130 STDOUT 30131 None. 30132 STDERR 30133 The standard error shall be used only for diagnostic messages. 30134 OUTPUT FILES 30135 None.	30113 ENVIR	ONMENT VA	ARIABLES
Solition Section Sec	30114	The followin	ng environment variables shall affect the execution of <i>qhold</i> :
Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. LC_CTYPE 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). LC_MESSAGES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. LOGNAME Determine the login name of the user. LOGNAME Default. None. The standard error shall be used only for diagnostic messages. None.	30115	LANG	Provide a default value for the internationalization variables that are unset or null.
used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. LC_CTYPE 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). LC_MESSAGES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. LOGNAME Determine the login name of the user. LOGNAME Determine the login name of the user. STDOUT None. The standard error shall be used only for diagnostic messages. None.	30116		
30119	30117		
internationalization variables. 30121	30118		used to determine the values of locale categories.)
characters (for example, single-byte as opposed to multi-byte characters in arguments). LC_MESSAGES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. LOGNAME Determine the login name of the user. LOGNAME Default. Default. Soli28 ASYNCHRONOUS EVENTS Default. None. STDOUT 30131 None. The standard error shall be used only for diagnostic messages. 30134 OUTPUT FILES 30135 None.		LC_ALL	
Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. **Today**	30122	LC_CTYPE	characters (for example, single-byte as opposed to multi-byte characters in
Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. **Today**	30124	LC_MESSA	GES
30127 LOGNAME Determine the login name of the user. 30128 ASYNCHRONOUS EVENTS 30129 Default. 30130 STDOUT 30131 None. 30132 STDERR 30133 The standard error shall be used only for diagnostic messages. 30134 OUTPUT FILES 30135 None.	30125		
30128 ASYNCHRONOUS EVENTS 30129 Default. 30130 STDOUT 30131 None. 30132 STDERR 30133 The standard error shall be used only for diagnostic messages. 30134 OUTPUT FILES 30135 None.	30126		diagnostic messages written to standard error.
30129 Default. 30130 STDOUT 30131 None. 30132 STDERR 30133 The standard error shall be used only for diagnostic messages. 30134 OUTPUT FILES 30135 None.	30127	LOGNAME	Determine the login name of the user.
30129 Default. 30130 STDOUT 30131 None. 30132 STDERR 30133 The standard error shall be used only for diagnostic messages. 30134 OUTPUT FILES 30135 None.	30128 ASYN (HRONOUS	EVENTS
None. 30132 STDERR 30133 The standard error shall be used only for diagnostic messages. 30134 OUTPUT FILES 30135 None.			
None. 30132 STDERR 30133 The standard error shall be used only for diagnostic messages. 30134 OUTPUT FILES 30135 None.	30130 STDOI	J T	
The standard error shall be used only for diagnostic messages. OUTPUT FILES None.			
The standard error shall be used only for diagnostic messages. OUTPUT FILES None.	20122 STDER	D	
30135 None.			d error shall be used only for diagnostic messages.
30135 None.	30134 OUTPI	T FILES	
20126 EVTENDED DESCRIPTION			
20130 EXTENDED DESCRIPTION	30136 EXTEN	DED DESCR	IPTION

None.

30137

qhold Utilities

30138 EXIT STATUS

The following exit values shall be returned:

30140 0 Successful completion.

30141 >0 An error occurred.

30142 CONSEQUENCES OF ERRORS

In addition to the default behavior, the *qhold* utility shall not be required to write a diagnostic message to standard error when the error reply received from a batch server indicates that the batch *job_identifier* does not exist on the server. Whether or not the *qhold* utility waits to output the diagnostic message while attempting to locate the job on other servers is implementation-

30147 defined.

30148 APPLICATION USAGE

30149 None.

30150 EXAMPLES

30151 None.

30152 RATIONALE

The *qhold* utility allows users to place a hold on one or more jobs. A hold makes a batch job ineligible for execution.

The *qhold* utility has options that allow the user to specify the type of hold. Should the user wish to place a hold on a set of jobs that meet a selection criteria, such a list of jobs can be acquired using the *qselect* utility.

The -h option allows the user to specify the type of hold that is to be placed on the job. This option allows for USER, SYSTEM, OPERATOR, and implementation-defined hold types. The USER and OPERATOR holds are distinct. The batch server that manages the batch job will verify that the user is authorized to set the specified hold for the batch job.

Mail is not required on hold because the administrator has the tools and libraries to build this option if he or she wishes.

Historically, the *qhold* utility has been a part of some existing batch systems, although it has not traditionally been a part of the NQS.

30166 FUTURE DIRECTIONS

30167 None.

30168 **SEE ALSO**

30169 Chapter 3 (on page 101), qselect

30170 CHANGE HISTORY

30171 Derived from IEEE Std 1003.2d-1994.

30172 **Issue 6**

30173 The *LC_TIME* and *TZ* entries are removed from the ENVIRONMENT VARIABLES section.

Utilities qmove

30174 **NAME**

30175 qmove — move batch jobs

30176 SYNOPSIS

30177 BE qmove destination job_identifier ..

30178

30179 **DESCRIPTION**

To move a batch job is to remove the batch job from the batch queue in which it resides and instantiate the batch job in another batch queue. A batch job is moved by a request to the batch server that manages the batch job. The *qmove* utility is a user-accessible batch client that requests the movement of one or more batch jobs.

The *qmove* utility shall move those batch jobs, and only those batch jobs, for which a batch *job_identifier* is presented to the utility.

The *qmove* utility shall move batch jobs in the order in which the corresponding batch job_identifiers are presented to the utility.

If the *qmove* utility fails to process a batch *job_identifier* successfully, the utility shall proceed to process the remaining batch *job_identifiers*, if any.

The *qmove* utility shall move batch jobs by sending a *Move Job Request* to the batch server that manages each batch job. The *qmove* utility shall not exit before the batch jobs corresponding to all successfully processed batch *job_identifiers* have been moved.

30193 OPTIONS

30194 None.

30195 OPERANDS

The *qmove* utility shall accept one operand that conforms to the syntax for a destination (see Section 3.3.2 (on page 123)).

The *qmove* utility shall accept one or more operands that conform to the syntax for a batch *job_identifier* (see Section 3.3.1 (on page 122)).

30200 STDIN

30201 Not used.

30202 INPUT FILES

30203 None.

30204 ENVIRONMENT VARIABLES

30205 The following environment variables shall affect the execution of *qmove*:

30206 LANG Provide a default value for the internationalization variables that are unset or null.
30207 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
30208 Internationalization Variables for the precedence of internationalization variables
30209 used to determine the values of locale categories.)

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

 LC_CTYPE 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).

30215 LC_MESSAGES

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

Utilities qmove

30218 *LOGNAME* Determine the login name of the user. 30219 ASYNCHRONOUS EVENTS 30220 Default. **30221 STDOUT** 30222 None. 30223 STDERR 30224 The standard error shall be used only for diagnostic messages. 30225 OUTPUT FILES 30226 None. 30227 EXTENDED DESCRIPTION 30228 None. 30229 EXIT STATUS 30230 The following exit values shall be returned: Successful completion. 30231 >0 An error occurred. 30232 30233 CONSEQUENCES OF ERRORS 30234 In addition to the default behavior, the *qmove* utility shall not be required to write a diagnostic 30235 message to standard error when the error reply received from a batch server indicates that the 30236 batch job_identifier does not exist on the server. Whether or not the qmove utility waits to output the diagnostic message while attempting to locate the job on other servers is implementation-30237 defined. 30238 30239 APPLICATION USAGE 30240 None. 30241 EXAMPLES 30242 None. 30243 RATIONALE 30244 The *qmove* utility allows users to move jobs between queues. 30245 The alternative to using the qmove utility—deleting the batch job and requeuing it—entails considerably more typing. 30246 Since the means of selecting jobs based on attributes has been encapsulated in the *qselect* utility, 30247 30248 the only option of the *qmove* utility concerns authorization. The $-\mathbf{u}$ option provides the user with the convenience of changing the user identifier under which the batch job will execute. 30249 Minimalism and consistency have taken precedence over convenience; the -u option has been 30250 deleted because the equivalent capability exists with the $-\mathbf{u}$ option of the *qalter* utility. 30251 30252 FUTURE DIRECTIONS 30253 None. 30254 SEE ALSO 30255 Chapter 3 (on page 101), qalter, qselect 30256 CHANGE HISTORY Derived from IEEE Std 1003.2d-1994.

30257

Utilities **qmove**

30258 **Issue 6**

30259

The *LC_TIME* and *TZ* entries are removed from the ENVIRONMENT VARIABLES section.

qmsg Utilities

```
30260 NAME
```

30261 qmsg — send message to batch jobs

30262 SYNOPSIS

30263 BE qmsg [-E][-O] message_string job_identifier ...
30264

30273

30274

30279

30280

30289

30290

30291

30292

30293

30297

30298

30299

30300

30265 **DESCRIPTION**

To send a message to a batch job is to request that a server write a message string into one or more output files of the batch job. A message is sent to a batch job by a request to the batch server that manages the batch job. The *qmsg* utility is a user-accessible batch client that requests the sending of messages to one or more batch jobs.

The *qmsg* utility shall write messages into the files of batch jobs by sending a *Job Message Request* to the batch server that manages the batch job. The *qmsg* utility shall not directly write the message into the files of the batch job.

The *qmsg* utility shall send a *Job Message Request* for those batch jobs, and only those batch jobs, for which a batch *job_identifier* is presented to the utility.

The *qmsg* utility shall send *Job Message Requests* for batch jobs in the order in which their batch *job_identifiers* are presented to the utility.

If the *qmsg* utility fails to process any batch *job_identifier* successfully, the utility shall proceed to process the remaining batch *job_identifiers*, if any.

The *qmsg* utility shall not exit before a *Job Message Request* has been sent to the server that manages the batch job that corresponds to each successfully processed batch *job_identifier*.

30281 OPTIONS

The *qmsg* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported by the implementation:

30285 -E Specify that the message is written to the standard error of each batch job.

30286 The *qmsg* utility shall write the message into the standard error of the batch job.

30287 –O Specify that the message is written to the standard output of each batch job.

The *qmsg* utility shall write the message into the standard output of the batch job.

If neither the $-\mathbf{O}$ nor the $-\mathbf{E}$ option is presented to the *qmsg* utility, the utility shall write the message into an implementation-defined file. The conformance document for the implementation shall describe the name and location of the implementation-defined file. If both the $-\mathbf{O}$ and the $-\mathbf{E}$ options are presented to the *qmsg* utility, then the utility shall write the messages to both standard output and standard error.

30294 OPERANDS

The *qmsg* utility shall accept a minimum of two operands, *message_string* and one or more batch *job_identifiers*.

The *message_string* operand shall be the string to be written to one or more output files of the batch job followed by a <newline>. If the string contains <blank>s, then the application shall ensure that the string is quoted. The *message_string* shall be encoded in the portable character set (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.1, Portable Character Set).

All remaining operands are batch *job_identifiers* that conform to the syntax for a batch *job_identifier* (see Section 3.3.1 (on page 122)).

Utilities qmsg

30303 **STDIN** 30304 Not used. 30305 INPUT FILES None. 30306 30307 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *qmsg*: 30308 **LANG** 30309 Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2. 30310 Internationalization Variables for the precedence of internationalization variables 30311 used to determine the values of locale categories.) 30312 LC ALL If set to a non-empty string value, override the values of all the other 30313 internationalization variables. 30314 Determine the locale for the interpretation of sequences of bytes of text data as LC_CTYPE 30315 characters (for example, single-byte as opposed to multi-byte characters in 30316 arguments). 30317 LC_MESSAGES 30318 Determine the locale that should be used to affect the format and contents of 30319 30320 diagnostic messages written to standard error. 30321 *LOGNAME* Determine the login name of the user. 30322 ASYNCHRONOUS EVENTS Default. 30323 30324 STDOUT None. 30325 30326 STDERR 30327 The standard error shall be used only for diagnostic messages. 30328 OUTPUT FILES None. 30329 30330 EXTENDED DESCRIPTION 30331 None. 30332 EXIT STATUS The following exit values shall be returned: 30333 30334 Successful completion. 30335 >0 An error occurred. **30336 CONSEQUENCES OF ERRORS** In addition to the default behavior, the qmsg utility shall not be required to write a diagnostic 30337 message to standard error when the error reply received from a batch server indicates that the 30338 batch job_identifier does not exist on the server. Whether or not the qmsg utility waits to output 30339

the diagnostic message while attempting to locate the job on other servers is implementation-

30340

30341

defined.

qmsg Utilities

30342 APPLICATION USAGE

30343 None.

30344 EXAMPLES

30345 None.

30346 RATIONALE

The *qmsg* utility allows users to write messages into the output files of running jobs. Users, including operators and administrators, have a number of occasions when they want to place messages in the output files of a batch job. For example, if a disk that is being used by a batch job is showing errors, the operator might note this in the standard error stream of the batch job.

The options of the *qmsg* utility provide users with the means of placing the message in the output stream of their choice. The default output stream for the message—if the user does not designate an output stream—is implementation-defined, since many implementations will provide, as an extension to this volume of IEEE Std 1003.1-2001, a log file that shows the history of utility execution.

If users wish to send a message to a set of jobs that meet a selection criteria, the *qselect* utility can be used to acquire the appropriate list of job identifiers.

30358 The –E option allows users to place the message in the standard error stream of the batch job.

30359 The **–O** option allows users to place the message in the standard output stream of the batch job.

Historically, the *qmsg* utility is an existing practice in the offerings of one or more implementors of an NQS-derived batch system. The utility has been found to be useful enough that it deserves to be included in this volume of IEEE Std 1003.1-2001.

30363 FUTURE DIRECTIONS

30364 None.

30365 SEE ALSO

30366 Chapter 3 (on page 101), qselect

30367 CHANGE HISTORY

30368 Derived from IEEE Std 1003.2d-1994.

30369 **Issue 6**

30370 The *LC_TIME* and *TZ* entries are removed from the ENVIRONMENT VARIABLES section.

Utilities qrerun

30371 **NAME**

30372 qrerun — rerun batch jobs

30373 SYNOPSIS

30374 BE qrerun job_identifier ...

30375

30376 **DESCRIPTION**

To rerun a batch job is to terminate the session leader of the batch job, delete any associated checkpoint files, and return the batch job to the batch queued state. A batch job is rerun by a request to the batch server that manages the batch job. The *qrerun* utility is a user-accessible batch client that requests the rerunning of one or more batch jobs.

The *qrerun* utility shall rerun those batch jobs for which a batch *job_identifier* is presented to the utility.

The *qrerun* utility shall rerun batch jobs in the order in which their batch *job_identifiers* are presented to the utility.

If the *qrerun* utility fails to process any batch *job_identifier* successfully, the utility shall proceed to process the remaining batch *job_identifiers*, if any.

The *qrerun* utility shall rerun batch jobs by sending a *Rerun Job Request* to the batch server that manages each batch job.

For each successfully processed batch *job_identifier*, the *qrerun* utility shall have rerun the corresponding batch job at the time the utility exits.

30391 OPTIONS

30392 None.

30393 OPERANDS

The *qrerun* utility shall accept one or more operands that conform to the syntax for a batch *job_identifier* (see Section 3.3.1 (on page 122)).

30396 **STDIN**

30397 Not used.

30398 INPUT FILES

30399 None.

30400 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *qrerun*:

30402 LANG Provide a default value for the internationalization variables that are unset or null.
30403 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
30404 Internationalization Variables for the precedence of internationalization variables
30405 used to determine the values of locale categories.)

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

30408 *LC_CTYPE* 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).

30411 LC MESSAGES

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

qrerun Utilities

30414 *LOGNAME* Determine the login name of the user. 30415 ASYNCHRONOUS EVENTS 30416 Default. **30417 STDOUT** 30418 None. 30419 STDERR 30420 The standard error shall be used only for diagnostic messages. 30421 OUTPUT FILES 30422 None. 30423 EXTENDED DESCRIPTION 30424 None. 30425 EXIT STATUS 30426 The following exit values shall be returned: Successful completion. 30427 >0 An error occurred. 30428 30429 CONSEQUENCES OF ERRORS 30430 In addition to the default behavior, the *qrerun* utility shall not be required to write a diagnostic 30431 message to standard error when the error reply received from a batch server indicates that the 30432 batch job_identifier does not exist on the server. Whether or not the qrerun utility waits to output the diagnostic message while attempting to locate the job on other servers is implementation-30433 defined. 30434 30435 APPLICATION USAGE 30436 None. 30437 EXAMPLES 30438 None. 30439 RATIONALE 30440 The *qrerun* utility allows users to cause jobs in the running state to exit and rerun. 30441 The *qrerun* utility is a new utility, *vis-a-vis* existing practice, that has been defined in this volume of IEEE Std 1003.1-2001 to correct user-perceived deficiencies in the existing practice. 30442 30443 FUTURE DIRECTIONS None. 30444 30445 SEE ALSO 30446 Chapter 3 (on page 101) 30447 CHANGE HISTORY Derived from IEEE Std 1003.2d-1994. 30448

30449 Issue 6

30450

The *LC_TIME* and *TZ* entries are removed from the ENVIRONMENT VARIABLES section.

qrls **Utilities**

30451 **NAME**

qrls — release batch jobs 30452

30453 SYNOPSIS

qrls [-h hold list] job identifier ... 30454 BE

30455

30464

30475 30476

30477

30478

30479 30480

30481 30482

30483

30484

30485

30486

30487

30488 30489

30490

30456 **DESCRIPTION**

A batch job might have one or more holds, which prevent the batch job from executing. A batch 30457 job from which all the holds have been removed becomes eligible for execution and is said to 30458 30459 have been released. A batch job hold is removed by sending a request to the batch server that 30460 manages the batch job. The *qrls* utility is a user-accessible client of batch services that requests holds be removed from one or more batch jobs. 30461

The qrls utility shall remove one or more holds from those batch jobs for which a batch 30462 30463 *job_identifier* is presented to the utility.

The *qrls* utility shall remove holds from batch jobs in the order in which their batch *job_identifiers* 30465 are presented to the utility.

30466 If the qrls utility fails to process a batch job_identifier successfully, the utility shall proceed to 30467 process the remaining batch *job_identifiers*, if any.

The *qrls* utility shall remove holds on each batch job by sending a *Release Job Request* to the batch 30468 30469 server that manages the batch job.

30470 The qrls utility shall not exit until the holds have been removed from the batch job corresponding to each successfully processed batch *job_identifier*. 30471

30472 OPTIONS

The qrls utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 30473 12.2, Utility Syntax Guidelines. 30474

The following option shall be supported by the implementation:

-h hold list Define the types of holds to be removed from the batch job.

The qrls -h option shall accept a value for the hold list option-argument that is a string of alphanumeric characters in the portable character set (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.1, Portable Character Set).

The *qrls* utility shall accept a value for the *hold_list* option-argument that is a string of one or more of the characters 'u', 's', or 'o', or the single character 'n'.

For each unique character in the *hold_list* option-argument, the *qrls* utility shall add a value to the *Hold_Types* attribute of the batch job as follows, each representing a different hold type:

USER u

SYSTEM

OPERATOR

If any of these characters are duplicated in the *hold_list* option-argument, the duplicates shall be ignored.

An existing *Hold_Types* attribute can be cleared by the following hold type:

30491 NO_HOLD **qrls** Utilities

30492 30493		The <i>qrls</i> utility shall consider it an error if any hold type other than 'n' is combined with hold type 'n'.
30494 30495 30496 30497		Strictly conforming applications shall not repeat any of the characters 'u', 's', 'o', or 'n' within the <i>hold_list</i> option-argument. The <i>qrls</i> utility shall permit the repetition of characters, but shall not assign additional meaning to the repeated characters.
30498 30499 30500		An implementation may define other hold types. The conformance document for an implementation shall describe any additional hold types, how they are specified, their internal behavior, and how they affect the behavior of the utility.
30501 30502		If the $-\mathbf{h}$ option is not presented to the <i>qrls</i> utility, the implementation shall remove the USER hold in the <i>Hold_Types</i> attribute.
30503 OI 30504 30505	_	lity shall accept one or more operands that conform to the syntax for a batch (see Section 3.3.1 (on page 122)).
30506 ST 30507	'DIN Not used.	
30508 IN 30509	PUT FILES None.	
30510 EN 30511	IVIRONMENT VA The followin	ARIABLES ag environment variables shall affect the execution of <i>qrls</i> :
30512 30513 30514 30515	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
30516 30517	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
30518 30519 30520	LC_CTYPE	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).
30521	LC_MESSAC	GES
30522 30523		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
30524	LOGNAME	Determine the login name of the user.
30525 AS 30526	SYNCHRONOUS I Default.	EVENTS
30527 ST 30528	DOUT None.	
30529 ST		
30530		d error shall be used only for diagnostic messages.
30531 OU 30532	U TPUT FILES None.	

Utilities qrls

30533 EXTENDED DESCRIPTION

30534 None.

30535 EXIT STATUS

30536 The following exit values shall be returned:

30537 0 Successful completion.

30538 >0 An error occurred.

30539 CONSEQUENCES OF ERRORS

In addition to the default behavior, the *qrls* utility shall not be required to write a diagnostic message to standard error when the error reply received from a batch server indicates that the batch *job_identifier* does not exist on the server. Whether or not the *qrls* utility waits to output the diagnostic message while attempting to locate the job on other servers is implementation-

30544 defined.

30545 APPLICATION USAGE

30546 None.

30547 EXAMPLES

30548 None.

30549 RATIONALE

30550 The *qrls* utility allows users, operators, and administrators to remove holds from jobs.

The *qrls* utility does not support any job selection options or wildcard arguments. Users may acquire a list of jobs selected by attributes using the *qselect* utility. For example, a user could select all of their held jobs.

The -h option allows the user to specify the type of hold that is to be removed. This option allows for USER, SYSTEM, OPERATOR, and implementation-defined hold types. The batch server that manages the batch job will verify whether the user is authorized to remove the specified hold for the batch job. If more than one type of hold has been placed on the batch job, a user may wish to remove only some of them.

Mail is not required on release because the administrator has the tools and libraries to build this option if required.

The *qrls* utility is a new utility *vis-a-vis* existing practice; it has been defined in this volume of IEEE Std 1003.1-2001 as the natural complement to the *qhold* utility.

30563 FUTURE DIRECTIONS

30564 None.

30565 SEE ALSO

30566 Chapter 3 (on page 101), qhold, qselect

30567 CHANGE HISTORY

30568 Derived from IEEE Std 1003.2d-1994.

30569 Issue 6

30570 The *LC_TIME* and *TZ* entries are removed from the ENVIRONMENT VARIABLES section.

qselect Utilities

```
30571 NAME
30572 qselect — select batch jobs
30573 SYNOPSIS
```

```
qselect [-a [op]date_time] [-A account_string] [-c [op]interval]

[-h hold_list] [-l resource_list] [-N name] [-p [op]priority]

[-q destination] [-r y|n] [-s states] [-u user list]
```

30578 **DESCRIPTION**

30577

30579 30580

30581

30585

30586

30587

30588

30589 30590

30591

30592

30593 30594

30595

30596 30597

30598

30600 30601

30602

30603

30604

30605 30606

30607

30608

To select a set of batch jobs is to return the batch *job_identifiers* for each batch job that meets a list of selection criteria. A set of batch jobs is selected by a request to a batch server. The *qselect* utility is a user-accessible batch client that requests the selection of batch jobs.

Upon successful completion, the *qselect* utility shall have returned a list of zero or more batch *job_identifiers* that meet the criteria specified by the options and option-arguments presented to the utility.

The *qselect* utility shall select batch jobs by sending a *Select Jobs Request* to a batch server. The *qselect* utility shall not exit until the server replies to each request generated.

For each option presented to the *qselect* utility, the utility shall restrict the set of selected batch jobs as described in the OPTIONS section.

The *qselect* utility shall not restrict selection of batch jobs except by authorization and as required by the options presented to the utility.

When an option is specified with a mandatory or optional *op* component to the optionargument, then *op* shall specify a relation between the value of a certain batch job attribute and the *value* component of the option-argument. If an *op* is allowable on an option, then the description of the option letter indicates the *op* as either mandatory or optional. Acceptable strings for the *op* component, and the relation the string indicates, are shown in the following list:

- .eq. The value represented by the attribute of the batch job is equal to the value represented by the option-argument.
- .ge. The value represented by the attribute of the batch job is greater than or equal to the value represented by the option-argument.
- .gt. The value represented by the attribute of the batch job is greater than the value represented by the option-argument.
- .1t. The value represented by the attribute of the batch job is less than the value represented by the option-argument.
- .le. The value represented by the attribute of the batch job is less than or equal to the value represented by the option-argument.
- .ne. The value represented by the attribute of the batch job is not equal to the value represented by the option-argument.

30609 OPTIONS

The *qselect* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported by the implementation:

30613 —**a** [op]date_time

Restrict selection to a specific time, or a range of times.

Utilities qselect

30615			lect utility shall select only batch jobs for which the value of the	
30616 30617		Execution_Time attribute is related to the Epoch equivalent of the local time expressed by the value of the date_time component of the option-argument in the		
30618	:	manner indicated by the value of the <i>op</i> component of the option-argument.		
30619 30620		The <i>qselect</i> utility shall accept a <i>date_time</i> component of the option-argument that conforms to the syntax of the <i>time</i> operand of the <i>touch</i> utility.		
30621 30622 30623	,	the utilit	component of the option-argument is not presented to the <i>qselect</i> utility, by shall select batch jobs for which the <i>Execution_Time</i> attribute is equal to <i>time</i> component of the option-argument.	
30624 30625			omparing times, the <i>qselect</i> utility shall use the following definitions for the onent of the option-argument:	
30626 30627 30628		.eq.	The time represented by value of the <i>Execution_Time</i> attribute of the batch job is equal to the time represented by the <i>date_time</i> component of the option-argument.	
30629 30630 30631		.ge.	The time represented by value of the <i>Execution_Time</i> attribute of the batch job is after or equal to the time represented by the <i>date_time</i> component of the option-argument.	
30632 30633 30634		.gt.	The time represented by value of the <i>Execution_Time</i> attribute of the batch job is after the time represented by the <i>date_time</i> component of the option-argument.	
30635 30636 30637		.lt.	The time represented by value of the <i>Execution_Time</i> attribute of the batch job is before the time represented by the <i>date_time</i> component of the option-argument.	
30638 30639 30640		.le.	The time represented by value of the <i>Execution_Time</i> attribute of the batch job is before or equal to the time represented by the <i>date_time</i> component of the option-argument.	
30641 30642 30643		.ne.	The time represented by value of the <i>Execution_Time</i> attribute of the batch job is not equal to the time represented by the <i>date_time</i> component of the option-argument.	
30644 30645			ect utility shall accept the defined character strings for the <i>op</i> component of on-argument.	
30646 30647	-A account_st	_	selection to the batch jobs charging a specified account.	
30648 30649 30650		Account_	<i>lect</i> utility shall select only batch jobs for which the value of the _Name attribute of the batch job matchs the value of the account_string rgument.	
30651	i	The synt	ax of the account_string option-argument is unspecified.	
30652 30653	−c [op]interval		selection to batch jobs within a range of checkpoint intervals.	
30654 30655 30656	;	attribute	ect utility shall select only batch jobs for which the value of the <i>Checkpoint</i> e relates to the value of the <i>interval</i> component of the option-argument in ner indicated by the value of the <i>op</i> component of the option-argument.	
30657 30658		_	o component of the option-argument is omitted, the <i>qselect</i> utility shall atch jobs for which the value of the <i>Checkpoint</i> attribute is equal to the value	

qselect Utilities

30659		of the in	nterval component of the option-argument.
30660 30661			comparing checkpoint intervals, the <i>qselect</i> utility shall use the following ons for the <i>op</i> component of the option-argument:
30662 30663		.eq.	The value of the <i>Checkpoint</i> attribute of the batch job equals the value of the <i>interval</i> component of the option-argument.
30664 30665		.ge.	The value of the <i>Checkpoint</i> attribute of the batch job is greater than or equal to the value of the <i>interval</i> component option-argument.
30666 30667		.gt.	The value of the <i>Checkpoint</i> attribute of the batch job is greater than the value of the <i>interval</i> component option-argument.
30668 30669		.lt.	The value of the <i>Checkpoint</i> attribute of the batch job is less than the value of the <i>interval</i> component option-argument.
30670 30671		.le.	The value of the <i>Checkpoint</i> attribute of the batch job is less than or equal to the value of the <i>interval</i> component option-argument.
30672 30673		.ne.	The value of the <i>Checkpoint</i> attribute of the batch job does not equal the value of the <i>interval</i> component option-argument.
30674 30675			<i>lect</i> utility shall accept the defined character strings for the <i>op</i> component of ion-argument.
30676 30677		The ord to be:	lering relationship for the values of the interval option-argument is defined
30678		`n' .g	gt. `s' .gt. `c=minutes' .ge. `c'
30679 30680			comparing <i>Checkpoint</i> attributes with an interval having the value of the haracter 'u', only equality or inequality are valid comparisons.
30681	- h hold_list	Restrict	selection to batch jobs that have a specific type of hold.
30682 30683		•	lect utility shall select only batch jobs for which the value of the <code>Hold_Types</code> e matches the value of the <code>hold_list</code> option-argument.
30684 30685 30686		string o	<i>lect</i> – h option shall accept a value for the <i>hold_list</i> option-argument that is a of alphanumeric characters in the portable character set (see the Base ons volume of IEEE Std 1003.1-2001, Section 6.1, Portable Character Set).
30687 30688 30689			<i>lect</i> utility shall accept a value for the <i>hold_list</i> option-argument that is a of one or more of the characters 'u', 's', or 'o', or the single character
30690 30691			nique character in the <i>hold_list</i> option-argument of the <i>qselect</i> utility is as follows, each representing a different hold type:
30692		u US	ER
30693		s SY	STEM
30694		o OP	PERATOR
30695 30696			of these characters are duplicated in the <i>hold_list</i> option-argument, the tes shall be ignored.
30697		Tr)	elect utility shall consider it an error if any hold type other than 'n' is

Utilities qselect

30699 Strictly conforming applications shall not repeat any of the characters 'u', 's', 30700 'o', or 'n' within the *hold_list* option-argument. The *qselect* utility shall permit the repetition of characters, but shall not assign additional meaning to the repeated 30701 characters. 30702 An implementation may define other hold types. The conformance document for 30703 an implementation shall describe any additional hold types, how they are 30704 30705 specified, their internal behavior, and how they affect the behavior of the utility. -l resource list 30706 Restrict selection to batch jobs with specified resource limits and attributes. 30707 The *qselect* utility shall accept a *resource_list* option-argument with the following 30708 30709 syntax: resource name op value [,,resource name op value,, ...] 30710 When comparing resource values, the *qselect* utility shall use the following 30711 definitions for the *op* component of the option-argument: 30712 The value of the resource of the same name in the *Resource List* attribute 30713 .eq. of the batch job equals the value of the value component of the option-30714 argument. 30715 30716 The value of the resource of the same name in the *Resource List* attribute .ge. 30717 of the batch job is greater than or equal to the value of the value 30718 component of the option-argument. The value of the resource of the same name in the *Resource_List* attribute 30719 .gt. 30720 of the batch job is greater than the value of the value component of the option-argument. 30721 30722 .lt. The value of the resource of the same name in the *Resource_List* attribute 30723 of the batch job is less than the value of the value component of the option-argument. 30724 The value of the resource of the same name in the *Resource List* attribute 30725 .ne. of the batch job does not equal the value of the value component of the 30726 option-argument. 30727 The value of the resource of the same name in the *Resource_List* attribute 30728 .le. of the batch job is less than or equal to the value of the value component 30729 of the option-argument. 30730 When comparing the limit of a *Resource_List* attribute with the *value* component of 30731 the option-argument, if the limit, the value, or both are non-numeric, only equality 30732 or inequality are valid comparisons. 30733 The *qselect* utility shall select only batch jobs for which the values of the 30734 30735 resource names listed in the resource list option-argument match the corresponding limits of the *Resource_List* attribute of the batch job. 30736 30737 Limits of resource_names present in the Resource_List attribute of the batch job that have no corresponding values in the resource_list option-argument shall not be 30738 considered when selecting batch jobs. 30739 -N name Restrict selection to batch jobs with a specified name. 30740 30741 The *qselect* utility shall select only batch jobs for which the value of the *Job_Name* attribute matches the value of the name option-argument. The string specified in 30742

qselect Utilities

30743 the *name* option-argument shall be passed, uninterpreted, to the server. This allows 30744 an implementation to match "wildcard" patterns against batch job names. 30745 An implementation shall describe in the conformance document the format it supports for matching against the *Job_Name* attribute. 30746 30747 -p [op]priority Restrict selection to batch jobs of the specified priority or range of priorities. 30748 30749 The *qselect* utility shall select only batch jobs for which the value of the *Priority* attribute of the batch job relates to the value of the priority component of the 30750 option-argument in the manner indicated by the value of the op component of the 30751 option-argument. 30752 30753 If the op component of the option-argument is omitted, the qselect utility shall 30754 select batch jobs for which the value of the *Priority* attribute of the batch job is equal to the value of the *priority* component of the option-argument. 30755 30756 When comparing priority values, the *qselect* utility shall use the following definitions for the *op* component of the option-argument: 30757 The value of the *Priority* attribute of the batch job equals the value of the 30758 .eq. *priority* component of the option-argument. 30759 The value of the *Priority* attribute of the batch job is greater than or equal 30760 .ge. to the value of the *priority* component option-argument. 30761 The value of the *Priority* attribute of the batch job is greater than the value 30762 .gt. 30763 of the *priority* component option-argument. The value of the *Priority* attribute of the batch job is less than the value of .lt. 30764 30765 the *priority* component option-argument. The value of the *Priority* attribute of the batch job is less than or equal to 30766 .1t.. the value of the *priority* component option-argument. 30767 .ne. The value of the *Priority* attribute of the batch job does not equal the value 30768 30769 of the *priority* component option-argument. 30770 -q destination Restrict selection to the specified batch queue or server, or both. 30771 The *qselect* utility shall select only batch jobs that are located at the destination 30772 indicated by the value of the *destination* option-argument. 30773 The destination defines a batch queue, a server, or a batch queue at a server. 30774 30775 The *qselect* utility shall accept an option-argument for the $-\mathbf{q}$ option that conforms 30776 to the syntax for a destination. If the $-\mathbf{q}$ option is not presented to the *qselect* utility, the utility shall select batch jobs from all batch queues at the default batch server. 30777 If the option-argument describes only a batch queue, the *qselect* utility shall select 30778 only batch jobs from the batch queue of the specified name at the default batch 30779 30780 server. The means by which qselect determines the default server is 30781 implementation-defined. If the option-argument describes only a batch server, the *qselect* utility shall select 30782 30783 batch jobs from all the batch queues at that batch server. 30784 If the option-argument describes both a batch queue and a batch server, the qselect utility shall select only batch jobs from the specified batch queue at the specified 30785

Utilities qselect

30786		server.
30787	-r y n	Restrict selection to batch jobs with the specified rerunability status.
30788 30789		The <i>qselect</i> utility shall select only batch jobs for which the value of the <i>Rerunable</i> attribute of the batch job matches the value of the option-argument.
30790 30791 30792		The <i>qselect</i> utility shall accept a value for the option-argument that consists of either the single character $'y'$ or the single character $'n'$. The character $'y'$ represents the value TRUE, and the character $'n'$ represents the value FALSE.
30793	-s states	Restrict selection to batch jobs in the specified states.
30794 30795		The <i>qselect</i> utility shall accept an option-argument that consists of any combination of the characters $' e' , ' q' , ' r' , ' w' , ' h' , and ' t' .$
30796 30797 30798		Conforming applications shall not repeat any character in the option-argument. The <i>qselect</i> utility shall permit the repetition of characters in the option-argument, but shall not assign additional meaning to repeated characters.
30799 30800		The <i>qselect</i> utility shall interpret the characters in the <i>states</i> option-argument as follows:
30801		e Represents the EXITING state.
30802		q Represents the QUEUED state.
30803		r Represents the RUNNING state.
30804		t Represents the TRANSITING state.
30805		h Represents the HELD state.
30806		w Represents the WAITING state.
30807 30808		For each character in the <i>states</i> option-argument, the <i>qselect</i> utility shall select batch jobs in the corresponding state.
30809	- u user_list	Restrict selection to batch jobs owned by the specified user names.
30810 30811		The <i>qselect</i> utility shall select only the batch jobs of those users specified in the <i>user_list</i> option-argument.
30812 30813		The <i>qselect</i> utility shall accept a <i>user_list</i> option-argument that conforms to the following syntax:
30814		username[@host][,,username[@host],,]
30815 30816		The <i>qselect</i> utility shall accept only one user name that is missing a corresponding host name. The <i>qselect</i> utility shall accept only one user name per named host.
30817 OPERA 30818	N DS None.	
30819 STDIN 30820	Not used.	
30821 INPUT 30822	FILES None.	

qselect Utilities

30823 30824	ENVIRONMENT VA The followin	RIABLES g environment variables shall affect the execution of <i>qselect</i> :		
30825 30826 30827 30828	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)		
30829 30830	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.		
30831 30832 30833	LC_CTYPE	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).		
30834	LC_MESSAC	GES .		
30835 30836	Ec_ivizion ic	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.		
30837	LOGNAME	Determine the login name of the user.		
30838 30839	TZ	Determine the timezone used to interpret the <i>date-time</i> option-argument. If <i>TZ</i> is unset or null, an unspecified default timezone shall be used.		
30840 30841	ASYNCHRONOUS I Default.	EVENTS		
30842 30843	STDOUT The qselect ut	tility shall write zero or more batch <i>job_identifier</i> s to standard output.		
30844 30845	The <i>qselect</i> u space.	The <i>qselect</i> utility shall separate the batch <i>job_identifiers</i> written to standard output by white space.		
30846	The <i>qselect</i> ut	The <i>qselect</i> utility shall write batch <i>job_identifiers</i> in the following format:		
30847	sequence_r	sequence_number.server_name@server		
30848 30849	STDERR The standard	l error shall be used only for diagnostic messages.		
30850	OUTPUT FILES			
30851	None.			
30852	EXTENDED DESCRI	IPTION		
30853	None.			
30854 30855	EXIT STATUS The followin	g exit values shall be returned:		
30856	0 Successi	ful completion.		
30857				
30858	CONSEQUENCES O	FERRORS		

30859

Default.

Utilities qselect

30860 APPLICATION USAGE

30861 None.

30862 EXAMPLES

The following example shows how a user might use the *qselect* utility in conjunction with the *qdel* utility to delete all of his or her jobs in the queued state without affecting any jobs that are already running:

30866 qdel \$(qselect -s q)

30867 or:

30868 qselect -s q | xargs qdel

30869 RATIONALE

30870

30871

30872 30873

30874

30875

30876

30877

30878 30879

30880

30881

30882 30883

30884

30885

30886 30887

30888

30889

30890 30891

30892

30893

30894

30895

30896

The *qselect* utility allows users to acquire a list of job identifiers that match user-specified selection criteria. The list of identifiers returned by the *qselect* utility conforms to the syntax of the batch job identifier list processed by a utility such as *qmove*, *qdel*, and *qrls*. The *qselect* utility is thus a powerful tool for causing another batch system utility to act upon a set of jobs that match a list of selection criteria.

The options of the *qselect* utility let the user apply a number of useful filters for selecting jobs. Each option further restricts the selection of jobs. Many of the selection options allow the specification of a relational operator. The FORTRAN-like syntax of the operator—that is, ".lt."—was chosen rather than the C-like "<=" meta-characters.

The -a option allows users to restrict the selected jobs to those that have been submitted (or altered) to wait until a particular time. The time period is determined by the argument of this option, which includes both a time and an operator—it is thus possible to select jobs waiting until a specific time, jobs waiting until after a certain time, or those waiting for a time before the specified time.

The **–A** option allows users to restrict the selected jobs to those that have been submitted (or altered) to charge a particular account.

The -c option allows users to restrict the selected jobs to those whose checkpointing interval falls within the specified range.

The -l option allows users to select those jobs whose resource limits fall within the range indicated by the value of the option. For example, a user could select those jobs for which the CPU time limit is greater than two hours.

The -N option allows users to select jobs by job name. For instance, all the parts of a task that have been divided in parallel jobs might be given the same name, and thus manipulated as a group by means of this option.

The $-\mathbf{q}$ option allows users to select jobs in a specified queue.

The -r option allows users to select only those jobs with a specified rerun criteria. For instance, a user might select only those jobs that can be rerun for use with the *qrerun* utility.

The -s option allows users to select only those jobs that are in a certain state.

The **–u** option allows users to select jobs that have been submitted to execute under a particular account.

The selection criteria provided by the options of the *qselect* utility allow users to select jobs based on all the appropriate attributes that can be assigned to jobs by the *qsub* utility.

Historically, the *qselect* utility has not been a part of existing practice; it is an improvement that has been introduced in this volume of IEEE Std 1003.1-2001.

qselect Utilities

30904 FUTURE DIRECTIONS

30905 None.

30906 SEE ALSO

qdel, qrerun, qrls, qselect, qsub, touch, Chapter 3 (on page 101)

30908 CHANGE HISTORY

30909 Derived from IEEE Std 1003.2d-1994.

Utilities qsig

30910 NAME

30911 qsig — signal batch jobs

30912 SYNOPSIS

30913 BE qsig [-s signal] job_identifier ..

30914

30915 **DESCRIPTION**

To signal a batch job is to send a signal to the session leader of the batch job. A batch job is signaled by sending a request to the batch server that manages the batch job. The *qsig* utility is a user-accessible batch client that requests the signaling of a batch job.

The *qsig* utility shall signal those batch jobs for which a batch *job_identifier* is presented to the utility. The *qsig* utility shall not signal any batch jobs whose batch *job_identifiers* are not presented to the utility.

The *qsig* utility shall signal batch jobs in the order in which the corresponding batch *job_identifiers* are presented to the utility. If the *qsig* utility fails to process a batch *job_identifier* successfully, the utility shall proceed to process the remaining batch *job_identifiers*, if any.

The *qsig* utility shall signal batch jobs by sending a *Signal Job Request* to the batch server that manages the batch job.

For each successfully processed batch *job_identifier*, the *qsig* utility shall have received a completion reply to each *Signal Job Request* sent to a batch server at the time the utility exits.

30929 OPTIONS

30932

30934 30935

30936 30937

30938

The *qsig* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following option shall be supported by the implementation:

30933 —s signal Define the signal to be sent to the batch job.

The *qsig* utility shall accept a *signal* option-argument that is either a symbolic signal name or an unsigned integer signal number (see the POSIX.1-1990 standard, Section 3.3.1.1). The *qsig* utility shall accept signal names for which the SIG prefix has been omitted.

If the *signal* option-argument is a signal name, the *qsig* utility shall send that name.

If the *signal* option-argument is a number, the *qsig* utility shall send the signal value represented by the number.

If the **–s** option is not presented to the *qsig* utility, the utility shall send the signal SIGTERM to each signaled batch job.

30943 OPERANDS

The *qsig* utility shall accept one or more operands that conform to the syntax for a batch *job_identifier* (see Section 3.3.1 (on page 122)).

30946 STDIN

Not used.

30948 INPUT FILES

30949 None.

qsig Utilities

30950 ENVIRONMENT VARIABLES 30951 The following environment variables shall affect the execution of *qsig*: 30952 Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 30953 Internationalization Variables for the precedence of internationalization variables 30954 used to determine the values of locale categories.) 30955 LC_ALL If set to a non-empty string value, override the values of all the other 30956 internationalization variables. 30957 Determine the locale for the interpretation of sequences of bytes of text data as LC_CTYPE 30958 characters (for example, single-byte as opposed to multi-byte characters in 30959 arguments). 30960 LC_MESSAGES 30961 Determine the locale that should be used to affect the format and contents of 30962 30963 diagnostic messages written to standard error. *LOGNAME* Determine the login name of the user. 30964 30965 ASYNCHRONOUS EVENTS Default. 30966 30967 **STDOUT** An implementation of the *qsig* utility may write informative messages to standard output. 30968 30969 STDERR 30970 The standard error shall be used only for diagnostic messages. 30971 OUTPUT FILES None. 30972 30973 EXTENDED DESCRIPTION None. 30974 30975 EXIT STATUS 30976 The following exit values shall be returned: Successful completion. 30977 >0 An error occurred. 30978 30979 CONSEQUENCES OF ERRORS In addition to the default behavior, the qsig utility shall not be required to write a diagnostic 30980 30981 message to standard error when the error reply received from a batch server indicates that the batch job_identifier does not exist on the server. Whether or not the qsig utility waits to output the 30982 diagnostic message while attempting to locate the batch job on other servers is implementation-30983 defined. 30984 30985 APPLICATION USAGE None. 30986 30987 EXAMPLES

796

30989 RATIONALE

30988

30990

30992

None.

The *qsig* utility allows users to signal batch jobs.

A user may be unable to signal a batch job with the *kill* utility of the operating system for a number of reasons. First, the process ID of the batch job may be unknown to the user. Second,

Utilities qsig

30993 30994	the processes of the batch job may be on a remote node. However, by virtue of communication between batch nodes, the <i>qsig</i> utility can arrange for the signaling of a process.				
30995 30996	Because a batch job that is not running cannot be signaled, and because the signal may not terminate the batch job, the <i>qsig</i> utility is not a substitute for the <i>qdel</i> utility.				
30997 30998	The options of the <i>qsig</i> utility allow the user to specify the signal that is to be sent to the batch job.				
30999 31000	The $-s$ option allows users to specify a signal by name or by number, and thus override the default signal. The POSIX.1-1990 standard defines signals by both name and number.				
31001 31002	The <i>qsig</i> utility is a new utility, <i>vis-a-vis</i> existing practice; it has been defined in this volume of IEEE Std 1003.1-2001 in response to user-perceived shortcomings in existing practice.				
31003 FUTUR	31003 FUTURE DIRECTIONS				
31004	None.				
31005 SEE AI 31006	SO Chapter 3 (on page 101), kill, qdel				
31007 CHAIN 31008	GE HISTORY Derived from IEEE Std 1003.2d-1994.				
31009 Issue 6					
31010	The LC_TIME and TZ entries are removed from the ENVIRONMENT VARIABLES section.				

qstat Utilities

31011 NAME				
31012	qstat — show status of batch jobs			
31013 SYNO	31013 SYNOPSIS			
31014 BE	qstat [-f] job_identifier			
31015	qstat -Q [-f] destination			
31016	qstat -B [-f] server_name			
31017	NUL ON L			
31018 DESCI 31019	The status of a batch job, batch queue, or batch server is obtained by a request to the server. The	e		
31020	qstat utility is a user-accessible batch client that requests the status of one or more batch jobs			
31021	batch queues, or servers, and writes the status information to standard output.			
31022 31023	For each successfully processed batch <i>job_identifier</i> , the <i>qstat</i> utility shall display information about the corresponding batch job.	1		
31024 31025	For each successfully processed destination, the <i>qstat</i> utility shall display information about the corresponding batch queue.			
31026 31027	For each successfully processed server name, the <i>qstat</i> utility shall display information about the corresponding server.			
31028	The qstat utility shall acquire batch job status information by sending a Job Status Request to a	a		
31029	batch server. The <i>qstat</i> utility shall acquire batch queue status information by sending a <i>Queue</i>			
31030 31031	Status Request to a batch server. The <i>qstat</i> utility shall acquire server status information by sending a <i>Server Status Request</i> to a batch server.	/		
31032 OPTIC	•			
31033 31034	The <i>qstat</i> utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.	1		
31035	The following options shall be supported by the implementation:			
31036	-f Specify that a full display is produced.			
31037	The minimum contents of a full display are specified in the STDOUT section.			
31038	Additional contents and format of a full display are implementation-defined.			
31039	−Q Specify that the operand is a destination.			
31040 31041	The <i>qstat</i> utility shall display information about each batch queue at each destination identified as an operand.	1		
31042	− B Specify that the operand is a server name.			
31043	The <i>qstat</i> utility shall display information about each server identified as ar	1		
31044	operand.			
31045 OPER				
31046 31047	If the $-\mathbf{Q}$ option is presented to the <i>qstat</i> utility, the utility shall accept one or more operands that conform to the syntax for a destination (see Section 3.3.2 (on page 123)).	t		
31048 31049	If the –B option is presented to the <i>qstat</i> utility, the utility shall accept one or more <i>server_name</i> operands.			
31050	If neither the -B nor the -Q option is presented to the <i>qstat</i> utility, the utility shall accept one or more operands that conform to the syntax for a batch ich identifier (see Section 2.2.1 (on page			

122)).

31051

31052

more operands that conform to the syntax for a batch job_identifier (see Section 3.3.1 (on page

Utilities **qstat**

31053 STDIN			
31054	Not used.		
31055 INPUT 31056	FILES None.		
31057 ENVIR 31058	ONMENT VA	ARIABLES ng environment variables shall affect the execution of <i>qstat</i> :	
31059	НОМЕ	Determine the pathname of the user's home directory.	
31060 31061 31062 31063	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)	
31064 31065	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.	
31066 31067 31068	LC_COLLAT	Determine the locale for the behavior of ranges, equivalence classes, and multi- character collating elements within regular expressions.	
31069 31070 31071	LC_CTYPE	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).	
31072	LC_MESSA	GES	
31073 31074		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.	
31075	LC_NUMERIC		
31076 31077	Determine the locale for selecting the radix character used when writing floating- point formatted output.		
31078 ASYNO 31079	CHRONOUS Default.	EVENTS	
31080 STDOU	J T		
31081		nd presented to the <i>qstat</i> utility is a batch <i>job_identifier</i> and the – f option is not	
31082 31083	specified, the <i>qstat</i> utility shall display the following items on a single line, in the stated order, with white space between each item, for each successfully processed operand:		
31084		h job_identifier	
31085	• The batc	h job name	
31086	• The <i>Job_Owner</i> attribute		
31087	The CPU time used by the batch job		
31088	• The batc	h job state	
31089		h job location	
31090 31091	If an operan	ad presented to the <i>qstat</i> utility is a batch <i>job_identifier</i> and the -f option is specified, ity shall display the following items for each success fully processed operand:	

- The batch $job_identifier$

• The batch job name

31092

31093

qstat Utilities

- The *Job_Owner* attribute
- The execution user ID
- The CPU time used by the batch job
- The batch job state

31099

31104

31107

31108

31109

31110

31111

31112

31113

31115

31117

31118

31119

31120 31121

31122

31123 31124

31125

31126 31127

31129

- The batch job location
 - Additional implementation-defined information, if any, about the batch job or batch queue

If an operand presented to the *qstat* utility is a destination, the **-Q** option is specified, and the **-f** option is not specified, the *qstat* utility shall display the following items on a single line, in the stated order, with white space between each item, for each successfully processed operand:

- The batch queue name
 - The maximum number of batch jobs that shall be run in the batch queue concurrently
- The total number of batch jobs in the batch queue
- The status of the batch queue
 - For each state, the number of batch jobs in that state in the batch queue and the name of the state
 - The type of batch queue (execution or routing)

If the operands presented to the *qstat* utility are destinations, the **-Q** option is specified, and the **-f** option is specified, the *qstat* utility shall display the following items for each successfully processed operand:

- The batch queue name
- The maximum number of batch jobs that shall be run in the batch queue concurrently
- The total number of batch jobs in the batch queue
- The status of the batch queue
 - For each state, the number of batch jobs in that state in the batch queue and the name of the state
 - The type of batch queue (execution or routing)
 - Additional implementation-defined information, if any, about the batch queue

If the operands presented to the *qstat* utility are batch server names, the $-\mathbf{B}$ option is specified, and the $-\mathbf{f}$ option is not specified, the *qstat* utility shall display the following items on a single line, in the stated order, with white space between each item, for each successfully processed operand:

- The batch server name
- The maximum number of batch jobs that shall be run in the batch queue concurrently
- The total number of batch jobs managed by the batch server
- The status of the batch server
 - For each state, the number of batch jobs in that state and the name of the state

If the operands presented to the *qstat* utility are server names, the **–B** option is specified, and the **–f** option is specified, the *qstat* utility shall display the following items for each successfully processed operand:

Utilities qstat

31133 The server name The maximum number of batch jobs that shall be run in the batch queue concurrently 31134 The total number of batch jobs managed by the server 31135 The status of the server 31136 For each state, the number of batch jobs in that state and the name of the state 31137 31138 Additional implementation-defined information, if any, about the server 31139 STDERR 31140 The standard error shall be used only for diagnostic messages. 31141 OUTPUT FILES None. 31149 31143 EXTENDED DESCRIPTION None. 31144 31145 EXIT STATUS 31146 The following exit values shall be returned: Successful completion. 31147 An error occurred. 31148 31149 CONSEQUENCES OF ERRORS 31150 In addition to the default behavior, the qstat utility shall not be required to write a diagnostic 31151 message to standard error when the error reply received from a batch server indicates that the 31152 batch job_identifier does not exist on the server. Whether or not the qstat utility waits to output the diagnostic message while attempting to locate the batch job on other servers is 31153 31154 implementation-defined. 31155 APPLICATION USAGE 31156 None. 31157 EXAMPLES 31158 None. 31159 RATIONALE The *qstat* utility allows users to display the status of jobs and list the batch jobs in queues. 31160 The operands of the *qstat* utility may be either job identifiers, queues (specified as destination 31161 31162 identifiers), or batch server names. The $-\mathbf{Q}$ and $-\mathbf{B}$ options, or absence thereof, indicate the nature of the operands. 31163 The other options of the *qstat* utility allow the user to control the amount of information 31164 31165 displayed and the format in which it is displayed. Should a user wish to display the status of a set of jobs that match a selection criteria, the *qselect* utility may be used to acquire such a list. 31166 31167 The -f option allows users to request a "full" display in an implementation-defined format. Historically, the *qstat* utility has been a part of the NQS and its derivatives, the existing practice 31168 on which it is based. 31169 31170 FUTURE DIRECTIONS

31171

None.

qstat Utilities

31172 SEE ALSO
31173 Chapter 3 (on page 101), qselect
31174 CHANGE HISTORY
31175 Derived from IEEE Std 1003.2d-1994.
31176 Issue 6
31177 IEEE PASC Interpretation 1003.2 #191 is applied, removing the following ENVIRONMENT
31178 VARIABLES listed as affecting qstat: COLUMNS, LINES, LOGNAME, TERM, and TZ.
31179 The LC_TIME entry is also removed from the ENVIRONMENT VARIABLES section.

Utilities qsub

```
31180 NAME
31181
              qsub — submit a script
31182 SYNOPSIS
              qsub [-a date time] [-A account string] [-c interval]
31183 BE
31184
                    [-C directive prefix] [-e path name] [-h] [-j join list] [-k keep list]
                    [-m mail options] [-M mail list] [-N name]
31185
                    [-o path name] [-p priority] [-q destination] [-r y|n]
31186
                    [-S path name list] [-u user list] [-v variable list] [-V]
31187
                    [-z][script]
31188
31189
31190 DESCRIPTION
31191
              To submit a script is to create a batch job that executes the script. A script is submitted by a
31192
              request to a batch server. The qsub utility is a user-accessible batch client that submits a script.
              Upon successful completion, the qsub utility shall have created a batch job that will execute the
31193
31194
              submitted script.
              The qsub utility shall submit a script by sending a Queue Job Request to a batch server.
31195
31196
              The qsub utility shall place the value of the following environment variables in the Variable_List
              attribute of the batch job: HOME, LANG, LOGNAME, PATH, MAIL, SHELL, and TZ. The name
31197
              of the environment variable shall be the current name prefixed with the string PBS O.
31198
31199
              Note:
                        If the current value of the HOME variable in the environment space of the qsub utility is
                        /aa/bb/cc, then qsub shall place PBS_O_HOME=/aa/bb/cc in the Variable_List attribute of the
31200
31201
                        batch job.
31202
              In addition to the variables described above, the qsub utility shall add the following variables
              with the indicated values to the variable list:
31203
              PBS_O_WORKDIR
                                    The absolute path of the current working directory of the qsub utility
31204
31205
              PBS_O_HOST
                                    The name of the host on which the qsub utility is running.
31206
31207 OPTIONS
              The qsub utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
31208
              12.2, Utility Syntax Guidelines.
31209
              The following options shall be supported by the implementation:
31210
31211
              -a date_time Define the time at which a batch job becomes eligible for execution.
31212
                           The qsub utility shall accept an option-argument that conforms to the syntax of the
```

time operand of the *touch* utility.

31213

qsub Utilities

31214			Table 4-18 Environme	ent Variable Values (Utilities)	
31215			Variable Name	Value at qsub Time	
31216			PBS_O_HOME	НОМЕ	
31217			PBS_O_HOST	Client host name	
31218			PBS_O_LANG	LANG	
31219			PBS_O_LOGNAME	LOGNAME	
31220			PBS_O_PATH	PATH	
31221			PBS_O_MAIL	MAIL	
31222			PBS_O_SHELL	SHELL	
31223			PBS_O_TZ	TZ	
31224			PBS_O_WORKDIR	Current working directory	
31225 31226				tion of the batch job will add one Section 3.2.2.1 (on page 106).	other variables to
31227		The <i>qsub</i> util	ity shall set the Execution_	<i>Time</i> attribute of the batch jol	o to the number
31228				quivalent to the local time ex	
31229				ment. The Epoch is define	d in the Base
31230		Definitions v	olume of IEEE Std 1003.1	-2001, Section 3.149, Epoch.	
31231				o the <i>qsub</i> utility, the utilit	
31232				job to a time (number of se	conds since the
31233		Epoch) that i	is earlier than the time at v	which the utility exits.	
31234	-A account_				
31235			ccount to which the resou	arce consumption of the bate	h job should be
31236		charged.			
31237		The syntax o	of the <i>account_string</i> option	n-argument is unspecified.	
31238 31239		•	ity shall set the <i>Account_1</i> string option-argument.	Name attribute of the batch job	to the value of
31240 31241			ption is not presented to ne attribute from the attrib	o the <i>qsub</i> utility, the utility butes of the batch job.	shall omit the
31242	–c interval	Define whetl	her the batch job should b	e checkpointed, and if so, hov	v often.
31243 31244		The <i>qsub</i> util		or the interval option-argume	nt that is one of
31245 31246		n	No checkpointing sh (NO_CHECKPOINT).	nall be performed on t	he batch job
31247 31248		S	Checkpointing shall be p down (CHECKPOINT_A	performed only when the batch	ch server is shut
31249 31250 31251		С	Minimum_Cpu_Interval a	neckpointing shall be perf attribute of the batch queue, T_AT_MIN_CPU_INTERVAL,	in units of CPU
31252 31253 31254 31255		c=minutes	of CPU time, or every N	ckpointing shall be performed Minimum_Cpu_Interval minute rgument shall conform to nall be greater than zero.	es, whichever is
31256 31257		-	ity shall set the <i>Checkpoin</i> on-argument.	t attribute of the batch job to	the value of the

Utilities qsub

31258 31259 31260		If the $-c$ option is not presented to the <i>qsub</i> utility, the utility shall set the <i>Checkpoint</i> attribute of the batch job to the single character 'u' (CHECKPOINT_UNSPECIFIED).
31261	-C directive_	
31262		Define the prefix that declares a directive to the <i>qsub</i> utility within the script.
31263 31264		The <i>directive_prefix</i> is not a batch job attribute; it affects the behavior of the <i>qsub</i> utility.
31265		If the -C option is presented to the <i>qsub</i> utility, and the value of the <i>directive_prefix</i>
31266 31267		option-argument is the null string, the utility shall not scan the script file for directives. If the –C option is not presented to the <i>qsub</i> utility, then the value of the
31268		PBS_DPREFIX environment variable is used. If the environment variable is not
31269		defined, then #PBS encoded in the portable character set is the default.
31270	-е path_nam	
31271		Define the path to be used for the standard error stream of the batch job.
31272 31273		The <i>qsub</i> utility shall accept a <i>path_name</i> option-argument which can be preceded by a host name element of the form <i>hostname</i> :.
31274		If the path_name option-argument constitutes an absolute pathname, the qsub
31275		utility shall set the <i>Error_Path</i> attribute of the batch job to the value of the
31276		path_name option-argument.
31277 31278		If the <i>path_name</i> option-argument constitutes a relative pathname and no host name element is specified, the <i>qsub</i> utility shall set the <i>Error_Path</i> attribute of the
31279		batch job to the value of the absolute pathname derived by expanding the
31280		path_name option-argument relative to the current directory of the process
31281		executing <i>qsub</i> .
31282		If the <i>path_name</i> option-argument constitutes a relative pathname and a host name element is specified, the <i>qsub</i> utility shall set the <i>Error_Path</i> attribute of the batch
31283 31284		job to the value of the <i>path_name</i> option-argument without expansion. The host
31285		name element shall be included.
31286		If the path_name option-argument does not include a host name element, the qsub
31287		utility shall prefix the pathname with <i>hostname</i> :, where <i>hostname</i> is the name of the
31288		host upon which the <i>qsub</i> utility is being executed.
31289		If the –e option is not presented to the <i>qsub</i> utility, the utility shall set the <i>Error_Path</i> attribute of the batch job to the host name and path of the current
31290 31291		directory of the submitting process and the default filename.
31292		The default filename for standard error has the following format:
31293		job_name.esequence_number
31294	-h	Specify that a USER hold is applied to the batch job.
31295 31296		The <i>qsub</i> utility shall set the value of the <i>Hold_Types</i> attribute of the batch job to the value USER.
31297 31298		If the -h option is not presented to the <i>qsub</i> utility, the utility shall set the <i>Hold_Types</i> attribute of the batch job to the value NO_HOLD.
31299	- j join_list	Define which streams of the batch job are to be merged. The <i>qsub</i> – j option shall
31300	JJ:	accept a value for the <i>join_list</i> option-argument that is a string of alphanumeric
31301		characters in the portable character set (see the Base Definitions volume of

qsub Utilities

31302		IEEE Std 1003.1-2001, Section 6.1, Portable Character Set).
31303 31304		The <i>qsub</i> utility shall accept a <i>join_list</i> option-argument that consists of one or more of the characters $'e'$ and $'o'$, or the single character $'n'$.
31305 31306		All of the other batch job output streams specified will be merged into the output stream represented by the character listed first in the <i>join_list</i> option-argument.
31307 31308 31309		For each unique character in the <i>join_list</i> option-argument, the <i>qsub</i> utility shall add a value to the <i>Join_Path</i> attribute of the batch job as follows, each representing a different batch job stream to join:
31310		e The standard error of the batch job (JOIN_STD_ERROR).
31311		 The standard output of the batch job (JOIN_STD_OUTPUT).
31312		An existing <i>Join_Path</i> attribute can be cleared by the following join type:
31313		n NO_JOIN
31314 31315		If 'n' is specified, then no files are joined. The <i>qsub</i> utility shall consider it an error if any join type other than 'n' is combined with join type 'n'.
31316 31317 31318 31319		Strictly conforming applications shall not repeat any of the characters $'e'$, $'o'$, or $'n'$ within the $join_list$ option-argument. The $qsub$ utility shall permit the repetition of characters, but shall not assign additional meaning to the repeated characters.
31320 31321 31322		An implementation may define other join types. The conformance document for an implementation shall describe any additional batch job streams, how they are specified, their internal behavior, and how they affect the behavior of the utility.
31323 31324		If the $-\mathbf{j}$ option is not presented to the <i>qsub</i> utility, the utility shall set the value of the <i>Join_Path</i> attribute of the batch job to NO_JOIN.
31325	-k keep_list	Define which output of the batch job to retain on the execution host.
31326 31327 31328		The <i>qsub</i> – k option shall accept a value for the <i>keep_list</i> option-argument that is a string of alphanumeric characters in the portable character set (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.1, Portable Character Set).
31329 31330		The <i>qsub</i> utility shall accept a <i>keep_list</i> option-argument that consists of one or more of the characters $'e'$ and $'o'$, or the single character $'n'$.
31331 31332 31333		For each unique character in the <i>keep_list</i> option-argument, the <i>qsub</i> utility shall add a value to the <i>Keep_Files</i> attribute of the batch job as follows, each representing a different batch job stream to keep:
31334		e The standard error of the batch job (KEEP_STD_ERROR).
31335		• The standard output of the batch job (KEEP_STD_OUTPUT).
31336 31337		If both 'e' and 'o' are specified, then both files are retained. An existing <i>Keep_Files</i> attribute can be cleared by the following keep type:
31338		n NO_KEEP
31339 31340		If 'n' is specified, then no files are retained. The <i>qsub</i> utility shall consider it an error if any keep type other than 'n' is combined with keep type 'n'.
31341 31342 31343		Strictly conforming applications shall not repeat any of the characters $' \in '$, $' \circ '$, or $' n'$ within the $keep_list$ option-argument. The $qsub$ utility shall permit the repetition of characters, but shall not assign additional meaning to the repeated

Utilities qsub

31344		characters.
31345 31346 31347 31348 31349		An implementation may define other keep types. The conformance document for an implementation shall describe any additional keep types, how they are specified, their internal behavior, and how they affect the behavior of the utility. If the $-\mathbf{k}$ option is not presented to the <i>qsub</i> utility, the utility shall set the <i>Keep_Files</i> attribute of the batch job to the value NO_KEEP.
31350 - 31351 31352	- m mail_optio	Define the points in the execution of the batch job at which the batch server that manages the batch job shall send mail about a change in the state of the batch job.
31353 31354 31355		The <i>qsub</i> – m option shall accept a value for the <i>mail_options</i> option-argument that is a string of alphanumeric characters in the portable character set (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.1, Portable Character Set).
31356 31357 31358		The $qsub$ utility shall accept a value for the $mail_options$ option-argument that is a string of one or more of the characters 'e', 'b', and 'a', or the single character 'n'.
31359 31360 31361		For each unique character in the <i>mail_options</i> option-argument, the <i>qsub</i> utility shall add a value to the <i>Mail_Users</i> attribute of the batch job as follows, each representing a different time during the life of a batch job at which to send mail:
31362		e MAIL_AT_EXIT
31363		b MAIL_AT_BEGINNING
31364		a MAIL_AT_ABORT
31365 31366		If any of these characters are duplicated in the ${\it mail_options}$ option-argument, the duplicates shall be ignored.
31367		An existing Mail_Points attribute can be cleared by the following mail type:
31368		n NO_MAIL
31369 31370		If 'n' is specified, then mail is not sent. The <i>qsub</i> utility shall consider it an error if any mail type other than 'n' is combined with mail type 'n'.
31371 31372		Strictly conforming applications shall not repeat any of the characters 'e', 'b', 'a', or 'n' within the <i>mail_options</i> option-argument.
31373 31374 31375 31376 31377		The <i>qsub</i> utility shall permit the repetition of characters, but shall not assign additional meaning to the repeated characters. An implementation may define other mail types. The conformance document for an implementation shall describe any additional mail types, how they are specified, their internal behavior, and how they affect the behavior of the utility.
31378 31379		If the $-\mathbf{m}$ option is not presented to the $qsub$ utility, the utility shall set the $Mail_Points$ attribute to the value MAIL_AT_ABORT.
31380 - 31381	-M mail_list	Define the list of users to which a batch server that executes the batch job shall send mail, if the server sends mail about the batch job.
31382		The syntax of the mail_list option-argument is unspecified.
31383 31384		If the implementation of the $qsub$ utility uses a name service to locate users, the utility should accept the syntax used by the name service.
31385 31386		If the implementation of the $qsub$ utility does not use a name service to locate users, the implementation should accept the following syntax for user names:

qsub Utilities

31387		mail_address[,,mail_address,,]
31388		The interpretation of <i>mail_address</i> is implementation-defined.
31389 31390		The <i>qsub</i> utility shall set the <i>Mail_Users</i> attribute of the batch job to the value of the <i>mail_list</i> option-argument.
31391 31392 31393		If the –M option is not presented to the <i>qsub</i> utility, the utility shall place only the user name and host name for the current process in the <i>Mail_Users</i> attribute of the batch job.
31394	-N name	Define the name of the batch job.
31395 31396 31397 31398		The <i>qsub</i> –N option shall accept a value for the <i>name</i> option-argument that is a string of up to 15 alphanumeric characters in the portable character set (see the Base Definitions volume of IEEE Std 1003.1-2001, Section 6.1, Portable Character Set) where the first character is alphabetic.
31399 31400		The <i>qsub</i> utility shall set the value of the <i>Job_Name</i> attribute of the batch job to the value of the <i>name</i> option-argument.
31401 31402 31403		If the –N option is not presented to the <i>qsub</i> utility, the utility shall set the <i>Job_Name</i> attribute of the batch job to the name of the <i>script</i> argument from which the directory specification if any, has been removed.
31404 31405 31406		If the $-N$ option is not presented to the <i>qsub</i> utility, and the script is read from standard input, the utility shall set the <i>Job_Name</i> attribute of the batch job to the value STDIN.
31407	-o path_nam	
31408		Define the path for the standard output of the batch job.
31409 31410 31411 31412		The <i>qsub</i> utility shall accept a <i>path_name</i> option-argument that conforms to the syntax of the <i>path_name</i> element defined in the System Interfaces volume of IEEE Std 1003.1-2001, which can be preceded by a host name element of the form <i>hostname</i> :.
31413 31414 31415		If the <i>path_name</i> option-argument constitutes an absolute pathname, the <i>qsub</i> utility shall set the <i>Output_Path</i> attribute of the batch job to the value of the <i>path_name</i> option-argument without expansion.
31416 31417 31418 31419		If the <code>path_name</code> option-argument constitutes a relative pathname and no host name element is specified, the <code>qsub</code> utility shall set the <code>Output_Path</code> attribute of the batch job to the pathname derived by expanding the value of the <code>path_name</code> option-argument relative to the current directory of the process executing the <code>qsub</code> .
31420 31421 31422		If the <code>path_name</code> option-argument constitutes a relative pathname and a host name element is specified, the <code>qsub</code> utility shall set the <code>Output_Path</code> attribute of the batch job to the value of the <code>path_name</code> option-argument without expansion.
31423 31424 31425		If the <code>path_name</code> option-argument does not specify a host name element, the <code>qsub</code> utility shall prefix the pathname with <code>hostname</code> :, where <code>hostname</code> is the name of the host upon which the <code>qsub</code> utility is executing.
31426 31427 31428		If the $-\mathbf{o}$ option is not presented to the <i>qsub</i> utility, the utility shall set the <i>Output_Path</i> attribute of the batch job to the host name and path of the current directory of the submitting process and the default filename.
31429		The default filename for standard output has the following format:

Utilities qsub

31430		job_name.osequence_number
31431 31432	- p priority	Define the priority the batch job should have relative to other batch jobs owned by the batch server.
31433 31434		The $qsub$ utility shall set the $Priority$ attribute of the batch job to the value of the $priority$ option-argument.
31435 31436		If the $-\mathbf{p}$ option is not presented to the $qsub$ utility, the value of the $Priority$ attribute is implementation-defined.
31437 31438 31439		The $qsub$ utility shall accept a value for the $priority$ option-argument that conforms to the syntax for signed decimal integers, and which is not less than -1024 and not greater than 1023 .
31440 31441	-q destination	n Define the destination of the batch job.
31442 31443		The destination is not a batch job attribute; it determines the batch server, and possibly the batch queue, to which the <i>qsub</i> utility batch queues the batch job.
31444 31445 31446		The <i>qsub</i> utility shall submit the script to the batch server named by the <i>destination</i> option-argument or the server that owns the batch queue named in the <i>destination</i> option-argument.
31447 31448		The $qsub$ utility shall accept an option-argument for the $-\mathbf{q}$ option that conforms to the syntax for a destination (see Section 3.3.2 (on page 123)).
31449 31450 31451		If the $-\mathbf{q}$ option is not presented to the <i>qsub</i> utility, the <i>qsub</i> utility shall submit the batch job to the default destination. The mechanism for determining the default destination is implementation-defined.
31452	$-\mathbf{r} y n$	Define whether the batch job is rerunnable.
31453 31454		If the value of the option-argument is y , the $qsub$ utility shall set the $Rerunable$ attribute of the batch job to TRUE.
31455 31456		If the value of the option-argument is n , the $qsub$ utility shall set the $Rerunable$ attribute of the batch job to FALSE.
31457 31458		If the $-\mathbf{r}$ option is not presented to the <i>qsub</i> utility, the utility shall set the <i>Rerunable</i> attribute of the batch job to TRUE.
31459 31460	-S path_name	e_list Define the pathname to the shell under which the batch job is to execute.
31461 31462		The <i>qsub</i> utility shall accept a <i>path_name_list</i> option-argument that conforms to the following syntax:
31463		<pre>pathname[@host][,,pathname[@host],,]</pre>
31464 31465		The <i>qsub</i> utility shall allow only one pathname for a given host name. The <i>qsub</i> utility shall allow only one pathname that is missing a corresponding host name.
31466 31467		The <i>qsub</i> utility shall add a value to the <i>Shell_Path_List</i> attribute of the batch job for each entry in the <i>path_name_list</i> option-argument.
31468 31469		If the $-\mathbf{S}$ option is not presented to the <i>qsub</i> utility, the utility shall set the <i>Shell_Path_List</i> attribute of the batch job to the null string.
31470 31471		The conformance document for an implementation shall describe the mechanism used to set the default shell and determine the current value of the default shell.

qsub Utilities

31472 31473 31474 31475		An implementation shall provide a means for the installation to set the default shell to the login shell of the user under which the batch job is to execute. See Section 3.3.3 (on page 123) for a means of removing <code>keyword=value</code> (and <code>value@keyword</code>) pairs and other general rules for list-oriented batch job attributes.
31476	- u user_list	Define the user name under which the batch job is to execute.
31477 31478		The <i>qsub</i> utility shall accept a <i>user_list</i> option-argument that conforms to the following syntax:
31479		username[@host][,,username[@host],,]
31480 31481		The <i>qsub</i> utility shall accept only one user name that is missing a corresponding host name. The <i>qsub</i> utility shall accept only one user name per named host.
31482 31483		The <i>qsub</i> utility shall add a value to the <i>User_List</i> attribute of the batch job for each entry in the <i>user_list</i> option-argument.
31484 31485 31486 31487		If the $-\mathbf{u}$ option is not presented to the <i>qsub</i> utility, the utility shall set the <i>User_List</i> attribute of the batch job to the user name from which the utility is executing. See Section 3.3.3 (on page 123) for a means of removing <i>keyword=value</i> (and <i>value@keyword</i>) pairs and other general rules for list-oriented batch job attributes.
31488	-v variable_l	
31489		Add to the list of variables that are exported to the session leader of the batch job.
31490 31491		A <i>variable_list</i> is a set of strings of either the form <i><variable></variable></i> or <i><variable=value></variable=value></i> , delimited by commas.
31492 31493 31494 31495		If the $-\mathbf{v}$ option is presented to the <i>qsub</i> utility, the utility shall also add, to the environment $Variable_List$ attribute of the batch job, every variable named in the environment $variable_list$ option-argument and, optionally, values of specified variables.
31496 31497 31498 31499		If a value is not provided on the command line, the <i>qsub</i> utility shall set the value of each variable in the environment <i>Variable_List</i> attribute of the batch job to the value of the corresponding environment variable for the process in which the utility is executing; see Table 4-18 (on page 804).
31500 31501		A conforming application shall not repeat a variable in the environment <code>variable_list</code> option-argument.
31502 31503 31504 31505		The <i>qsub</i> utility shall not repeat a variable in the environment <i>Variable_List</i> attribute of the batch job. See Section 3.3.3 (on page 123) for a means of removing <i>keyword=value</i> (and <i>value@keyword</i>) pairs and other general rules for list-oriented batch job attributes.
31506 31507	- V	Specify that all of the environment variables of the process are exported to the context of the batch job.
31508 31509 31510		The <i>qsub</i> utility shall place every environment variable in the process in which the utility is executing in the list and shall set the value of each variable in the attribute to the value of that variable in the process.
31511 31512	-z	Specify that the utility does not write the batch $job_identifier$ of the created batch job to standard output.
31513 31514		If the $-\mathbf{z}$ option is presented to the <i>qsub</i> utility, the utility shall not write the batch <i>job_identifier</i> of the created batch job to standard output.

Utilities qsub

31515 If the -z option is not presented to the *qsub* utility, the utility shall write the 31516 identifier of the created batch job to standard output. 31517 **OPERANDS** The *qsub* utility shall accept a *script* operand that indicates the path to the script of the batch job. 31518 31519 If the *script* operand is not presented to the *qsub* utility, or if the operand is the single-character string ' - ', the utility shall read the script from standard input. 31520 31521 If the script represents a partial path, the qsub utility shall expand the path relative to the current directory of the process executing the utility. 31522 31523 **STDIN** The qsub utility reads the script of the batch job from standard input if the script operand is 31524 omitted or is the single character '-'. 31525 31526 INPUT FILES In addition to binding the file indicated by the *script* operand to the batch job, the *qsub* utility 31527 reads the script file and acts on directives in the script. 31528 31529 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *qsub*: 31530 LANG Provide a default value for the internationalization variables that are unset or null. 31531 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 31532 Internationalization Variables for the precedence of internationalization variables 31533 31534 used to determine the values of locale categories.) LC ALL If set to a non-empty string value, override the values of all the other 31535 31536 internationalization variables. Determine the locale for the interpretation of sequences of bytes of text data as 31537 LC_CTYPE 31538 characters (for example, single-byte as opposed to multi-byte characters in 31539 arguments). 31540 LC_MESSAGES Determine the locale that should be used to affect the format and contents of 31541 31542 diagnostic messages written to standard error. *LOGNAME* Determine the login name of the user. 31543 31544 PBS DPREFIX Determine the default prefix for directives within the script. 31545 31546 SHELL Determine the pathname of the preferred command language interpreter of the 31547 user. TZ31548 Determine the timezone used to interpret the *date-time* option-argument. If TZ is unset or null, an unspecified default timezone shall be used. 31549 31550 ASYNCHRONOUS EVENTS Once created, a batch job exists until it exits, aborts, or is deleted. 31551 31552 After a batch job is created by the qsub utility, batch servers might route, execute, modify, or 31553 delete the batch job. 31554 STDOUT The *qsub* utility writes the batch *job_identifier* assigned to the batch job to standard output, unless 31555 31556 the -z option is specified.

qsub Utilities

31557 STDERR

31567

31568

31569

31570 31571

31572

31573 31574

31575

31576

31577

31578

31579

31580

31581 31582

31583

31584 31585

31586

31587

31588

31589

31590 31591

31592

31558 The standard error shall be used only for diagnostic messages.

31559 OUTPUT FILES

31560 None.

31561 EXTENDED DESCRIPTION

31562 Script Preservation

The *qsub* utility shall make the script available to the server executing the batch job in such a way that the server executes the script as it exists at the time of submission.

The *qsub* utility can send a copy of the script to the server with the *Queue Job Request* or store a temporary copy of the script in a location specified to the server.

Option Specification

A script can contain directives to the *qsub* utility.

The *qsub* utility shall scan the lines of the script for directives, skipping blank lines, until the first line that begins with a string other than the directive string; if directives occur on subsequent lines, the utility shall ignore those directives.

Lines are separated by a <newline>. If the first line of the script begins with "#!" or a colon (':'), then it is skipped. The *qsub* utility shall process a line in the script as a directive if and only if the string of characters from the first non-white-space character on the line until the first <space> or <tab> on the line match the directive prefix. If a line in the script contains a directive and the final characters of the line are backslash (' $\$ ') and <newline>, then the next line shall be interpreted as a continuation of that directive.

The *qsub* utility shall process the options and option-arguments contained on the directive prefix line using the same syntax as if the options were input on the *qsub* utility.

The *qsub* utility shall continue to process a directive prefix line until after a <newline> is encountered. An implementation may ignore lines which, according to the syntax of the shell that will interpret the script, are comments. An implementation shall describe in the conformance document the format of any shell comments that it will recognize.

If an option is present in both a directive and the arguments to the *qsub* utility, the utility shall ignore the option and the corresponding option-argument, if any, in the directive.

If an option that is present in the directive is not present in the arguments to the *qsub* utility, the utility shall process the option and the option-argument, if any.

In order of preference, the *qsub* utility shall select the directive prefix from one of the following sources:

- If the -C option is presented to the utility, the value of the *directive_prefix* option-argument
- If the environment variable *PBS_DPREFIX* is defined, the value of that variable
- The four-character string "#PBS" encoded in the portable character set
- 31593 If the **–C** option is present in the script file it shall be ignored.

31594 EXIT STATUS

31595 The following exit values shall be returned:

31596 0 Successful completion.

Utilities qsub

31597 >0 An error occurred.

31598 CONSEQUENCES OF ERRORS

Default.

31600 APPLICATION USAGE

31601 None.

31602 EXAMPLES

31603 None.

31604 RATIONALE

The *qsub* utility allows users to create a batch job that will process the script specified as the operand of the utility.

The options of the *qsub* utility allow users to control many aspects of the queuing and execution of a batch job.

The –a option allows users to designate the time after which the batch job will become eligible to run. By specifying an execution time, users can take advantage of resources at off-peak hours, synchronize jobs with chronologically predictable events, and perhaps take advantage of off-peak pricing of computing time. For these reasons and others, a timing option is existing practice on the part of almost every batch system, including NQS.

The **–A** option allows users to specify the account that will be charged for the batch job. Support for account is not mandatory for conforming batch servers.

The -C option allows users to prescribe the prefix for directives within the script file. The default prefix "#PBS" may be inappropriate if the script will be interpreted with an alternate shell, as specified by the -S option.

The –c option allows users to establish the checkpointing interval for their jobs. A checkpointing system, which is not defined by this volume of IEEE Std 1003.1-2001, allows recovery of a batch job at the most recent checkpoint in the event of a crash. Checkpointing is typically used for jobs that consume expensive computing time or must meet a critical schedule. Users should be allowed to make the tradeoff between the overhead of checkpointing and the risk to the timely completion of the batch job; therefore, this volume of IEEE Std 1003.1-2001 provides the checkpointing interval option. Support for checkpointing is optional for batch servers.

The —e option allows users to redirect the standard error streams of their jobs to a non-default path. For example, if the submitted script generally produces a great deal of useless error output, a user might redirect the standard error output to the null device. Or, if the file system holding the default location (the home directory of the user) has too little free space, the user might redirect the standard error stream to a file in another file system.

The **–h** option allows users to create a batch job that is held until explicitly released. The ability to create a held job is useful when some external event must complete before the batch job can execute. For example, the user might submit a held job and release it when the system load has dropped.

The -j option allows users to merge the standard error of a batch job into its standard output stream, which has the advantage of showing the sequential relationship between output and error messages.

The -m option allows users to designate those points in the execution of a batch job at which mail will be sent to the submitting user, or to the account(s) indicated by the -M option. By requesting mail notification at points of interest in the life of a job, the submitting user, or other designated users, can track the progress of a batch job.

qsub Utilities

The –N option allows users to associate a name with the batch job. The job name in no way affects the processing of the batch job, but rather serves as a mnemonic handle for users. For example, the batch job name can help the user distinguish between multiple jobs listed by the *qstat* utility.

The $-\mathbf{o}$ option allows users to redirect the standard output stream. A user might, for example, wish to redirect to the null device the standard output stream of a job that produces copious yet superfluous output.

The **–P** option allows users to designate the relative priority of a batch job for selection from a queue.

The $-\mathbf{q}$ option allows users to specify an initial queue for the batch job. If the user specifies a routing queue, the batch server routes the batch job to another queue for execution or further routing. If the user specifies a non-routing queue, the batch server of the queue eventually executes the batch job.

The **-r** option allows users to control whether the submitted job will be rerun if the controlling batch node fails during execution of the batch job. The **-r** option likewise allows users to indicate whether or not the batch job is eligible to be rerun by the *qrerun* utility. Some jobs cannot be correctly rerun because of changes they make in the state of databases or other aspects of their environment. This volume of IEEE Std 1003.1-2001 specifies that the default, if the **-r** option is not presented to the utility, will be that the batch job cannot be rerun, since the result of rerunning a non-rerunnable job might be catastrophic.

The -S option allows users to specify the program (usually a shell) that will be invoked to process the script of the batch job. This option has been modified to allow a list of shell names and locations associated with different hosts.

The $-\mathbf{u}$ option is useful when the submitting user is authorized to use more than one account on a given host, in which case the $-\mathbf{u}$ option allows the user to select from among those accounts. The option-argument is a list of user-host pairs, so that the submitting user can provide different user identifiers for different nodes in the event the batch job is routed. The $-\mathbf{u}$ option provides a lot of flexibility to accommodate sites with complex account structures. Users that have the same user identifier on all the hosts they are authorized to use will not need to use the $-\mathbf{u}$ option.

The –V option allows users to export all their current environment variables, as of the time the batch job is submitted, to the context of the processes of the batch job.

The $-\mathbf{v}$ option allows users to export specific environment variables from their current process to the processes of the batch job.

The -z option allows users to suppress the writing of the batch job identifier to standard output. The -z option is an existing NQS practice that has been standardized.

Historically, the *qsub* utility has served the batch job-submission function in the NQS system, the existing practice on which it is based. Some changes and additions have been made to the *qsub* utility in this volume of IEEE Std 1003.1-2001, *vis-a-vis* NQS, as a result of the growing pool of experience with distributed batch systems.

The set of features of the *qsub* utility as defined in this volume of IEEE Std 1003.1-2001 appears to incorporate all the common existing practice on potentially conforming platforms.

31683 FUTURE DIRECTIONS

31684 None.

 Utilities qsub

31685 SEE ALSO

31686 Chapter 3 (on page 101), qrerun, qstat, touch

31687 CHANGE HISTORY

31688 Derived from IEEE Std 1003.2d-1994.

31689 **Issue 6**

31690 The -1 option has been removed as there is no portable description of the resources that are

allowed or required by the batch job.

read Utilities

```
31692 NAME
31693
              read — read a line from standard input
31694 SYNOPSIS
31695
              read [-r] var...
31696 DESCRIPTION
              The read utility shall read a single line from standard input.
31697
31698
              By default, unless the -\mathbf{r} option is specified, backslash ('\') shall act as an escape character, as
              described in Section 2.2.1 (on page 30). If standard input is a terminal device and the invoking
31699
              shell is interactive, read shall prompt for a continuation line when:
31700

    The shell reads an input line ending with a backslash, unless the -r option is specified.

31701
31702

    A here-document is not terminated after a <newline> is entered.

              The line shall be split into fields as in the shell (see Section 2.6.5 (on page 42)); the first field shall
31703
              be assigned to the first variable var, the second field to the second variable var, and so on. If
31704
              there are fewer var operands specified than there are fields, the leftover fields and their
31705
              intervening separators shall be assigned to the last var. If there are fewer fields than vars, the
31706
              remaining vars shall be set to empty strings.
31707
              The setting of variables specified by the var operands shall affect the current shell execution
31708
              environment; see Section 2.12 (on page 61). If it is called in a subshell or separate utility
31709
31710
              execution environment, such as one of the following:
31711
               (read foo)
31712
              nohup read ...
31713
              find . -exec read ... \;
              it shall not affect the shell variables in the caller's environment.
31714
31715 OPTIONS
31716
              The read utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
31717
              12.2, Utility Syntax Guidelines.
31718
              The following option is supported:
                            Do not treat a backslash character in any special way. Consider each backslash to
31719
              -\mathbf{r}
31720
                            be part of the input line.
31721 OPERANDS
31722
              The following operand shall be supported:
                            The name of an existing or nonexisting shell variable.
31793
               var
31724 STDIN
              The standard input shall be a text file.
31725
31726 INPUT FILES
              None
31728 ENVIRONMENT VARIABLES
              The following environment variables shall affect the execution of read:
31729
              IFS
                            Determine the internal field separators used to delimit fields; see Section 2.5.3 (on
31730
                            page 34).
31731
```

31732

31733

31734

LANG

Provide a default value for the internationalization variables that are unset or null.

(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,

Internationalization Variables for the precedence of internationalization variables

Utilities read

31735		used to determine the values of locale categories.)
31736 31737	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
31738 31739 31740	LC_CTYPE	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).
31741	LC_MESSA	GES
31742		Determine the locale that should be used to affect the format and contents of
31743		diagnostic messages written to standard error.
31744 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .
31745	PS2	Provide the prompt string that an interactive shell shall write to standard error
31746		when a line ending with a backslash is read and the -r option was not specified, or
31747		if a here-document is not terminated after a <newline> is entered.</newline>
31748 ASYN C	CHRONOUS	EVENTS
31749	Default.	
31750 STDOU	J T	
31751	Not used.	
31752 STDER	R	
31753	The standar	d error shall be used for diagnostic messages and prompts for continued input.
31754 OUTPU	T FILES	
31755	None.	
31756 EXTEN	DED DESCR	IPTION
31757	None.	
31758 EXIT S	TATUS	
31759	The following	ng exit values shall be returned:
31760	0 Success	ful completion.
31761	>0 End-of-	file was detected or an error occurred.
31762 CONSE 31763	EQUENCES C Default.	OF ERRORS
31764 APPLIC	CATION USA	GE
31765	The -r optio	on is included to enable <i>read</i> to subsume the purpose of the <i>line</i> utility, which is not
31766	included in	IEEE Std 1003.1-2001.
31767 31768	The results a when - r is n	are undefined if an end-of-file is detected following a backslash at the end of a line of specified.
31769 EXAMI	PLES	
31770	The following	ng command:
31771	while rea	d -r xx yy
31772	do 	
31773	-	f "%s %s\n" "\$yy" "\$xx"
31774	done < in	
31775	prints a file	with the first field of each line moved to the end of the line.

read Utilities

31776 RATIONALE 31777 The read utility historically has been a shell built-in. It was separated off into its own utility to take advantage of the richer description of functionality introduced by this volume of 31778 IEEE Std 1003.1-2001. 31779 31780 Since read affects the current shell execution environment, it is generally provided as a shell regular built-in. If it is called in a subshell or separate utility execution environment, such as one 31781 31782 of the following: 31783 (read foo) 31784 nohup read ... 31785 find . -exec read ... \; it does not affect the shell variables in the environment of the caller. 31786 31787 FUTURE DIRECTIONS 31788 None. 31789 SEE ALSO 31790 Chapter 2 (on page 29) 31791 CHANGE HISTORY First released in Issue 2. 31792

Utilities renice

31793 NAME 31794	ranica — sat :	nice values of running processes
		ince values of running processes
31795 SYNOP 31796 UP		increment [-g -p -u] ID
31797		
31798 DESCR	IPTION	
31799		itility shall request that the nice values (see the Base Definitions volume of
31800 31801		3.1-2001, Section 3.239, Nice Value) of one or more running processes be changed. ne applicable processes are specified by their process IDs. When a process group is
31802		e –g), the request shall apply to all processes in the process group.
31803	The nice val	lue shall be bounded in an implementation-defined manner. If the requested
31804		ould raise or lower the nice value of the executed utility beyond implementation-
31805		s, then the limit whose value was exceeded shall be used.
31806 31807		is <i>renice</i> d, the request applies to all processes whose saved set-user-ID matches the sponding to the user.
31808		f which options are supplied or any other factor, <i>renice</i> shall not alter the nice values
31809 31810		ss unless the user requesting such a change has appropriate privileges to do so for process. If the user lacks appropriate privileges to perform the requested action, the
31811	•	eturn an error status.
31812	The saved se	t-user-ID of the user's process shall be checked instead of its effective user ID when
31813		its to determine the user ID of the process in order to determine whether the user
31814		ate privileges.
31815 OPTIO N 31816		ility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
31817		Syntax Guidelines.
31818	The following	g options shall be supported:
31819	- g	Interpret all operands as unsigned decimal integer process group IDs.
31820 31821	-n increment	Specify how the nice value of the specified process or processes is to be adjusted. The <i>increment</i> option-argument is a positive or negative decimal integer that shall
31822		be used to modify the nice value of the specified process or processes.
31823		Positive <i>increment</i> values shall cause a lower nice value. Negative <i>increment</i> values
31824		may require appropriate privileges and shall cause a higher nice value.
31825 31826	-p	Interpret all operands as unsigned decimal integer process IDs. The $-\mathbf{p}$ option is the default if no options are specified.
31827	-u	Interpret all operands as users. If a user exists with a user name equal to the
31828		operand, then the user ID of that user is used in further processing. Otherwise, if
31829 31830		the operand represents an unsigned decimal integer, it shall be used as the numeric user ID of the user.
31831 OPERA	NDS	
31832		g operands shall be supported:

A process ID, process group ID, or user name/user ID, depending on the option

selected.

ID

31833

31834

renice Utilities

31835 STDIN		
31836	Not used.	
31837 INPUT F 31838	TILES None.	
31839 ENVIRO 31840		RIABLES g environment variables shall affect the execution of <i>renice</i> :
31841 31842 31843 31844	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
31845 31846	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
31847 31848 31849	LC_CTYPE	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).
31850 31851 31852	LC_MESSAC	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
31853 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
31854 ASYNCH 31855	HRONOUS I Default.	EVENTS
31856 STDOU 7 31857	Γ Not used.	
31858 STDERR	2	
31859	The standard	l error shall be used only for diagnostic messages.
31860 OUTPU T 31861	Γ FILES None.	
31862 EXTEND 31863	DED DESCRI None.	PTION
31864 EXIT ST A 31865		g exit values shall be returned:
31866	0 Successf	ful completion.
31867	>0 An error	occurred.
31868 CONSEC	QUENCES O	F ERRORS

31869

Default.

Utilities renice

31870 APPLICATION USAGE

None. None.

31872 EXAMPLES

1. Adjust the nice value so that process IDs 987 and 32 would have a lower nice value:

```
31874 renice -n 5 -p 987 32
```

2. Adjust the nice value so that group IDs 324 and 76 would have a higher nice value, if the user has the appropriate privileges to do so:

```
renice -n -4 -g 324 76
```

3. Adjust the nice value so that numeric user ID 8 and user **sas** would have a lower nice value:

```
31880 renice -n 4 -u 8 sas
```

Useful nice value increments on historical systems include 19 or 20 (the affected processes run only when nothing else in the system attempts to run) and any negative number (to make processes run faster).

31884 RATIONALE

The *gid*, *pid*, and *user* specifications do not fit either the definition of operand or optionargument. However, for clarity, they have been included in the OPTIONS section, rather than the OPERANDS section.

The definition of nice value is not intended to suggest that all processes in a system have priorities that are comparable. Scheduling policy extensions such as the realtime priorities in the System Interfaces volume of IEEE Std 1003.1-2001 make the notion of a single underlying priority for all scheduling policies problematic. Some implementations may implement the *nice*-related features to affect all processes on the system, others to affect just the general time-sharing activities implied by this volume of IEEE Std 1003.1-2001, and others may have no effect at all. Because of the use of "implementation-defined" in *nice* and *renice*, a wide range of implementation strategies are possible.

Originally, this utility was written in the historical manner, using the term "nice value". This was always a point of concern with users because it was never intuitively obvious what this meant. With a newer version of *renice*, which used the term "system scheduling priority", it was hoped that novice users could better understand what this utility was meant to do. Also, it would be easier to document what the utility was meant to do. Unfortunately, the addition of the POSIX realtime scheduling capabilities introduced the concepts of process and thread scheduling priorities that were totally unaffected by the *nice/renice* utilities or the *nice()/setpriority()* functions. Continuing to use the term "system scheduling priority" would have incorrectly suggested that these utilities and functions were indeed affecting these realtime priorities. It was decided to revert to the historical term "nice value" to reference this unrelated process attribute.

Although this utility has use by system administrators (and in fact appears in the system administration portion of the BSD documentation), the standard developers considered that it was very useful for individual end users to control their own processes.

31910 FUTURE DIRECTIONS

31911 None.

renice Utilities

31912 **SEE ALSO** 31913 *nice*

31914 CHANGE HISTORY

First released in Issue 4.

31916 **Issue 5**

In the SYNOPSIS, an ellipsis is added to the $-\mathbf{u}$ option in all three obsolescent forms.

31918 **Issue 6**

This utility is marked as part of the User Portability Utilities option.

The APPLICATION USAGE section is added.

The obsolescent forms of the SYNOPSIS are removed.

Text previously conditional on POSIX_SAVED_IDS is mandatory in this issue. This is a FIPS

31923 requirement.

Utilities rm

```
31924 NAME
31925 rm — remove directory entries
31926 SYNOPSIS
31927 rm [-fiRr] file...
```

31928 DESCRIPTION

The *rm* utility shall remove the directory entry specified by each *file* argument.

If either of the files dot or dot-dot are specified as the basename portion of an operand (that is, the final pathname component), *rm* shall write a diagnostic message to standard error and do nothing more with such operands.

For each *file* the following steps shall be taken:

- 1. If the *file* does not exist:
 - a. If the –f option is not specified, *rm* shall write a diagnostic message to standard error.
 - b. Go on to any remaining *files*.
- 2. If *file* is of type directory, the following steps shall be taken:
 - a. If neither the $-\mathbf{R}$ option nor the $-\mathbf{r}$ option is specified, rm shall write a diagnostic message to standard error, do nothing more with *file*, and go on to any remaining files.
 - b. If the **–f** option is not specified, and either the permissions of *file* do not permit writing and the standard input is a terminal or the **–i** option is specified, *rm* shall write a prompt to standard error and read a line from the standard input. If the response is not affirmative, *rm* shall do nothing more with the current file and go on to any remaining files.
 - c. For each entry contained in *file*, other than dot or dot-dot, the four steps listed here (1 to 4) shall be taken with the entry as if it were a *file* operand. The *rm* utility shall not traverse directories by following symbolic links into other parts of the hierarchy, but shall remove the links themselves.
 - d. If the **-i** option is specified, *rm* shall write a prompt to standard error and read a line from the standard input. If the response is not affirmative, *rm* shall do nothing more with the current file, and go on to any remaining files.
- 3. If *file* is not of type directory, the **-f** option is not specified, and either the permissions of *file* do not permit writing and the standard input is a terminal or the **-i** option is specified, *rm* shall write a prompt to the standard error and read a line from the standard input. If the response is not affirmative, *rm* shall do nothing more with the current file and go on to any remaining files.
- 4. If the current file is a directory, rm shall perform actions equivalent to the rmdir() function defined in the System Interfaces volume of IEEE Std 1003.1-2001 called with a pathname of the current file used as the path argument. If the current file is not a directory, rm shall perform actions equivalent to the unlink() function defined in the System Interfaces volume of IEEE Std 1003.1-2001 called with a pathname of the current file used as the path argument.

If this fails for any reason, *rm* shall write a diagnostic message to standard error, do nothing more with the current file, and go on to any remaining files.

The *rm* utility shall be able to descend to arbitrary depths in a file hierarchy, and shall not fail due to path length limitations (unless an operand specified by the user exceeds system

rm Utilities

31968	limitations).	
31969 OPTIO	NS	
31970 31971	The <i>rm</i> utilit	y shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, ax Guidelines.
31972	The followin	ng options shall be supported:
31973 31974 31975	−f	Do not prompt for confirmation. Do not write diagnostic messages or modify the exit status in the case of nonexistent operands. Any previous occurrences of the $-\mathbf{i}$ option shall be ignored.
31976 31977	- i	Prompt for confirmation as described previously. Any previous occurrences of the —f option shall be ignored.
31978	$-\mathbf{R}$	Remove file hierarchies. See the DESCRIPTION.
31979	-r	Equivalent to – R .
31980 OPERA	NDS	
31981		ng operand shall be supported:
31982	file	A pathname of a directory entry to be removed.
31983 STDIN		
31984 31985		d input shall be used to read an input line in response to each prompt specified in section. Otherwise, the standard input shall not be used.
31986 INPUT 31987	FILES None.	
	ONMENT VA	
31988 ENVIR 31989 31990 31991 31992 31993		ARIABLES ag environment variables shall affect the execution of <i>rm</i> : Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
31989 31990 31991 31992	The following	ng environment variables shall affect the execution of <i>rm</i> : Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables
31989 31990 31991 31992 31993 31994 31995	The followin	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables.
31989 31990 31991 31992 31993 31994	The followin	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables.
31989 31990 31991 31992 31993 31994 31995 31996 31997 31998	The followin	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. The Determine the locale for the behavior of ranges, equivalence classes, and multicharacter collating elements used in the extended regular expression defined for
31989 31990 31991 31992 31993 31994 31995 31996 31997 31998 31999 32000 32001 32002 32003	The following LANG LC_ALL LC_COLLAT	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) If set to a non-empty string value, override the values of all the other internationalization variables. E Determine the locale for the behavior of ranges, equivalence classes, and multicharacter collating elements used in the extended regular expression defined for the yesexpr locale keyword in the LC_MESSAGES category. 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 the behavior of character classes within regular expressions used in the extended regular expression defined for the yesexpr locale keyword in the LC_MESSAGES category.

Utilities rm

32010 ASYNCHRONOUS EVENTS 32011 Default.

32012 STDOUT

32013 Not used.

32014 STDERR

Prompts shall be written to standard error under the conditions specified in the DESCRIPTION and OPTIONS sections. The prompts shall contain the *file* pathname, but their format is otherwise unspecified. The standard error also shall be used for diagnostic messages.

32018 OUTPUT FILES

32019 None.

32020 EXTENDED DESCRIPTION

32021 None.

32022 EXIT STATUS

32023 The following exit values shall be returned:

32024 0 All of the named directory entries for which *rm* performed actions equivalent to the *rmdir*() or *unlink*() functions were removed.

32026 >0 An error occurred.

32027 CONSEQUENCES OF ERRORS

32028 Default.

32029 APPLICATION USAGE

The *rm* utility is forbidden to remove the names dot and dot-dot in order to avoid the consequences of inadvertently doing something like:

32032 rm -r .*

Some implementations do not permit the removal of the last link to an executable binary file that is being executed; see the [EBUSY] error in the *unlink()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001. Thus, the *rm* utility can fail to remove such files.

The $-\mathbf{i}$ option causes rm to prompt and read the standard input even if the standard input is not a terminal, but in the absence of $-\mathbf{i}$ the mode prompting is not done when the standard input is not a terminal.

32039 EXAMPLES

32036

32037 32038

32040

32043

1. The following command:

32041 rm a.out core

32042 removes the directory entries: **a.out** and **core**.

2. The following command:

32044 rm -Rf junk

32045 removes the directory **junk** and all its contents, without prompting.

32046 RATIONALE

For absolute clarity, paragraphs (2b) and (3) in the DESCRIPTION of *rm* describing the behavior when prompting for confirmation, should be interpreted in the following manner:

```
32049 if ((NOT f_option) AND 
32050 ((not_writable AND input_is_terminal) OR i_option))
```

rm Utilities

The exact format of the interactive prompts is unspecified. Only the general nature of the contents of prompts are specified because implementations may desire more descriptive prompts than those used on historical implementations. Therefore, an application not using the –f option, or using the –i option, relies on the system to provide the most suitable dialog directly with the user, based on the behavior specified.

The -r option is historical practice on all known systems. The synonym -R option is provided for consistency with the other utilities in this volume of IEEE Std 1003.1-2001 that provide options requesting recursive descent through the file hierarchy.

The behavior of the **–f** option in historical versions of *rm* is inconsistent. In general, along with "forcing" the unlink without prompting for permission, it always causes diagnostic messages to be suppressed and the exit status to be unmodified for nonexistent operands and files that cannot be unlinked. In some versions, however, the **–f** option suppresses usage messages and system errors as well. Suppressing such messages is not a service to either shell scripts or users.

It is less clear that error messages regarding files that cannot be unlinked (removed) should be suppressed. Although this is historical practice, this volume of IEEE Std 1003.1-2001 does not permit the –f option to suppress such messages.

When given the **-r** and **-i** options, historical versions of *rm* prompt the user twice for each directory, once before removing its contents and once before actually attempting to delete the directory entry that names it. This allows the user to "prune" the file hierarchy walk. Historical versions of *rm* were inconsistent in that some did not do the former prompt for directories named on the command line and others had obscure prompting behavior when the **-i** option was specified and the permissions of the file did not permit writing. The POSIX Shell and Utilities *rm* differs little from historic practice, but does require that prompts be consistent. Historical versions of *rm* were also inconsistent in that prompts were done to both standard output and standard error. This volume of IEEE Std 1003.1-2001 requires that prompts be done to standard error, for consistency with *cp* and *mv*, and to allow historical extensions to *rm* that provide an option to list deleted files on standard output.

The *rm* utility is required to descend to arbitrary depths so that any file hierarchy may be deleted. This means, for example, that the *rm* utility cannot run out of file descriptors during its descent (that is, if the number of file descriptors is limited, *rm* cannot be implemented in the historical fashion where one file descriptor is used per directory level). Also, *rm* is not permitted to fail because of path length restrictions, unless an operand specified by the user is longer than {PATH_MAX}.

The rm utility removes symbolic links themselves, not the files they refer to, as a consequence of the dependence on the unlink() functionality, per the DESCRIPTION. When removing hierarchies with $-\mathbf{r}$ or $-\mathbf{R}$, the prohibition on following symbolic links has to be made explicit.

32087 FUTURE DIRECTIONS

32088 None.

32089 SEE ALSO

32090 rmdir, the System Interfaces volume of IEEE Std 1003.1-2001, remove(), rmdir(), unlink()

32091 CHANGE HISTORY

32092 First released in Issue 2.

32093 Issue 5

32094 The FUTURE DIRECTIONS section is added.

Utilities rm

32095 Issue 6

Text is added to clarify actions relating to symbolic links as specified in the IEEE P1003.2b draft standard.

rmdel Utilities

32098 NAME

32099 rmdel — remove a delta from an SCCS file (**DEVELOPMENT**)

32100 SYNOPSIS

32101 XSI rmdel -r SID file...

32102

32103 DESCRIPTION

The *rmdel* utility shall remove the delta specified by the SID from each named SCCS file. The delta to be removed shall be the most recent delta in its branch in the delta chain of each named SCCS file. In addition, the application shall ensure that the SID specified is not that of a version being edited for the purpose of making a delta; that is, if a *p-file* (see *get*) exists for the named SCCS file, the SID specified shall not appear in any entry of the *p-file*.

32109 Removal of a delta shall be restricted to:

- 32110 1. The user who made the delta
- 32111 2. The owner of the SCCS file
- 32112 3. The owner of the directory containing the SCCS file

32113 OPTIONS

The *rmdel* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

in the state of th

32116 The following option shall be supported:

32117 — r SID Specify the SCCS identification string (SID) of the delta to be deleted.

32118 **OPERANDS**

32119 The following operand shall be supported:

32120 *file* A pathname of an existing SCCS file or a directory. If *file* is a directory, the *rmdel*32121 utility shall behave as though each file in the directory were specified as a named
32122 file, except that non-SCCS files (last component of the pathname does not begin
32123 with **s.**) and unreadable files shall be silently ignored.

If exactly one *file* operand appears, and it is '-', the standard input shall be read; each line of the standard input is taken to be the name of an SCCS file to be processed. Non-SCCS files and unreadable files shall be silently ignored.

32127 **STDIN**

32128 32129

32133

The standard input shall be a text file used only when the *file* operand is specified as '-'. Each line of the text file shall be interpreted as an SCCS pathname.

32130 INPUT FILES

32131 The SCCS files shall be files of unspecified format.

32132 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *rmdel*:

32134 LANG Provide a default value for the internationalization variables that are unset or null.
32135 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
32136 Internationalization Variables for the precedence of internationalization variables
32137 used to determine the values of locale categories.)

32138 *LC_ALL* If set to a non-empty string value, override the values of all the other internationalization variables.

Utilities rmdel

32140 32141 32142	LC_CTYPE	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).
32143	LC_MESSA	GES
32144		Determine the locale that should be used to affect the format and contents of
32145		diagnostic messages written to standard error.
32146	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
32147 ASYN	CHRONOUS	EVENTS
32148	Default.	
32149 STDO	UT	
32150	Not used.	
32151 STDE I	RR	
32152	The standar	d error shall be used only for diagnostic messages.
32153 OUTP	UT FILES	
32154	The SCCS fi	les shall be files of unspecified format. During processing of a <i>file</i> , a temporary <i>x-file</i> ,
32155		I in admin, may be created and deleted; a locking z-file, as described in get, may be
32156	created and	deleted.
32157 EXTEN	NDED DESCR	IPTION
32158	None.	
32159 EXIT S	TATUS	
32160	The following	ng exit values shall be returned:
32161	0 Success	ful completion.
32162	>0 An erro	or occurred.
32163 CONS	EQUENCES C	OF ERRORS
32164	Default.	
32165 APPLI	CATION USA	.GE
32166	None.	
32167 EXAM	PLES	
32168	None.	
32169 RATIC	NALE	
32170	None.	
39171 FI TI TI	RE DIRECTIO	NS
32172	None.	
32173 SEE A l		
32173 SEE A 1	admin, delta,	get, prs
	GE HISTORY	•
32175 CHAIN	First release	
		
32177 Issue 6)	

The normative text is reworded to avoid use of the term "must" for application requirements.

32178

rmdir Utilities

	NAME			
32180		- remove directories		
32181 32182	SYNOPSIS rmdir	[-p] dir		
	DESCRIPTION			
32184	The <i>rma</i>	<i>lir</i> utility shall remove the directory entry specified by each <i>dir</i> operand.		
32185 32186		n <i>dir</i> operand, the <i>rmdir</i> utility shall perform actions equivalent to the <i>rmdir</i> () function with the <i>dir</i> operand as its only argument.		
32187 32188 32189 32190	director subdire	Directories shall be processed in the order specified. If a directory and a subdirectory of that directory are specified in a single invocation of the <i>rmdir</i> utility, the application shall specify the subdirectory before the parent directory so that the parent directory will be empty when the <i>rmdir</i> utility tries to remove it.		
32191	OPTIONS			
32192 32193		<i>lir</i> utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section lity Syntax Guidelines.		
32194	The foll	owing option shall be supported:		
32195	-p	Remove all directories in a pathname. For each dir operand:		
32196		1. The directory entry it names shall be removed.		
32197 32198		2. If the <i>dir</i> operand includes more than one pathname component, effects equivalent to the following command shall occur:		
32199		rmdir -p \$(dirname dir)		
32200	OPERANDS			
32201	The follo	owing operand shall be supported:		
32202	dir	A pathname of an empty directory to be removed.		
	STDIN			
32204	Not use	d.		
32205 32206	INPUT FILES None.			
32207 32208	ENVIRONMEN' The follo	F VARIABLES by owing environment variables shall affect the execution of <i>rmdir</i> :		
32209	LANG	Provide a default value for the internationalization variables that are unset or null.		
32210		(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,		
32211 32212		Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)		
32213 32214	LC_ALL	<u> </u>		
32215 32216 32217	LC_CTY	The 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).		
32218	LC_ME			
32219 32220	LC_iviEx	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.		

Utilities rmdir

32221 XSI **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 32222 ASYNCHRONOUS EVENTS 32223 Default. 32224 STDOUT 32225 Not used. 32226 STDERR 32227 The standard error shall be used only for diagnostic messages. 32228 OUTPUT FILES 32229 None. 32230 EXTENDED DESCRIPTION 32231 None. 32232 EXIT STATUS 32233 The following exit values shall be returned: Each directory entry specified by a *dir* operand was removed successfully. 32234 >0 An error occurred. 32235 32236 CONSEQUENCES OF ERRORS 32237 Default. 32238 APPLICATION USAGE 32239 The definition of an empty directory is one that contains, at most, directory entries for dot and 32240 dot-dot. 32241 EXAMPLES 32242 If a directory **a** in the current directory is empty except it contains a directory **b** and **a/b** is empty 32243 except it contains a directory **c**: 32244 rmdir -p a/b/c removes all three directories. 32245 32246 RATIONALE On historical System V systems, the -p option also caused a message to be written to the 32247 32248 standard output. The message indicated whether the whole path was removed or whether part 32249 of the path remained for some reason. The STDERR section requires this diagnostic when the entire path specified by a dir operand is not removed, but does not allow the status message 32250 32251 reporting success to be written as a diagnostic. The rmdir utility on System V also included a -s option that suppressed the informational 32252 32253 message output by the -p option. This option has been omitted because the informational 32254 message is not specified by this volume of IEEE Std 1003.1-2001. 32255 FUTURE DIRECTIONS None. 32256 32257 SEE ALSO 32258 rm, the System Interfaces volume of IEEE Std 1003.1-2001, remove(), rmdir(), unlink() 32259 CHANGE HISTORY First released in Issue 2. 32260

rmdir Utilities

32261 **Issue 6**

32262 The norm

The normative text is reworded to avoid use of the term "must" for application requirements.

Utilities sact

32263 **NAME** 32264 sact — print current SCCS file-editing activity (**DEVELOPMENT**) 32265 SYNOPSIS sact file ... 32266 XSI 32267 32268 DESCRIPTION The sact utility shall inform the user of any impending deltas to a named SCCS file by writing a 32269 list to standard output. This situation occurs when get -e has been executed previously without 32270 a subsequent execution of delta, unget, or sccs unedit. 32271 32272 OPTIONS None. 32273 32274 OPERANDS 32275 The following operand shall be supported: file A pathname of an existing SCCS file or a directory. If file is a directory, the sact 32276 utility shall behave as though each file in the directory were specified as a named 32277 32278 file, except that non-SCCS files (last component of the pathname does not begin 32279 with **s.**) and unreadable files shall be silently ignored. If exactly one *file* operand appears, and it is '-', the standard input shall be read; 32280 each line of the standard input shall be taken to be the name of an SCCS file to be 32281 processed. Non-SCCS files and unreadable files shall be silently ignored. 32282 32283 STDIN The standard input shall be a text file used only when the *file* operand is specified as '-'. Each 32284 line of the text file shall be interpreted as an SCCS pathname. 32285 32286 INPUT FILES 32287 Any SCCS files interrogated are files of an unspecified format. 32288 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *sact*: 32289 LANG Provide a default value for the internationalization variables that are unset or null. 32290 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 32291 32292 Internationalization Variables for the precedence of internationalization variables 32293 used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other 32294 32295 internationalization variables. LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 32296 characters (for example, single-byte as opposed to multi-byte characters in 32297

32299 LC_MESSAGES

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

32302 NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

32303 ASYNCHRONOUS EVENTS

32304 Default.

32298

arguments and input files).

sact Utilities

32305 STDO I	U T		
32306	The output for each named file shall consist of a line in the following format:		
32307	"%s Δ %s Δ %s	s Δ %s Δ %s n ", < SID >, < n ew SID >, < l ogin>, < d ate>, < t ime>	
32308 32309	<sid></sid>	Specifies the SID of a delta that currently exists in the SCCS file to which changes are made to make the new delta.	
32310	<new sid=""></new>	Specifies the SID for the new delta to be created.	
32311 32312	<login></login>	Contains the login name of the user who makes the delta (that is, who executed a <i>get</i> for editing).	
32313 32314	<date></date>	Contains the date that get – e was executed, in the format used by the prs : D : data keyword.	
32315 32316	<time></time>	Contains the time that get – e was executed, in the format used by the prs : T : data keyword.	
32317 32318		ore than one named file or if a directory or standard input is named, each pathname tten before each of the preceding lines:	
32319	"\n%s:\n"	, <pathname></pathname>	
32320 STDER	RR		
32321 32322			
32323 OUTPUT FILES			
32324	None.		
32325 EXTENDED DESCRIPTION 32326 None.			
32327 EXIT S			
32328	The following exit values shall be returned:		
32329	0 Success	sful completion.	
32330	>0 An erro	or occurred.	
32331 CONSEQUENCES OF ERRORS 32332 Default.			
32333 APPLI	CATION USA	AGE	
32334	None.		
32335 EXAM 32336	PLES None.		
32337 RATIO 32338	NALE None.		
32339 FUTUF 32340	RE DIRECTIO None.	ONS	
32341 SEE AI			
32342	delta, get, sc	cs, unget	

Utilities sact

32343 **CHANGE HISTORY**

First released in Issue 2.

SCCS Utilities

32345 NAME				
32346	sccs — front	t end for the SCCS subsystem (DEVELOPMENT)		
32347 SYNOP				
32348 XSI 32349	sccs [-r]	[-d path] [-p path] command [options] [operands]		
32350 DESCR	IPTION			
32351 32352		lity is a front end to the SCCS programs. It also includes the capability to run set- nother user to provide additional protection.		
32353 32354		lity shall invoke the specified <i>command</i> with the specified <i>options</i> and <i>operands</i> . By n of the <i>operands</i> shall be modified by prefixing it with the string "SCCS/s.".		
32355 32356 32357	(admin, delta	The <i>command</i> can be the name of one of the SCCS utilities in this volume of IEEE Std 1003.1-2001 (<i>admin, delta, get, prs, rmdel, sact, unget, val,</i> or <i>what</i>) or one of the pseudo-utilities listed in the EXTENDED DESCRIPTION section.		
32358 OPTIO				
32359 32360 32361	The <i>sccs</i> utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines, except that <i>options</i> operands are actually options to be passed to the utility named by <i>command</i> . When the portion of the command:			
32362	command [options] [operands]		
32363 32364 32365	is considered, all of the pseudo-utilities used as <i>command</i> shall support the Utility Syntax Guidelines. Any of the other SCCS utilities that can be invoked in this manner support the Guidelines to the extent indicated by their individual OPTIONS sections.			
32366	The following options shall be supported preceding the <i>command</i> operand:			
32367 32368 32369	-d path	A pathname of a directory to be used as a root directory for the SCCS files. The default shall be the current directory. The $-\mathbf{d}$ option shall take precedence over the <i>PROJECTDIR</i> variable. See $-\mathbf{p}$.		
32370 32371	-p path	A pathname of a directory in which the SCCS files are located. The default shall be the SCCS directory.		
32372 32373 32374		The $-\mathbf{p}$ option differs from the $-\mathbf{d}$ option in that the $-\mathbf{d}$ option-argument shall be prefixed to the entire pathname and the $-\mathbf{p}$ option-argument shall be inserted before the final component of the pathname. For example:		
32375		sccs -d /x -p y get a/b		
32376		converts to:		
32377		get /x/a/y/s.b		
32378		This allows the creation of aliases such as:		
32379		alias syssccs="sccs -d /usr/src"		
32380		which is used as:		
32381		syssccs get cmd/who.c		
32382 32383 32384 32385	- r	Invoke <i>command</i> with the real user ID of the process, not any effective user ID that the <i>sccs</i> utility is set to. Certain commands (<i>admin</i> , check , clean , diffs , info , <i>rmdel</i> , and tell) cannot be run set-user-ID by all users, since this would allow anyone to change the authorizations. These commands are always run as the real user.		

Utilities SCCS

anna ODEDA	NDC	
32386 OPERA 32387		ng operands shall be supported:
32388 32389	command	An SCCS utility name or the name of one of the pseudo-utilities listed in the EXTENDED DESCRIPTION section.
32390	options	An option or option-argument to be passed to command.
32391	operands	An operand to be passed to command.
32392 STDIN 32393	See the utilit	ty description for the specified <i>command</i> .
32394 INPUT 32395		ty description for the specified <i>command</i> .
	ONMENT VA	
32397		ng environment variables shall affect the execution of <i>sccs</i> :
32398 32399 32400 32401	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
32402 32403	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
32404 32405 32406	LC_CTYPE	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).
32407 32408 32409	LC_MESSA	GES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
32410	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
32411	PROJECTDI	
32412 32413 32414 32415 32416		Provide a default value for the –d <i>path</i> option. If the value of <i>PROJECTDIR</i> begins with a slash, it shall be considered an absolute pathname; otherwise, the value of <i>PROJECTDIR</i> is treated as a user name and that user's initial working directory shall be examined for a subdirectory src or source . If such a directory is found, it shall be used. Otherwise, the value shall be used as a relative pathname.
32417 32418	Additional ecommand.	environment variable effects may be found in the utility description for the specified
32419 ASYNC 32420	CHRONOUS Default.	EVENTS
32421 STDOU 32422		ty description for the specified <i>command</i> .
32423 STDER		
32424	See the utilit	ty description for the specified <i>command</i> .

See the utility description for the specified *command*.

32425 OUTPUT FILES

32426

Utilities SCCS

32427 EXTENDED DESCRIPTION 32428 The following pseudo-utilities shall be supported as command operands. All options referred to 32429 in the following list are values given in the *options* operands following *command*. Equivalent to **info**, except that nothing shall be printed if nothing is being edited, and a check 32430 non-zero exit status shall be returned if anything is being edited. The intent is to have 32431 this included in an "install" entry in a makefile to ensure that everything is included 32432 32433 into the SCCS file before a version is installed. clean Remove everything from the current directory that can be recreated from SCCS files, 32434 32435 but do not remove any files being edited. If the -b option is given, branches shall be 32436 ignored in the determination of whether they are being edited; this is dangerous if branches are kept in the same directory. 32437 create Create an SCCS file, taking the initial contents from the file of the same name. Any 32438 options to admin are accepted. If the creation is successful, the original files shall be 32439 renamed by prefixing the basenames with a comma. These renamed files should be 32440 32441 removed after it has been verified that the SCCS files have been created successfully. delget Perform a delta on the named files and then get new versions. The new versions shall 32442 have ID keywords expanded and shall not be editable. Any -m, -p, -r, -s, and -y32443 options shall be passed to *delta*, and any $-\mathbf{b}$, $-\mathbf{c}$, $-\mathbf{e}$, $-\mathbf{i}$, $-\mathbf{k}$, $-\mathbf{l}$, $-\mathbf{s}$, and $-\mathbf{x}$ options shall be 32444 passed to *get*. 32445 **deledit** Equivalent to **delget**, except that the *get* phase shall include the -e option. This option 32446 32447 is useful for making a checkpoint of the current editing phase. The same options shall be passed to *delta* as described above, and all the options listed for *get* above except -e 32448 shall be passed to **edit**. 32449 diffs 32450 Write a difference listing between the current version of the files checked out for editing and the versions in SCCS format. Any $-\mathbf{r}$, $-\mathbf{c}$, $-\mathbf{i}$, $-\mathbf{x}$, and $-\mathbf{t}$ options shall be 32451 passed to get; any $-\mathbf{l}$, $-\mathbf{s}$, $-\mathbf{e}$, $-\mathbf{f}$, $-\mathbf{h}$, and $-\mathbf{b}$ options shall be passed to diff. A $-\mathbf{C}$ option 32452 32453 shall be passed to diff as $-\mathbf{c}$. edit 32454 Equivalent to $get - \mathbf{e}$. fix 32455 Remove the named delta, but leave a copy of the delta with the changes that were in it. It is useful for fixing small compiler bugs, and so on. The application shall ensure that it 32456 is followed by a -r SID option. Since fix does not leave audit trails, it should be used 32457 32458 carefully. info 32459 Write a listing of all files being edited. If the $-\mathbf{b}$ option is given, branches (that is, SIDs 32460 with two or fewer components) shall be ignored. If a $-\mathbf{u}$ user option is given, then only 32461 files being edited by the named user shall be listed. A -**U** option shall be equivalent to -u<current user>. 32462 Write out verbose information about the named files, equivalent to sccs prs. 32463 print tell Write a <newline>-separated list of the files being edited to standard output. Takes the 32464 $-\mathbf{b}$, $-\mathbf{u}$, and $-\mathbf{U}$ options like **info** and **check**. 32465 32466 unedit This is the opposite of an edit or a get -e. It should be used with caution, since any 32467 changes made since the *get* are lost. 32468 EXIT STATUS 32469

The following exit values shall be returned:

32470 Successful completion. Utilities SCCS

```
32471 >0 An error occurred.
```

32472 CONSEQUENCES OF ERRORS

Default.

32474 APPLICATION USAGE

Many of the SCCS utilities take directory names as operands as well as specific filenames. The pseudo-utilities supported by *sccs* are not described as having this capability, but are not prohibited from doing so.

32478 EXAMPLES

32473

32483

32487

32489

32491

32493

32498

32499

32500

32501 32502

32503 32504

32505 32506

32507

32508

1. To get a file for editing, edit it and produce a new delta:

```
32480 sccs get -e file.c
32481 ex file.c
32482 sccs delta file.c
```

2. To get a file from another directory:

```
32484 sccs -p /usr/src/sccs/s. get cc.c
32485 or:
32486 sccs get /usr/src/sccs/s.cc.c
```

3. To make a delta of a large number of files in the current directory:

```
32488 sccs delta *.c
```

4. To get a list of files being edited that are not on branches:

```
32490 sccs info -b
```

5. To delta everything being edited by the current user:

```
32492 sccs delta $(sccs tell -U)
```

6. In a makefile, to get source files from an SCCS file if it does not already exist:

```
32494 SRCS = st of source files>
32495 $ (SRCS):
32496 sccs get $ (REL) $@
```

32497 RATIONALE

SCCS and its associated utilities are part of the XSI Development Utilities option within the XSI extension.

SCCS is an abbreviation for Source Code Control System. It is a maintenance and enhancement tracking tool. When a file is put under SCCS, the source code control system maintains the file and, when changes are made, identifies and stores them in the file with the original source code and/or documentation. As other changes are made, they too are identified and retained in the file

Retrieval of the original and any set of changes is possible. Any version of the file as it develops can be reconstructed for inspection or additional modification. History data can be stored with each version, documenting why the changes were made, who made them, and when they were made.

SCCS Utilities

32509 FUTURE DIRECTIONS 32510 None. 32511 **SEE ALSO** 32512 admin, delta, get, make, prs, rmdel, sact, unget, val, what 32513 CHANGE HISTORY 32514 First released in Issue 4. 32515 **Issue 6** 32516 In the ENVIRONMENT VARIABLES section, the PROJECTDIR description is updated from "otherwise, the home directory of a user of that name is examined" to "otherwise, the value of 32517 *PROJECTDIR* is treated as a user name and that user's initial working directory is examined". 32518 32519 The normative text is reworded to avoid use of the term "must" for application requirements.

Utilities sed

22520 T	NAME			
32521	sed — stream	m editor		
32522	SYNOPSIS			
32523	sed [-n]	sed [-n] script[file]		
32524	sed [-n][-e script][-f script_file][file]		
32525]	DESCRIPTION			
32526		ity is a stream editor that shall read one or more text files, make editing changes		
32527 32528		o a script of editing commands, and write the results to standard output. The script nined from either the <i>script</i> operand string or a combination of the option-arguments		
32529		script and –f script_file options.		
32530 (OPTIONS			
32531	The <i>sed</i> utility	ty shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2,		
32532		ax Guidelines, except that the order of presentation of the -e and -f options is		
32533	significant.			
32534		ng options shall be supported:		
32535 32536	− e script	Add the editing commands specified by the <i>script</i> option-argument to the end of the script of editing commands. The <i>script</i> option-argument shall have the same		
32537		properties as the <i>script</i> operand, described in the OPERANDS section.		
32538	-f script_file	Add the editing commands in the file <i>script_file</i> to the end of the script.		
32539	-n	Suppress the default output (in which each line, after it is examined for editing, is		
32540		written to standard output). Only lines explicitly selected for output are written.		
32541 32542		and -f options may be specified. All commands shall be added to the script in the ied, regardless of their origin.		
32543	OPERANDS			
32544	The following	ng operands shall be supported:		
32545	file	A pathname of a file whose contents are read and edited. If multiple <i>file</i> operands		
32546 32547		are specified, the named files shall be read in the order specified and the concatenation shall be edited. If no <i>file</i> operands are specified, the standard input		
32548		shall be used.		
32549	script	A string to be used as the script of editing commands. The application shall not		
32550		present a <i>script</i> that violates the restrictions of a text file except that the final		
32551		character need not be a <newline>.</newline>		
32552 \$ 32553	STDIN The standar	d input shall be used only if no <i>file</i> operands are specified. See the INPUT FILES		
32554	section.	d input shall be used only if no me operands are specified. See the fivi of Fills		
32555]	INPUT FILES			
32556	The input fi	les shall be text files. The <i>script_files</i> named by the -f option shall consist of editing		
32557	commands.			
32558] 32559	ENVIRONMENT VA The followin	ARIABLES ng environment variables shall affect the execution of <i>sed</i> :		
32560	LANG	Provide a default value for the internationalization variables that are unset or null.		
32561		(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,		
32562 32563		Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)		
02000		and to determine the range of found enterporteen,		

sed Utilities

32564 LC_ALL If set to a non-empty string value, override the values of all the other 32565 internationalization variables. 32566 LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multi-32567 character collating elements within regular expressions. 32568 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 32569 characters (for example, single-byte as opposed to multi-byte characters in 32570 arguments and input files), and the behavior of character classes within regular 32571 32572 expressions. LC_MESSAGES 32573 Determine the locale that should be used to affect the format and contents of 32574 32575 diagnostic messages written to standard error. NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 32576 XSI

32577 ASYNCHRONOUS EVENTS

32578 Default.

32579 STDOUT

The input files shall be written to standard output, with the editing commands specified in the script applied. If the -n option is specified, only those input lines selected by the script shall be written to standard output.

32583 STDERR

32586

32594 32595

32596

32597

32598

32599

32600

32584 The standard error shall be used only for diagnostic messages.

32585 OUTPUT FILES

The output files shall be text files whose formats are dependent on the editing commands given.

32587 EXTENDED DESCRIPTION

32588 The *script* shall consist of editing commands of the following form:

32589 [address[,address]] function

where *function* represents a single-character command verb from the list in **Editing Commands** in **sed** (on page 843), followed by any applicable arguments.

The command can be preceded by
blank>s and/or semicolons. The function can be preceded by
by
blank>s. These optional characters shall have no effect.

In default operation, *sed* cyclically shall append a line of input, less its terminating <newline>, into the pattern space. Normally the pattern space will be empty, unless a **D** command terminated the last cycle. The *sed* utility shall then apply in sequence all commands whose addresses select that pattern space, and at the end of the script copy the pattern space to standard output (except when -**n** is specified) and delete the pattern space. Whenever the pattern space is written to standard output or a named file, *sed* shall immediately follow it with a <newline>.

Some of the editing commands use a hold space to save all or part of the pattern space for subsequent retrieval. The pattern and hold spaces shall each be able to hold at least 8 192 bytes.

sed **Utilities**

Addresses in sed

32603

32604

32605

32607

32608

32609 32610

32611

32612

32613

32614

32615

32616 32617

32618

32619 32620

32621

32622

32623 32624

32625

32626 32627

32628

32629

32630

32631 32632

32633

32634

32635 32636

32637

32638 32639

32640

32641

32642

32643

32644 32645

An address is either a decimal number that counts input lines cumulatively across files, a '\$' character that addresses the last line of input, or a context address (which consists of a BRE, as described in **Regular Expressions in sed**, preceded and followed by a delimiter, usually a slash). 32606

An editing command with no addresses shall select every pattern space.

An editing command with one address shall select each pattern space that matches the address.

An editing command with two addresses shall select the inclusive range from the first pattern space that matches the first address through the next pattern space that matches the second. (If the second address is a number less than or equal to the line number first selected, only one line shall be selected.) Starting at the first line following the selected range, sed shall look again for the first address. Thereafter, the process shall be repeated. Omitting either or both of the address components in the following form produces undefined results:

[address[,address]]

Regular Expressions in sed

The sed utility shall support the BREs described in the Base Definitions volume of IEEE Std 1003.1-2001, Section 9.3, Basic Regular Expressions, with the following additions:

- ullet In a context address, the construction "\cBREc", where c is any character other than backslash or <newline>, shall be identical to "/BRE/". If the character designated by cappears following a backslash, then it shall be considered to be that literal character, which shall not terminate the BRE. For example, in the context address "\xabc\xdefx", the second x stands for itself, so that the BRE is "abcxdef".
- The escape sequence '\n' shall match a <newline> embedded in the pattern space. A literal <newline> shall not be used in the BRE of a context address or in the substitute function.
- If an RE is empty (that is, no pattern is specified) sed shall behave as if the last RE used in the last command applied (either as an address or as part of a substitute command) was specified.

Editing Commands in sed

In the following list of editing commands, the maximum number of permissible addresses for each function is indicated by [0addr], [1addr], or [2addr], representing zero, one, or two addresses.

The argument *text* shall consist of one or more lines. Each embedded <newline> in the text shall be preceded by a backslash. Other backslashes in text shall be removed, and the following character shall be treated literally.

The **r** and **w** command verbs, and the *w* flag to the **s** command, take an optional *rfile* (or *wfile*) parameter, separated from the command verb letter or flag by one or more

 blank>s; implementations may allow zero separation as an extension.

The argument rfile or the argument wfile shall terminate the editing command. Each wfile shall be created before processing begins. Implementations shall support at least ten wfile arguments in the script; the actual number (greater than or equal to 10) that is supported by the implementation is unspecified. The use of the *wfile* parameter shall cause that file to be initially created, if it does not exist, or shall replace the contents of an existing file.

The b, r, s, t, w, y, and: command verbs shall accept additional arguments. The following synopses indicate which arguments shall be separated from the command verbs by a single

sed Utilities

32646 <space>.

The $\bf a$ and $\bf r$ commands schedule text for later output. The text specified for the $\bf a$ command, and the contents of the file specified for the $\bf r$ command, shall be written to standard output just before the next attempt to fetch a line of input when executing the $\bf N$ or $\bf n$ commands, or when reaching the end of the script. If written when reaching the end of the script, and the $-\bf n$ option was not specified, the text shall be written after copying the pattern space to standard output. The contents of the file specified for the $\bf r$ command shall be as of the time the output is written, not the time the $\bf r$ command is applied. The text shall be output in the order in which the $\bf a$ and $\bf r$ commands were applied to the input.

Command verbs other than {, a, b, c, i, r, t, w, :, and # can be followed by a semicolon, optional <blank>s, and another command verb. However, when the s command verb is used with the w flag, following it with another command in this manner produces undefined results.

A function can be preceded by one or more '!' characters, in which case the function shall be applied if the addresses do not select the pattern space. Zero or more <blank>s shall be accepted before the first '!' character. It is unspecified whether <blank>s can follow a '!' character, and conforming applications shall not follow a '!' character with <blank>s.

[2addr] {function

function

32664 ... 32665 }

Execute a list of *sed* functions only when the pattern space is selected. The list of *sed* functions shall be surrounded by braces and separated by <newline>s, and conform to the following rules. The braces can be preceded or followed by <blank>s. The functions can be preceded by <blank>s, but shall not be followed by <blank>s. The <right-brace> shall be preceded by a <newline> and can be preceded or followed by <blank>s.

[1addr]**a**\

text Write text to standard output as described previously.

[2addr]**b** [label]

Branch to the: function bearing the *label*. If *label* is not specified, branch to the end of the script. The implementation shall support *labels* recognized as unique up to at least 8 characters; the actual length (greater than or equal to 8) that shall be supported by the implementation is unspecified. It is unspecified whether exceeding a label length causes an error or a silent truncation.

[2addr]c\

text Delete the pattern space. With a 0 or 1 address or at the end of a 2-address range, place *text* on the output and start the next cycle.

[2addr]d Delete the pattern space and start the next cycle.

[2addr]**D** Delete the initial segment of the pattern space through the first <newline> and start the next cycle.

[2addr]g Replace the contents of the pattern space by the contents of the hold space.

[2addr]G Append to the pattern space a <newline> followed by the contents of the hold space.

[2addr]h Replace the contents of the hold space with the contents of the pattern space.

[2addr]H Append to the hold space a <newline> followed by the contents of the pattern space.

Utilities sed

32691 32692	[1addr] i \ text	Write <i>text</i> to standard output.
32693 32694 32695 32696 32697 32698 32699 32700 32701	[2addr]l	(The letter ell.) Write the pattern space to standard output in a visually unambiguous form. The characters listed in the Base Definitions volume of IEEE Std 1003.1-2001, Table 5-1, Escape Sequences and Associated Actions ('\\', '\a', '\b', '\f', '\r', '\t', '\v') shall be written as the corresponding escape sequence; the '\n' in that table is not applicable. Non-printable characters not in that table shall be written as one three-digit octal number (with a preceding backslash) for each byte in the character (most significant byte first). If the size of a byte on the system is greater than 9 bits, the format used for non-printable characters is implementation-defined.
32702 32703 32704 32705		Long lines shall be folded, with the point of folding indicated by writing a backslash followed by a <newline>; the length at which folding occurs is unspecified, but should be appropriate for the output device. The end of each line shall be marked with a '$\\$'.</newline>
32706 32707 32708	[2addr]n	Write the pattern space to standard output if the default output has not been suppressed, and replace the pattern space with the next line of input, less its terminating <newline>.</newline>
32709 32710		If no next line of input is available, the ${\bf n}$ command verb shall branch to the end of the script and quit without starting a new cycle.
32711 32712 32713	[2addr]N	Append the next line of input, less its terminating <newline>, to the pattern space, using an embedded <newline> to separate the appended material from the original material. Note that the current line number changes.</newline></newline>
32714 32715 32716		If no next line of input is available, the N command verb shall branch to the end of the script and quit without starting a new cycle or copying the pattern space to standard output.
32717	[2addr] p	Write the pattern space to standard output.
32718	[2addr] P	Write the pattern space, up to the first <newline>, to standard output.</newline>
32719	[1addr]q	Branch to the end of the script and quit without starting a new cycle.
32720 32721 32722	[1addr]r rfile	Copy the contents of <i>rfile</i> to standard output as described previously. If <i>rfile</i> does not exist or cannot be read, it shall be treated as if it were an empty file, causing no error condition.
32723 32724 32725 32726 32727 32728	[2addr]s/BRE	Substitute the replacement string for instances of the BRE in the pattern space. Any character other than backslash or <newline> can be used instead of a slash to delimit the BRE and the replacement. Within the BRE and the replacement, the BRE delimiter itself can be used as a literal character if it is preceded by a backslash.</newline>
32729 32730 32731 32732 32733 32734 32735 32736		The replacement string shall be scanned from beginning to end. An ampersand $('\&')$ appearing in the replacement shall be replaced by the string matching the BRE. The special meaning of $'\&'$ in this context can be suppressed by preceding it by a backslash. The characters $"\n"$, where n is a digit, shall be replaced by the text matched by the corresponding backreference expression. The special meaning of $"\n"$ where n is a digit in this context, can be suppressed by preceding it by a backslash. For each other backslash $('\n')$ encountered, the following character shall lose its special meaning (if any). The meaning of a $'\n'$ immediately followed

sed Utilities

32737 32738			acter other than ' &', ' \ ', a digit, or the delimiter character used for d, is unspecified.
32739 32740 32741 32742 32743 32744		the <newline shall be con identical to meaning of</newline 	esplit by substituting a <newline> into it. The application shall escape e> in the replacement by preceding it by a backslash. A substitution sidered to have been performed even if the replacement string is the string that it replaces. Any backslash used to alter the default a subsequent character shall be discarded from the BRE or the before evaluating the BRE or using the replacement.</newline>
32745		The value of	flags shall be zero or more of:
32746 32747		n	Substitute for the n th occurrence only of the BRE found within the pattern space.
32748 32749 32750		g	Globally substitute for all non-overlapping instances of the BRE rather than just the first one. If both ${\bf g}$ and ${\bf n}$ are specified, the results are unspecified.
32751 32752		p	Write the pattern space to standard output if a replacement was made.
32753 32754 32755 32756		w wfile	Write. Append the pattern space to <i>wfile</i> if a replacement was made. A conforming application shall precede the <i>wfile</i> argument with one or more slank>s. If the w flag is not the last flag value given in a concatenation of multiple flag values, the results are undefined.
32757 32758 32759 32760	[2addr]t [labe	Test. Branch made since t	to the: command verb bearing the <i>label</i> if any substitutions have been he most recent reading of an input line or execution of a t . If <i>label</i> is , branch to the end of the script.
32761 32762	[2addr]w wfil		te) the pattern space to <i>wfile</i> .
32763	[2addr] x	Exchange the	e contents of the pattern and hold spaces.
32764 32765 32766 32767 32768 32769 32770 32771 32772 32773 32774 32775 32776	[2addr]y/strin	Replace all occurrences of characters in <i>string1</i> with the corresponding characters in <i>string2</i> . If a backslash followed by an 'n' appear in <i>string1</i> or <i>string2</i> , the two characters shall be handled as a single <newline>. If the number of characters in <i>string1</i> and <i>string2</i> are not equal, or if any of the characters in <i>string1</i> appear more than once, the results are undefined. Any character other than backslash or <newline> can be used instead of slash to delimit the strings. If the delimiter is not <i>n</i>, within <i>string1</i> and <i>string2</i>, the delimiter itself can be used as a literal character if it is preceded by a backslash. If a backslash character is immediately followed by a backslash character in <i>string1</i> or <i>string2</i>, the two backslash characters shall be counted as a single literal backslash character. The meaning of a backslash followed by any character that is not 'n', a backslash, or the delimiter character is undefined.</newline></newline>	
32777	[0addr]:label	Do nothing.	This command bears a <i>label</i> to which the ${f b}$ and ${f t}$ commands branch.
32778	[1addr]=	Write the foll	lowing to standard output:
32779		"%d\n", <	current line number>
32780	[0addr]	Ignore this en	mpty command.

Utilities sed

32781 [0addr]# Ignore the '#' and the remainder of the line (treat them as a comment), with the 32782 single exception that if the first two characters in the script are "#n", the default 32783 output shall be suppressed; this shall be the equivalent of specifying $-\mathbf{n}$ on the command line. 32784 32785 EXIT STATUS The following exit values shall be returned:

32786

Successful completion. 32787

>0 An error occurred. 32788

32789 CONSEQUENCES OF ERRORS

Default. 32790

32791 APPLICATION USAGE

Regular expressions match entire strings, not just individual lines, but a <newline> is matched by ' \n' in a sed RE; a <newline> is not allowed by the general definition of regular expression in IEEE Std 1003.1-2001. Also note that '\n' cannot be used to match a <newline> at the end of an arbitrary input line; <newline>s appear in the pattern space as a result of the N editing command.

32797 EXAMPLES

32792 32793

32794

32795 32796

32798 32799 This sed script simulates the BSD cat -s command, squeezing excess blank lines from standard input.

```
32800
            sed -n '
            # Write non-empty lines.
32801
32802
            /./ {
32803
                р
32804
                 d
32805
            # Write a single empty line, then look for more empty lines.
32806
32807
            /^$/
                     р
            # Get next line, discard the held <newline> (empty line),
32808
32809
            # and look for more empty lines.
            :Empty
32810
            /^$/
32811
32812
                Ν
32813
                 s/.//
32814
                b Empty
32815
32816
            # Write the non-empty line before going back to search
            # for the first in a set of empty lines.
32817
32818
32819
```

32820 RATIONALE

32821 32822

32823

32824 32825

32826

This volume of IEEE Std 1003.1-2001 requires implementations to support at least ten distinct wfiles, matching historical practice on many implementations. Implementations are encouraged to support more, but conforming applications should not exceed this limit.

The exit status codes specified here are different from those in System V. System V returns 2 for garbled *sed* commands, but returns zero with its usage message or if the input file could not be opened. The standard developers considered this to be a bug.

sed Utilities

The manner in which the **l** command writes non-printable characters was changed to avoid the historical backspace-overstrike method, and other requirements to achieve unambiguous output were added. See the RATIONALE for *ed* for details of the format chosen, which is the same as that chosen for *sed*.

This volume of IEEE Std 1003.1-2001 requires implementations to provide pattern and hold spaces of at least 8192 bytes, larger than the 4000 bytes spaces used by some historical implementations, but less than the 20480 bytes limit used in an early proposal. Implementations are encouraged to allocate dynamically larger pattern and hold spaces as needed.

The treatment of '#' comments differs from the SVID which only allows a comment as the first line of the script, but matches BSD-derived implementations. The comment character is treated as a command, and it has the same properties in terms of being accepted with leading
blank>s; the BSD implementation has historically supported this.

Early proposals required that a *script_file* have at least one non-comment line. Some historical implementations have behaved in unexpected ways if this were not the case. The standard developers considered that this was incorrect behavior and that application developers should not have to avoid this feature. A correct implementation of this volume of IEEE Std 1003.1-2001 shall permit *script_files* that consist only of comment lines.

Early proposals indicated that if -e and -f options were intermixed, all -e options were processed before any -f options. This has been changed to process them in the order presented because it matches historical practice and is more intuitive.

The treatment of the \mathbf{p} flag to the \mathbf{s} command differs between System V and BSD-based systems when the default output is suppressed. In the two examples:

```
echo a | sed 's/a/A/p'
echo a | sed -n 's/a/A/p'
```

this volume of IEEE Std 1003.1-2001, BSD, System V documentation, and the SVID indicate that the first example should write two lines with $\bf A$, whereas the second should write one. Some System V systems write the $\bf A$ only once in both examples because the $\bf p$ flag is ignored if the $\bf -n$ option is not specified.

This is a case of a diametrical difference between systems that could not be reconciled through the compromise of declaring the behavior to be unspecified. The SVID/BSD/System V documentation behavior was adopted for this volume of IEEE Std 1003.1-2001 because:

- No known documentation for any historic system describes the interaction between the **p** flag and the **-n** option.
- The selected behavior is more correct as there is no technical justification for any interaction between the $\bf p$ flag and the $\bf -n$ option. A relationship between $\bf -n$ and the $\bf p$ flag might imply that they are only used together, but this ignores valid scripts that interrupt the cyclical nature of the processing through the use of the $\bf D$, $\bf d$, $\bf q$, or branching commands. Such scripts rely on the $\bf p$ suffix to write the pattern space because they do not make use of the default output at the "bottom" of the script.

 Utilities sed

• Because the $-\mathbf{n}$ option makes the \mathbf{p} flag unnecessary, any interaction would only be useful if sed scripts were written to run both with and without the $-\mathbf{n}$ option. This is believed to be unlikely. It is even more unlikely that programmers have coded the \mathbf{p} flag expecting it to be unnecessary. Because the interaction was not documented, the likelihood of a programmer discovering the interaction and depending on it is further decreased.

• Finally, scripts that break under the specified behavior produce too much output instead of too little, which is easier to diagnose and correct.

The form of the substitute command that uses the n suffix was limited to the first 512 matches in an early proposal. This limit has been removed because there is no reason an editor processing lines of {LINE_MAX} length should have this restriction. The command s/a/A/2047 should be able to substitute the 2 047th occurrence of a on a line.

The **b**, **t**, and : commands are documented to ignore leading white space, but no mention is made of trailing white space. Historical implementations of *sed* assigned different locations to the labels 'x' and "x ". This is not useful, and leads to subtle programming errors, but it is historical practice, and changing it could theoretically break working scripts. Implementors are encouraged to provide warning messages about labels that are never used or jumps to labels that do not exist.

Historically, the *sed*! and } editing commands did not permit multiple commands on a single line using a semicolon as a command delimiter. Implementations are permitted, but not required, to support this extension.

32893 FUTURE DIRECTIONS

32894 None.

32895 SEE ALSO

32878 32879

32880 32881

32882

32883

32884

32885

32886 32887

32888

32889

32890

32891 32892

32896 awk, ed, grep

32897 CHANGE HISTORY

32898 First released in Issue 2.

32899 Issue 5

32900 The FUTURE DIRECTIONS section is added.

32901 Issue 6

32902

32903

32904 32905 The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:

• Implementations are required to support at least ten *wfile* arguments in an editing command.

The EXTENDED DESCRIPTION is changed to align with the IEEE P1003.2b draft standard.

32906 IEEE PASC Interpretation 1003.2 #190 is applied.

IEEE PASC Interpretation 1003.2 #203 is applied, clarifying the meaning of the backslash escape sequences in a replacement string for a BRE.

sh Utilities

32909 NAME 32910 sh — shell, the standard command language interpreter 32911 SYNOPSIS 32912 sh [-abCefhimnuvx] [-o option] [+abCefhimnuvx] [+o option] 32913 [command file [argument...]] sh -c[-abCefhimnuvx][-o option][+abCefhimnuvx][+o option]command string 32914 32915 [command name [argument...]] 32916 sh -s[-abCefhimnuvx][-o option][+abCefhimnuvx][+o option][argument] 32917 **DESCRIPTION** The sh utility is a command language interpreter that shall execute commands read from a 32918 32919 command line string, the standard input, or a specified file. The application shall ensure that the 32920 commands to be executed are expressed in the language described in Chapter 2 (on page 29). Pathname expansion shall not fail due to the size of a file. 32921 Shell input and output redirections have an implementation-defined offset maximum that is 32922 established in the open file description. 32923 32924 OPTIONS The *sh* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, 32925 32926 Utility Syntax Guidelines, with an extension for support of a leading plus sign ('+') as noted below. 32927 32928 The -a, -b, -C, -e, -f, -m, -n, -o option, -u, -v, and -x options are described as part of the set utility in Section 2.14 (on page 64). The option letters derived from the set special built-in shall 32929 32930 also be accepted with a leading plus sign ('+') instead of a leading hyphen (meaning the reverse case of the option as described in this volume of IEEE Std 1003.1-2001). 32931 32932 The following additional options shall be supported: Read commands from the command_string operand. Set the value of special 32933 -C 32934 parameter 0 (see Section 2.5.2 (on page 34)) from the value of the command_name operand and the positional parameters (\$1, \$2, and so on) in sequence from the 32935 32936 remaining argument operands. No commands shall be read from the standard 32937 input. -iSpecify that the shell is interactive; see below. An implementation may treat 32938 specifying the -i option as an error if the real user ID of the calling process does 32939 32940 not equal the effective user ID or if the real group ID does not equal the effective 32941 group ID. Read commands from the standard input. 32942 32943 If there are no operands and the -c option is not specified, the -s option shall be assumed. If the -i option is present, or if there are no operands and the shell's standard input and standard 32944 error are attached to a terminal, the shell is considered to be *interactive*. 32945 32946 OPERANDS 32947 The following operands shall be supported: A single hyphen shall be treated as the first operand and then ignored. If both '-' 32948 and "--" are given as arguments, or if other operands precede the single hyphen, 32949 the results are undefined. 32950

argument

32951

The positional parameters (\$1, \$2, and so on) shall be set to *arguments*, if any.

Utilities sh

32952 command_file The pathname of a file containing commands. If the pathname contains one or more slash characters, the implementation attempts to read that file; the file need not be executable. If the pathname does not contain a slash character:

- The implementation shall attempt to read that file from the current working directory; the file need not be executable.
- If the file is not in the current working directory, the implementation may perform a search for an executable file using the value of *PATH*, as described in Section 2.9.1.1 (on page 48).

Special parameter 0 (see Section 2.5.2 (on page 34)) shall be set to the value of *command_file*. If *sh* is called using a synopsis form that omits *command_file*, special parameter 0 shall be set to the value of the first argument passed to *sh* from its parent (for example, *argv*[0] for a C program), which is normally a pathname used to execute the *sh* utility.

command name

A string assigned to special parameter 0 when executing the commands in *command_string*. If *command_name* is not specified, special parameter 0 shall be set to the value of the first argument passed to *sh* from its parent (for example, *argv*[0] for a C program), which is normally a pathname used to execute the *sh* utility.

command string

A string that shall be interpreted by the shell as one or more commands, as if the string were the argument to the <code>system()</code> function defined in the System Interfaces volume of IEEE Std 1003.1-2001. If the <code>command_string</code> operand is an empty string, <code>sh</code> shall exit with a zero exit status.

STDIN

The standard input shall be used only if one of the following is true:

- The –**s** option is specified.
- \bullet The –c option is not specified and no operands are specified.
 - The script executes one or more commands that require input from standard input (such as a read command that does not redirect its input).

See the INPUT FILES section.

When the shell is using standard input and it invokes a command that also uses standard input, the shell shall ensure that the standard input file pointer points directly after the command it has read when the command begins execution. It shall not read ahead in such a manner that any characters intended to be read by the invoked command are consumed by the shell (whether interpreted by the shell or not) or that characters that are not read by the invoked command are not seen by the shell. When the command expecting to read standard input is started asynchronously by an interactive shell, it is unspecified whether characters are read by the command or interpreted by the shell.

If the standard input to *sh* is a FIFO or terminal device and is set to non-blocking reads, then *sh* shall enable blocking reads on standard input. This shall remain in effect when the command completes.

32993 INPUT FILES

The input file shall be a text file, except that line lengths shall be unlimited. If the input file is empty or consists solely of blank lines or comments, or both, *sh* shall exit with a zero exit status.

sh Utilities

32996 ENVIRONMENT VARIABLES 32997 The following environment variables shall affect the execution of *sh*: 32998 This variable, when and only when an interactive shell is invoked, shall be subjected to parameter expansion (see Section 2.6.2 (on page 37)) by the shell, and 32999 the resulting value shall be used as a pathname of a file containing shell 33000 commands to execute in the current environment. The file need not be executable. 33001 If the expanded value of ENV is not an absolute pathname, the results are 33002 33003 unspecified. ENV shall be ignored if the real and effective user IDs or real and effective group IDs of the process are different. 33004 33005 **FCEDIT** This variable, when expanded by the shell, shall determine the default value for the -e editor option's editor option-argument. If FCEDIT is null or unset, ed shall be 33006 used as the editor. This volume of IEEE Std 1003.1-2001 specifies the effects of this 33007 variable only for systems supporting the User Portability Utilities option. 33008 HISTFILE Determine a pathname naming a command history file. If the HISTFILE variable is 33009 33010 33011 33012 33013 33014 33015

not set, the shell may attempt to access or create a file .sh_history in the directory referred to by the *HOME* environment variable. If the shell cannot obtain both read and write access to, or create, the history file, it shall use an unspecified mechanism that allows the history to operate properly. (References to history "file" in this section shall be understood to mean this unspecified mechanism in such cases.) An implementation may choose to access this variable only when initializing the history file; this initialization shall occur when fc or sh first attempt to retrieve entries from, or add entries to, the file, as the result of commands issued by the user, the file named by the ENV variable, or implementation-defined system start-up files. Implementations may choose to disable the history list mechanism for users with appropriate privileges who do not set HISTFILE; the specific circumstances under which this occurs are implementation-defined. If more than one instance of the shell is using the same history file, it is unspecified how updates to the history file from those shells interact. As entries are deleted from the history file, they shall be deleted oldest first. It is unspecified when history file entries are physically removed from the history file. This volume of IEEE Std 1003.1-2001 specifies the effects of this variable only for systems supporting the User Portability Utilities option.

HISTSIZE Determine a decimal number representing the limit to the number of previous commands that are accessible. If this variable is unset, an unspecified default greater than or equal to 128 shall be used. The maximum number of commands in the history list is unspecified, but shall be at least 128. An implementation may choose to access this variable only when initializing the history file, as described under HISTFILE. Therefore, it is unspecified whether changes made to HISTSIZE after the history file has been initialized are effective.

Determine the pathname of the user's home directory. The contents of *HOME* are used in tilde expansion as described in Section 2.6.1 (on page 37). This volume of IEEE Std 1003.1-2001 specifies the effects of this variable only for systems supporting the User Portability Utilities option.

(Input Field Separators.) A string treated as a list of characters that shall be used for field splitting and to split lines into words with the *read* command. See Section 2.6.5 (on page 42). If *IFS* is not set, the shell shall behave as if the value of *IFS* were <space>, <tab>, and <newline>. Implementations may ignore the value of *IFS* in the environment at the time *sh* is invoked, treating *IFS* as if it were not set.

33016 33017

33018

33019

33020

33021 33022

33023

33024

33025

33026

33027

33028 33029

33030

33031

33032

33033

33034

33035

33036

33037

33038

33039

33040

33041

33042 33043 **HOME**

IFS

Utilities sh

33044 33045 33046 33047	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
33048 33049	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
33050 33051 33052	LC_COLLAT	E Determine the behavior of range expressions, equivalence classes, and multicharacter collating elements within pattern matching.
33053 33054 33055 33056	LC_CTYPE	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), which characters are defined as letters (character class alpha), and the behavior of character classes within pattern matching.
33057 33058 33059	LC_MESSAC	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
33060 33061 33062 33063 33064 33065 33066 33067 33068	MAIL	Determine a pathname of the user's mailbox file for purposes of incoming mail notification. If this variable is set, the shell shall inform the user if the file named by the variable is created or if its modification time has changed. Informing the user shall be accomplished by writing a string of unspecified format to standard error prior to the writing of the next primary prompt string. Such check shall be performed only after the completion of the interval defined by the <i>MAILCHECK</i> variable after the last such check. The user shall be informed only if <i>MAIL</i> is set and <i>MAILPATH</i> is not set. This volume of IEEE Std 1003.1-2001 specifies the effects of this variable only for systems supporting the User Portability Utilities option.
33069 33070 33071 33072 33073 33074 33075	MAILCHEC	Establish a decimal integer value that specifies how often (in seconds) the shell shall check for the arrival of mail in the files specified by the <i>MAILPATH</i> or <i>MAIL</i> variables. The default value shall be 600 seconds. If set to zero, the shell shall check before issuing each primary prompt. This volume of IEEE Std 1003.1-2001 specifies the effects of this variable only for systems supporting the User Portability Utilities option.
33076 33077 33078 33079 33080 33081 33082 33083	MAILPATH	Provide a list of pathnames and optional messages separated by colons. If this variable is set, the shell shall inform the user if any of the files named by the variable are created or if any of their modification times change. (See the preceding entry for <i>MAIL</i> for descriptions of mail arrival and user informing.) Each pathname can be followed by '%' and a string that shall be subjected to parameter expansion and written to standard error when the modification time changes. If a '%' character in the pathname is preceded by a backslash, it shall be treated as a literal '%' in the pathname. The default message is unspecified.
33084 33085 33086		The <i>MAILPATH</i> environment variable takes precedence over the <i>MAIL</i> variable. This volume of IEEE Std 1003.1-2001 specifies the effects of this variable only for systems supporting the User Portability Utilities option.
33087 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.
33088 33089 33090	PATH	Establish a string formatted as described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables, used to effect command interpretation; see Section 2.9.1.1 (on page 48).

sh Utilities

This variable shall represent an absolute pathname of the current working directory. Assignments to this variable may be ignored unless the value is an absolute pathname of the current working directory and there are no filename components of dot or dot-dot.

33095 ASYNCHRONOUS EVENTS

33096 Default.

33097 **STDOUT**

33098 See the STDERR section.

33099 **STDERR**

33107

33113

33118

33119 33120

33121

Except as otherwise stated (by the descriptions of any invoked utilities or in interactive mode), standard error shall be used only for diagnostic messages.

33102 OUTPUT FILES

33103 None.

33104 EXTENDED DESCRIPTION

See Chapter 2. The following additional capabilities are supported on systems supporting the User Portability Utilities option.

Command History List

When the *sh* utility is being used interactively, it shall maintain a list of commands previously entered from the terminal in the file named by the *HISTFILE* environment variable. The type, size, and internal format of this file are unspecified. Multiple *sh* processes can share access to the file for a user, if file access permissions allow this; see the description of the *HISTFILE* environment variable.

Command Line Editing

When *sh* is being used interactively from a terminal, the current command and the command history (see *fc*) can be edited using *vi*-mode command line editing. This mode uses commands, described below, similar to a subset of those described in the *vi* utility. Implementations may offer other command line editing modes corresponding to other editing utilities.

The command *set* **–o** *vi* shall enable *vi*-mode editing and place *sh* into *vi* insert mode (see **Command Line Editing (vi-mode)** (on page 855)). This command also shall disable any other editing mode that the implementation may provide. The command *set* **+o** *vi* disables *vi*-mode editing.

Certain block-mode terminals may be unable to support shell command line editing. If a terminal is unable to provide either edit mode, it need not be possible to *set* –**o** *vi* when using the shell on this terminal.

In the following sections, the characters *erase*, *interrupt*, *kill*, and *end-of-file* are those set by the stty utility.

sh **Utilities**

Command Line Editing (vi-mode)

33127

33128 33129

33130 33131

33132 33133

33134

33135

33136 33137

33138

33139

33140

33141

33142

33143

33144

33145 33146

33147

33148

33149

33150

33151

33152

33160 33161

33162 33163

33164 33165

33166

33167

33168

33169

end-of-file

In vi editing mode, there shall be a distinguished line, the edit line. All the editing operations which modify a line affect the edit line. The edit line is always the newest line in the command history buffer.

With *vi*-mode enabled, *sh* can be switched between insert mode and command mode.

When in insert mode, an entered character shall be inserted into the command line, except as noted in **vi Line Editing Insert Mode**. Upon entering *sh* and after termination of the previous command. *sh* shall be in insert mode.

Typing an escape character shall switch sh into command mode (see vi Line Editing Command Mode (on page 856)). In command mode, an entered character shall either invoke a defined operation, be used as part of a multi-character operation, or be treated as an error. A character that is not recognized as part of an editing command shall terminate any specific editing command and shall alert the terminal. Typing the interrupt character in command mode shall cause *sh* to terminate command line editing on the current command line, reissue the prompt on the next line of the terminal, and reset the command history (see fc) so that the most recently executed command is the previous command (that is, the command that was being edited when it was interrupted is not reentered into the history).

In the following sections, the phrase "move the cursor to the beginning of the word" shall mean "move the cursor to the first character of the current word" and the phrase "move the cursor to the end of the word" shall mean "move the cursor to the last character of the current word". The phrase "beginning of the command line" indicates the point between the end of the prompt string issued by the shell (or the beginning of the terminal line, if there is no prompt string) and the first character of the command text.

vi Line Editing Insert Mode

While in insert mode, any character typed shall be inserted in the current command line, unless it is from the following set.

Execute the current command line. If the current command line is not empty, this 33153 <newline> 33154 line shall be entered into the command history (see fc). Delete the character previous to the current cursor position and move the current 33155 erase 33156 cursor position back one character. In insert mode, characters shall be erased from 33157 both the screen and the buffer when backspacing. 33158 interrupt Terminate command line editing with the same effects as described for 33159

interrupting command mode; see **Command Line Editing (vi-mode)**.

kill Clear all the characters from the input line.

<control>-V Insert the next character input, even if the character is otherwise a special insert mode character.

<control>-W Delete the characters from the one preceding the cursor to the preceding word boundary. The word boundary in this case is the closer to the cursor of either the beginning of the line or a character that is in neither the **blank** nor **punct** character classification of the current locale.

> Interpreted as the end of input in sh. This interpretation shall occur only at the beginning of an input line. If end-of-file is entered other than at the beginning of the line, the results are unspecified.

sh Utilities

<ESC> 33170 Place *sh* into command mode. vi Line Editing Command Mode 33171 In command mode for the command line editing feature, decimal digits not beginning with 0 33172 that precede a command letter shall be remembered. Some commands use these decimal digits 33173 as a count number that affects the operation. 33174 The term *motion command* represents one of the commands: 33175 <space> \$ E f T 33176 В e h 33177 If the current line is not the edit line, any command that modifies the current line shall cause the content of the current line to replace the content of the edit line, and the current line shall 33178 become the edit line. This replacement cannot be undone (see the u and U commands below). 33179 The modification requested shall then be performed to the edit line. When the current line is the 33180 edit line, the modification shall be done directly to the edit line. 33181 Any command that is preceded by *count* shall take a count (the numeric value of any preceding 33182 decimal digits). Unless otherwise noted, this count shall cause the specified operation to repeat 33183 by the number of times specified by the count. Also unless otherwise noted, a count that is out of 33184 range is considered an error condition and shall alert the terminal, but neither the cursor 33185 position, nor the command line, shall change. 33186 The terms word and bigword are used as defined in the vi description. The term save buffer 33187 corresponds to the term *unnamed buffer* in *vi*. 33188 The following commands shall be recognized in command mode: 33189 33190 <newline> Execute the current command line. If the current command line is not empty, this line shall be entered into the command history (see *fc*). 33191 33192 <control>-L Redraw the current command line. Position the cursor at the same location on the 33193 redrawn line. 33194 # Insert the character '#' at the beginning of the current command line and treat the resulting edit line as a comment. This line shall be entered into the command 33195 33196 history; see fc. Display the possible shell word expansions (see Section 2.6 (on page 36)) of the 33197 bigword at the current command line position. 33198 This does not modify the content of the current line, and therefore does not Note: 33199 cause the current line to become the edit line. 33200 These expansions shall be displayed on subsequent terminal lines. If the bigword 33201 contains none of the characters '?', '*', or '[', an asterisk ('*') shall be 33202 implicitly assumed at the end. If any directories are matched, these expansions 33203 33204 shall have a '/' character appended. After the expansion, the line shall be redrawn, the cursor repositioned at the current cursor position, and sh shall be 33205 placed in command mode. 33206 / Perform pathname expansion (see Section 2.6.6 (on page 42)) on the current 33207 bigword, up to the largest set of characters that can be matched uniquely. If the 33208 33209 bigword contains none of the characters '?', '*', or '[', an asterisk ('*') shall be implicitly assumed at the end. This maximal expansion then shall replace the 33210 original bigword in the command line, and the cursor shall be placed after this 33211 expansion. If the resulting bigword completely and uniquely matches a directory, a 33212 '/' character shall be inserted directly after the bigword. If some other file is 33213 33214 completely matched, a single <space> shall be inserted after the bigword. After

Utilities sh

33215		this operation, <i>sh</i> shall be placed in insert mode.
33216 33217 33218 33219 33220 33221 33222 33222	*	Perform pathname expansion on the current bigword and insert all expansions into the command to replace the current bigword, with each expansion separated by a single <space>. If at the end of the line, the current cursor position shall be moved to the first column position following the expansions and <i>sh</i> shall be placed in insert mode. Otherwise, the current cursor position shall be the last column position of the first character after the expansions and <i>sh</i> shall be placed in insert mode. If the current bigword contains none of the characters '?', '*', or '[', before the operation, an asterisk shall be implicitly assumed at the end.</space>
33224 33225 33226 33227 33228	@letter	Insert the value of the alias named _letter. The symbol letter represents a single alphabetic character from the portable character set; implementations may support additional characters as an extension. If the alias _letter contains other editing commands, these commands shall be performed as part of the insertion. If no alias _letter is enabled, this command shall have no effect.
33229 33230 33231 33232 33233 33234 33235 33236 33237	[count]~	Convert, if the current character is a lowercase letter, to the equivalent uppercase letter and <i>vice versa</i> , as prescribed by the current locale. The current cursor position then shall be advanced by one character. If the cursor was positioned on the last character of the line, the case conversion shall occur, but the cursor shall not advance. If the '~' command is preceded by a <i>count</i> , that number of characters shall be converted, and the cursor shall be advanced to the character position after the last character converted. If the <i>count</i> is larger than the number of characters after the cursor, this shall not be considered an error; the cursor shall advance to the last character on the line.
33238 33239 33240 33241 33242 33243 33244	[count].	Repeat the most recent non-motion command, even if it was executed on an earlier command line. If the previous command was preceded by a <i>count</i> , and no count is given on the '.' command, the count from the previous command shall be included as part of the repeated command. If the '.' command is preceded by a <i>count</i> , this shall override any <i>count</i> argument to the previous command. The <i>count</i> specified in the '.' command shall become the count for subsequent '.' commands issued without a count.
33245 33246 33247 33248	[number]v	Invoke the <i>vi</i> editor to edit the current command line in a temporary file. When the editor exits, the commands in the temporary file shall be executed and placed in the command history. If a <i>number</i> is included, it specifies the command number in the command history to be edited, rather than the current command line.
33249 33250 33251 33252 33253 33254 33255	[count]l (e [count] <spac< td=""><td>ell) ce> Move the current cursor position to the next character position. If the cursor was positioned on the last character of the line, the terminal shall be alerted and the cursor shall not be advanced. If the <i>count</i> is larger than the number of characters after the cursor, this shall not be considered an error; the cursor shall advance to the last character on the line.</td></spac<>	ell) ce> Move the current cursor position to the next character position. If the cursor was positioned on the last character of the line, the terminal shall be alerted and the cursor shall not be advanced. If the <i>count</i> is larger than the number of characters after the cursor, this shall not be considered an error; the cursor shall advance to the last character on the line.
33256 33257 33258 33259 33260	[count]h	Move the current cursor position to the <i>count</i> th (default 1) previous character position. If the cursor was positioned on the first character of the line, the terminal shall be alerted and the cursor shall not be moved. If the count is larger than the number of characters before the cursor, this shall not be considered an error; the cursor shall move to the first character on the line.
33261 33262	[count]w	Move to the start of the next word. If the cursor was positioned on the last character of the line, the terminal shall be alerted and the cursor shall not be

sh Utilities

33263 33264 33265		advanced. If the <i>count</i> is larger than the number of words after the cursor, this shall not be considered an error; the cursor shall advance to the last character on the line.
33266 33267 33268 33269 33270	[count]W	Move to the start of the next bigword. If the cursor was positioned on the last character of the line, the terminal shall be alerted and the cursor shall not be advanced. If the <i>count</i> is larger than the number of bigwords after the cursor, this shall not be considered an error; the cursor shall advance to the last character on the line.
33271 33272 33273 33274 33275	[count]e	Move to the end of the current word. If at the end of a word, move to the end of the next word. If the cursor was positioned on the last character of the line, the terminal shall be alerted and the cursor shall not be advanced. If the <i>count</i> is larger than the number of words after the cursor, this shall not be considered an error; the cursor shall advance to the last character on the line.
33276 33277 33278 33279 33280	[count]E	Move to the end of the current bigword. If at the end of a bigword, move to the end of the next bigword. If the cursor was positioned on the last character of the line, the terminal shall be alerted and the cursor shall not be advanced. If the <i>count</i> is larger than the number of bigwords after the cursor, this shall not be considered an error; the cursor shall advance to the last character on the line.
33281 33282 33283 33284 33285 33286	[count] b	Move to the beginning of the current word. If at the beginning of a word, move to the beginning of the previous word. If the cursor was positioned on the first character of the line, the terminal shall be alerted and the cursor shall not be moved. If the <i>count</i> is larger than the number of words preceding the cursor, this shall not be considered an error; the cursor shall return to the first character on the line.
33287 33288 33289 33290 33291 33292	[count]B	Move to the beginning of the current bigword. If at the beginning of a bigword, move to the beginning of the previous bigword. If the cursor was positioned on the first character of the line, the terminal shall be alerted and the cursor shall not be moved. If the <i>count</i> is larger than the number of bigwords preceding the cursor, this shall not be considered an error; the cursor shall return to the first character on the line.
33293 33294	^	Move the current cursor position to the first character on the input line that is not a lank>.
33295	\$	Move to the last character position on the current command line.
33296	0	(Zero.) Move to the first character position on the current command line.
33297 33298 33299 33300 33301	[count]	Move to the <i>count</i> th character position on the current command line. If no number is specified, move to the first position. The first character position shall be numbered 1. If the count is larger than the number of characters on the line, this shall not be considered an error; the cursor shall be placed on the last character on the line.
33302 33303 33304 33305 33306	[count]fc	Move to the first occurrence of the character 'c' that occurs after the current cursor position. If the cursor was positioned on the last character of the line, the terminal shall be alerted and the cursor shall not be advanced. If the character 'c' does not occur in the line after the current cursor position, the terminal shall be alerted and the cursor shall not be moved.
33307 33308 33309	[count]Fc	Move to the first occurrence of the character 'c' that occurs before the current cursor position. If the cursor was positioned on the first character of the line, the terminal shall be alerted and the cursor shall not be moved. If the character 'c'

Utilities sh

33310 33311		does not occur in the line before the current cursor position, the terminal shall be alerted and the cursor shall not be moved.
33312 33313 33314 33315 33316	[count] t c	Move to the character before the first occurrence of the character 'c' that occurs after the current cursor position. If the cursor was positioned on the last character of the line, the terminal shall be alerted and the cursor shall not be advanced. If the character 'c' does not occur in the line after the current cursor position, the terminal shall be alerted and the cursor shall not be moved.
33317 33318 33319 33320 33321	[count]Tc	Move to the character after the first occurrence of the character 'c' that occurs before the current cursor position. If the cursor was positioned on the first character of the line, the terminal shall be alerted and the cursor shall not be moved. If the character 'c' does not occur in the line before the current cursor position, the terminal shall be alerted and the cursor shall not be moved.
33322 33323 33324	[count];	Repeat the most recent f , F , t , or T command. Any number argument on that previous command shall be ignored. Errors are those described for the repeated command.
33325 33326 33327	[count],	Repeat the most recent f , F , t , or T command. Any number argument on that previous command shall be ignored. However, reverse the direction of that command.
33328 33329	a	Enter insert mode after the current cursor position. Characters that are entered shall be inserted before the next character.
33330	A	Enter insert mode after the end of the current command line.
33331 33332	i	Enter insert mode at the current cursor position. Characters that are entered shall be inserted before the current character.
33333	I	Enter insert mode at the beginning of the current command line.
33334 33335	R	Enter insert mode, replacing characters from the command line beginning at the current cursor position.
33336 33337 33338 33339 33340 33341	[count]cmot	Delete the characters between the current cursor position and the cursor position that would result from the specified motion command. Then enter insert mode before the first character following any deleted characters. If <i>count</i> is specified, it shall be applied to the motion command. A <i>count</i> shall be ignored for the following motion commands:
33342		0 ^ \$ c
33343 33344 33345 33346 33347 33348 33349 33350 33351 33352 33353 33353		If the motion command is the character 'c', the current command line shall be cleared and insert mode shall be entered. If the motion command would move the current cursor position toward the beginning of the command line, the character under the current cursor position shall not be deleted. If the motion command would move the current cursor position toward the end of the command line, the character under the current cursor position shall be deleted. If the <i>count</i> is larger than the number of characters between the current cursor position and the end of the command line toward which the motion command would move the cursor, this shall not be considered an error; all of the remaining characters in the aforementioned range shall be deleted and insert mode shall be entered. If the motion command is invalid, the terminal shall be alerted, the cursor shall not be moved, and no text shall be deleted.

sh Utilities

 \mathbf{C} 33355 Delete from the current character to the end of the line and enter insert mode at the 33356 new end-of-line. S 33357 Clear the entire edit line and enter insert mode. [count]rc Replace the current character with the character 'c'. With a number count, 33358 33359 replace the current and the following *count*-1 characters. After this command, the current cursor position shall be on the last character that was changed. If the count 33360 is larger than the number of characters after the cursor, this shall not be considered 33361 an error; all of the remaining characters shall be changed. 33362 [count]_ Append a <space> after the current character position and then append the last 33363 bigword in the previous input line after the <space>. Then enter insert mode after 33364 the last character just appended. With a number count, append the countth 33365 bigword in the previous line. 33366 [count]x Delete the character at the current cursor position and place the deleted characters 33367 in the save buffer. If the cursor was positioned on the last character of the line, the 33368 character shall be deleted and the cursor position shall be moved to the previous 33369 character (the new last character). If the count is larger than the number of 33370 characters after the cursor, this shall not be considered an error; all the characters 33371 from the cursor to the end of the line shall be deleted. 33372 [count]X Delete the character before the current cursor position and place the deleted 33373 characters in the save buffer. The character under the current cursor position shall 33374 not change. If the cursor was positioned on the first character of the line, the 33375 terminal shall be alerted, and the X command shall have no effect. If the line 33376 contained a single character, the X command shall have no effect. If the line 33377 contained no characters, the terminal shall be alerted and the cursor shall not be 33378 moved. If the *count* is larger than the number of characters before the cursor, this 33379 shall not be considered an error; all the characters from before the cursor to the 33380 beginning of the line shall be deleted. 33381 33382 [count]dmotion Delete the characters between the current cursor position and the character 33383 position that would result from the motion command. A number *count* repeats the 33384 motion command count times. If the motion command would move toward the 33385 beginning of the command line, the character under the current cursor position 33386 shall not be deleted. If the motion command is d, the entire current command line 33387 shall be cleared. If the *count* is larger than the number of characters between the 33388 current cursor position and the end of the command line toward which the motion 33389 command would move the cursor, this shall not be considered an error; all of the 33390 remaining characters in the aforementioned range shall be deleted. The deleted 33391 characters shall be placed in the save buffer. 33392 D Delete all characters from the current cursor position to the end of the line. The 33393 33394 deleted characters shall be placed in the save buffer. [count]ymotion 33395 33396 Yank (that is, copy) the characters from the current cursor position to the position resulting from the motion command into the save buffer. A number *count* shall be 33397 applied to the motion command. If the motion command would move toward the 33398 beginning of the command line, the character under the current cursor position 33399 shall not be included in the set of yanked characters. If the motion command is y, 33400 the entire current command line shall be yanked into the save buffer. The current 33401 cursor position shall be unchanged. If the count is larger than the number of 33402

Utilities sh

33403 33404 33405 33406		characters between the current cursor position and the end of the command line toward which the motion command would move the cursor, this shall not be considered an error; all of the remaining characters in the aforementioned range shall be yanked.
33407 33408	Y	Yank the characters from the current cursor position to the end of the line into the save buffer. The current character position shall be unchanged.
33409 33410 33411 33412	[count]p	Put a copy of the current contents of the save buffer after the current cursor position. The current cursor position shall be advanced to the last character put from the save buffer. A <i>count</i> shall indicate how many copies of the save buffer shall be put.
33413 33414 33415 33416	[count]P	Put a copy of the current contents of the save buffer before the current cursor position. The current cursor position shall be moved to the last character put from the save buffer. A <i>count</i> shall indicate how many copies of the save buffer shall be put.
33417 33418	u	Undo the last command that changed the edit line. This operation shall not undo the copy of any command line to the edit line.
33419 33420	U	Undo all changes made to the edit line. This operation shall not undo the copy of any command line to the edit line.
33421 33422 33423 33424 33425 33426 33427	[count]k [count]–	Set the current command line to be the <i>count</i> th previous command line in the shell command history. If <i>count</i> is not specified, it shall default to 1. The cursor shall be positioned on the first character of the new command. If a \mathbf{k} or – command would retreat past the maximum number of commands in effect for this shell (affected by the <i>HISTSIZE</i> environment variable), the terminal shall be alerted, and the command shall have no effect.
33428 33429 33430 33431 33432 33433	[count] j [count]+	Set the current command line to be the <i>count</i> th next command line in the shell command history. If <i>count</i> is not specified, it shall default to 1. The cursor shall be positioned on the first character of the new command. If a \mathbf{j} or $+$ command advances past the edit line, the current command line shall be restored to the edit line and the terminal shall be alerted.
33434 33435 33436 33437	[number]G	Set the current command line to be the oldest command line stored in the shell command history. With a number <i>number</i> , set the current command line to be the command line <i>number</i> in the history. If command line <i>number</i> does not exist, the terminal shall be alerted and the command line shall not be changed.
33438 33439 33440 33441 33442 33443 33444 33445 33446 33447	/pattern <new< td=""><td>Move backwards through the command history, searching for the specified pattern, beginning with the previous command line. Patterns use the pattern matching notation described in Section 2.13 (on page 62), except that the '^' character shall have special meaning when it appears as the first character of pattern. In this case, the '^' is discarded and the characters after the '^' shall be matched only at the beginning of a line. Commands in the command history shall be treated as strings, not as filenames. If the pattern is not found, the current command line shall be unchanged and the terminal is alerted. If it is found in a previous line, the current command line shall be set to the first character of the new command line.</td></new<>	Move backwards through the command history, searching for the specified pattern, beginning with the previous command line. Patterns use the pattern matching notation described in Section 2.13 (on page 62), except that the '^' character shall have special meaning when it appears as the first character of pattern. In this case, the '^' is discarded and the characters after the '^' shall be matched only at the beginning of a line. Commands in the command history shall be treated as strings, not as filenames. If the pattern is not found, the current command line shall be unchanged and the terminal is alerted. If it is found in a previous line, the current command line shall be set to the first character of the new command line.

sh **Utilities**

33449		If pattern is empty, the last non-empty pattern provided to / or ? shall be used. If	
33450		there is no previous non-empty pattern, the terminal shall be alerted and the	
33451		current command line shall remain unchanged.	
33452	?pattern <newline></newline>		
33453	· paccorr	Move forwards through the command history, searching for the specified pattern,	
33454		beginning with the next command line. Patterns use the pattern matching notation	
33455		described in Section 2.13 (on page 62), except that the '^' character shall have	
33456		special meaning when it appears as the first character of <i>pattern</i> . In this case, the	
33457		'^' is discarded and the characters after the '^' shall be matched only at the	
33458		beginning of a line. Commands in the command history shall be treated as strings,	
33459		not as filenames. If the pattern is not found, the current command line shall be	
33460		unchanged and the terminal alerted. If it is found in a following line, the current	
33461		command line shall be set to that line and the cursor shall be set to the fist	
33462		character of the new command line.	
33463		If pattern is empty, the last non-empty pattern provided to / or ? shall be used. If	
33464		there is no previous non-empty pattern, the terminal shall be alerted and the	
33465		current command line shall remain unchanged.	
33466 33467	n	Repeat the most recent / or ? command. If there is no previous / or ?, the terminal shall be alerted and the current command line shall remain unchanged.	
00400	N		
33468	11	Repeat the most recent / or ? command, reversing the direction of the search. If	
33469 33470		there is no previous / or ?, the terminal shall be alerted and the current command line shall remain unchanged.	
33471 EXIT S	TATUS		
33472	The foll	owing exit values shall be returned:	
33473	0	The script to be executed consisted solely of zero or more blank lines or comments, or	
33474		both.	
33475	1-125	A non-interactive shell detected a syntax, redirection, or variable assignment error.	
33476	127	A specified <i>command_file</i> could not be found by a non-interactive shell.	
33477	Otherwise, the shell shall return the exit status of the last command it invoked or attempted to		
33478	invoke (see also the <i>exit</i> utility in Section 2.14 (on page 64)).		
33479 CONSI	EQUENC	ES OF ERRORS	
33480			
33481 APPLI	CATION	USAGE	

33481 APPLICATION USAGE

Standard input and standard error are the files that determine whether a shell is interactive 33482 when -i is not specified. For example: 33483

33484 sh > file

33485 and:

sh 2> file 33486

create interactive and non-interactive shells, respectively. Although both accept terminal input, 33487 33488 the results of error conditions are different, as described in Section 2.8.1 (on page 46); in the 33489 second example a redirection error encountered by a special built-in utility aborts the shell.

33490 A conforming application must protect its first operand, if it starts with a plus sign, by preceding 33491 it with the "--" argument that denotes the end of the options.

Utilities sh

33492 Applications should note that the standard *PATH* to the shell cannot be assumed to be either /bin/sh or /usr/bin/sh, and should be determined by interrogation of the PATH returned by 33493 getconf PATH, ensuring that the returned pathname is an absolute pathname and not a shell 33494 built-in. 33495 33496 For example, to determine the location of the standard *sh* utility: command -v sh 33497 33498 On some implementations this might return: 33499 /usr/xpg4/bin/sh Furthermore, on systems that support executable scripts (the "#!" construct), it is 33500 recommended that applications using executable scripts install them using getconf -v to 33501 determine the shell pathname and update the "#!" script appropriately as it is being installed 33502 (for example, with *sed*). For example: 33503 33504 # Installation time script to install correct POSIX shell pathname 33505 33506 # # Get list of paths to check 33507 33508 # Sifs=\$IFS 33509 33510 IFS=: 33511 set \$ (getconf PATH) IFS=\$Sifs 33512 33513 # Check each path for 'sh' 33514 33515 33516 for i in \$@ 33517 if $[-f ${i}/sh];$ 33518 33519 then Pshell=\${i}/sh 33520 fi 33521 33522 done 33523 # This is the list of scripts to update. They should be of the 33524 # form '\${name}.source' and will be transformed to '\${name}'. 33525 # Each script should begin: 33526 33527 # !INSTALLSHELLPATH -p 33528 33529 scripts="a b c" 33530 33531 33532 # Transform each script 33533 33534 for i in \${scripts} 33535 sed -e "s|INSTALLSHELLPATH|\${Pshell}|" < \${i}.source > \${i} 33536

33537

done

sh Utilities

33538 EXAMPLES

33539 1. Execute a shell command from a string:

33540 sh -c "cat myfile"

33541 2. Execute a shell script from a file in the current directory:

sh my_shell_cmds

33543 RATIONALE

The *sh* utility and the *set* special built-in utility share a common set of options.

The KornShell ignores the contents of *IFS* upon entry to the script. A conforming application cannot rely on importing *IFS*. One justification for this, beyond security considerations, is to assist possible future shell compilers. Allowing *IFS* to be imported from the environment prevents many optimizations that might otherwise be performed via dataflow analysis of the script itself.

The text in the STDIN section about non-blocking reads concerns an instance of *sh* that has been invoked, probably by a C-language program, with standard input that has been opened using the O_NONBLOCK flag; see *open()* in the System Interfaces volume of IEEE Std 1003.1-2001. If the shell did not reset this flag, it would immediately terminate because no input data would be available yet and that would be considered the same as end-of-file.

The options associated with a *restricted shell* (command name rsh and the $-\mathbf{r}$ option) were excluded because the standard developers considered that the implied level of security could not be achieved and they did not want to raise false expectations.

On systems that support set-user-ID scripts, a historical trapdoor has been to link a script to the name $-\mathbf{i}$. When it is called by a sequence such as:

33560 sh -33561 or by:

33562 #! usr/bin/sh -

the historical systems have assumed that no option letters follow. Thus, this volume of IEEE Std 1003.1-2001 allows the single hyphen to mark the end of the options, in addition to the use of the regular "--" argument, because it was considered that the older practice was so pervasive. An alternative approach is taken by the KornShell, where real and effective user/group IDs must match for an interactive shell; this behavior is specifically allowed by this volume of IEEE Std 1003.1-2001.

Note: There are other problems with set-user-ID scripts that the two approaches described here do not resolve.

The initialization process for the history file can be dependent on the system start-up files, in that they may contain commands that effectively preempt the user's settings of *HISTFILE* and *HISTSIZE*. For example, function definition commands are recorded in the history file, unless the *set* –**o** *nolog* option is set. If the system administrator includes function definitions in some system start-up file called before the *ENV* file, the history file is initialized before the user gets a chance to influence its characteristics. In some historical shells, the history file is initialized just after the *ENV* file has been processed. Therefore, it is implementation-defined whether changes made to *HISTFILE* after the history file has been initialized are effective.

The default messages for the various *MAIL*-related messages are unspecified because they vary across implementations. Typical messages are:

Utilities sh

33581 "you have mail\n"
33582 or:
33583 "you have new mail\n"

33584

33585

33586

33587

33588

33590

33591 33592

33593 33594

33595

33596

33597

33598

33599

33600 33601

33602

33603

33604 33605

33606

33607

33608

33609

33610

33611

33612

33613 33614

33615

33616

It is important that the descriptions of command line editing refer to the same shell as that in IEEE Std 1003.1-2001 so that interactive users can also be application programmers without having to deal with programmatic differences in their two environments. It is also essential that the utility name *sh* be specified because this explicit utility name is too firmly rooted in historical practice of application programs for it to change.

Consideration was given to mandating a diagnostic message when attempting to set *vi*-mode on terminals that do not support command line editing. However, it is not historical practice for the shell to be cognizant of all terminal types and thus be able to detect inappropriate terminals in all cases. Implementations are encouraged to supply diagnostics in this case whenever possible, rather than leaving the user in a state where editing commands work incorrectly.

In early proposals, the KornShell-derived *emacs* mode of command line editing was included, even though the emacs editor itself was not. The community of emacs proponents was adamant that the full emacs editor not be standardized because they were concerned that an attempt to standardize this very powerful environment would encourage vendors to ship strictly conforming versions lacking the extensibility required by the community. The author of the original *emacs* program also expressed his desire to omit the program. Furthermore, there were a number of historical systems that did not include emacs, or included it without supporting it, but there were very few that did not include and support vi. The shell emacs command line editing mode was finally omitted because it became apparent that the KornShell version and the editor being distributed with the GNU system had diverged in some respects. The author of emacs requested that the POSIX emacs mode either be deleted or have a significant number of unspecified conditions. Although the KornShell author agreed to consider changes to bring the shell into alignment, the standard developers decided to defer specification at that time. At the time, it was assumed that convergence on an acceptable definition would occur for a subsequent draft, but that has not happened, and there appears to be no impetus to do so. In any case, implementations are free to offer additional command line editing modes based on the exact models of editors their users are most comfortable with.

Early proposals had the following list entry in vi Line Editing Insert Mode (on page 855):

\ If followed by the *erase* or *kill* character, that character shall be inserted into the input line. Otherwise, the backslash itself shall be inserted into the input line.

However, this is not actually a feature of *sh* command line editing insert mode, but one of some historical terminal line drivers. Some conforming implementations continue to do this when the *stty* **iexten** flag is set.

33617 FUTURE DIRECTIONS

33618 None.

33619 **SEE ALSO**

Chapter 2 (on page 29), cd, echo, exit, fc, pwd, read, set, stty, test, umask, vi, the System Interfaces volume of IEEE Std 1003.1-2001, dup(), exec, exit(), fork(), open(), pipe(), signal(), system(), ulimit(), umask(), wait()

33623 CHANGE HISTORY

First released in Issue 2.

sh Utilities

33625 Issue 5	
33626	The FUTURE DIRECTIONS section is added.
33627	Text is added to the DESCRIPTION for the Large File Summit proposal.
33628 Issue 6	
33629	The Open Group Corrigendum U029/2 is applied, correcting the second SYNOPSIS.
33630	The Open Group Corrigendum $U027/3$ is applied, correcting a typographical error.
33631 33632	The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:
33633 33634	• The option letters derived from the set special built-in are also accepted with a leading plus sign ('+').
33635	Large file extensions are added:
33636	 Pathname expansion does not fail due to the size of a file.
33637 33638	 Shell input and output redirections have an implementation-defined offset maximum that is established in the open file description.
33639 33640	In the ENVIRONMENT VARIABLES section, the text "user's home directory" is updated to "directory referred to by the <i>HOME</i> environment variable".
33641 33642	Descriptions for the $\it ENV$ and $\it PWD$ environment variables are included to align with the IEEE P1003.2b draft standard.
33643	The normative text is reworded to avoid use of the term "must" for application requirements.

Utilities sleep

33644 **NAME** 33645 sleep — suspend execution for an interval 33646 SYNOPSIS 33647 sleep time 33648 **DESCRIPTION** The *sleep* utility shall suspend execution for at least the integral number of seconds specified by 33649 the *time* operand. 33650 33651 OPTIONS 33652 None. 33653 OPERANDS 33654 The following operand shall be supported: time A non-negative decimal integer specifying the number of seconds for which to 33655 33656 suspend execution. 33657 **STDIN** Not used. 33658 33659 INPUT FILES 33660 None. 33661 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *sleep*: 33662 LANG Provide a default value for the internationalization variables that are unset or null. 33663 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 33664 Internationalization Variables for the precedence of internationalization variables 33665 used to determine the values of locale categories.) 33666 LC ALL If set to a non-empty string value, override the values of all the other 33667 internationalization variables. 33668 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 33669 33670 characters (for example, single-byte as opposed to multi-byte characters in 33671 arguments). 33672 LC_MESSAGES 33673 Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. 33674 33675 XSI NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 33676 ASYNCHRONOUS EVENTS 33677 If the *sleep* utility receives a SIGALRM signal, one of the following actions shall be taken: Terminate normally with a zero exit status. 33678

- 33679 2. Effectively ignore the signal.
- 3. Provide the default behavior for signals described in the ASYNCHRONOUS EVENTS section of Section 1.11 (on page 20). This could include terminating with a non-zero exit status.
- 33683 The *sleep* utility shall take the standard action for all other signals.

sleep Utilities

```
33684 STDOUT
33685
             Not used.
33686 STDERR
33687
             The standard error shall be used only for diagnostic messages.
33688 OUTPUT FILES
             None.
33689
33690 EXTENDED DESCRIPTION
             None.
33691
33692 EXIT STATUS
             The following exit values shall be returned:
33693
                  The execution was successfully suspended for at least time seconds, or a SIGALRM signal
33694
                  was received. See the ASYNCHRONOUS EVENTS section.
33695
33696
             >0 An error occurred.
33697 CONSEQUENCES OF ERRORS
             Default.
33698
33699 APPLICATION USAGE
33700
             None.
33701 EXAMPLES
33702
             The sleep utility can be used to execute a command after a certain amount of time, as in:
33703
              (sleep 105; command) &
             or to execute a command every so often, as in:
33704
33705
             while true
33706
              do
33707
                   command
33708
                   sleep 37
33709
             done
33710 RATIONALE
33711
             The exit status is allowed to be zero when sleep is interrupted by the SIGALRM signal because
             most implementations of this utility rely on the arrival of that signal to notify them that the
33712
             requested finishing time has been successfully attained. Such implementations thus do not
33713
             distinguish this situation from the successful completion case. Other implementations are
33714
33715
             allowed to catch the signal and go back to sleep until the requested time expires or to provide
             the normal signal termination procedures.
33716
33717
             As with all other utilities that take integral operands and do not specify subranges of allowed
             values, sleep is required by this volume of IEEE Std 1003.1-2001 to deal with time requests of up
33718
             to 2 147 483 647 seconds. This may mean that some implementations have to make multiple calls
33719
             to the delay mechanism of the underlying operating system if its argument range is less than
33720
             this.
33721
33722 FUTURE DIRECTIONS
33723
             None.
33724 SEE ALSO
```

wait, the System Interfaces volume of IEEE Std 1003.1-2001, alarm(), sleep()

33725

Utilities sleep

33726 CHANGE HISTORY

First released in Issue 2.

sort Utilities

```
33728 NAME
33729
              sort — sort, merge, or sequence check text files
              sort [-m] [-o output] [-bdfinru] [-t char] [-k keydef] ... [file...]
33731
33732
              sort -c [-bdfinru] [-t char] [-k keydef] [file]
33733 DESCRIPTION
33734
              The sort utility shall perform one of the following functions:
33735
               1. Sort lines of all the named files together and write the result to the specified output.
               2. Merge lines of all the named (presorted) files together and write the result to the specified
33736
33737
                   output.
               3. Check that a single input file is correctly presorted.
33738
              Comparisons shall be based on one or more sort keys extracted from each line of input (or, if no
33739
33740
              sort keys are specified, the entire line up to, but not including, the terminating <newline>), and
              shall be performed using the collating sequence of the current locale.
33741
33742 OPTIONS
              The sort utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
33743
              12.2, Utility Syntax Guidelines, and the -\mathbf{k} keydef option should follow the -\mathbf{b}, -\mathbf{d}, -\mathbf{f}, -\mathbf{i}, -\mathbf{n}, and
33744
33745
              -r options.
33746
              The following options shall be supported:
                           Check that the single input file is ordered as specified by the arguments and the
33747
33748
                           collating sequence of the current locale. No output shall be produced; only the exit
                           code shall be affected.
33749
33750
              -m
                           Merge only; the input file shall be assumed to be already sorted.
                           Specify the name of an output file to be used instead of the standard output. This
33751
              -o output
33752
                           file can be the same as one of the input files.
33753
              -11
                           Unique: suppress all but one in each set of lines having equal keys. If used with
33754
                           the -c option, check that there are no lines with duplicate keys, in addition to
                           checking that the input file is sorted.
33755
33756
              The following options shall override the default ordering rules. When ordering options appear
              independent of any key field specifications, the requested field ordering rules shall be applied
33757
              globally to all sort keys. When attached to a specific key (see -k), the specified ordering options
33758
              shall override all global ordering options for that key.
33759
              -\mathbf{d}
                           33760
                           setting of LC_CTYPE, shall be significant in comparisons. The behavior is
33761
                           undefined for a sort key to which -\mathbf{i} or -\mathbf{n} also applies.
33762
              -\mathbf{f}
                           Consider all lowercase characters that have uppercase equivalents, according to
33763
                           the current setting of LC_CTYPE, to be the uppercase equivalent for the purposes
33764
33765
                           of comparison.
                           Ignore all characters that are non-printable, according to the current setting of
33766
              -i
                           LC_CTYPE.
33767
                           33768
              -n
```

33769

33770

optional minus sign, and zero or more digits with an optional radix character and thousands separators (as defined in the current locale), which shall be sorted by

Utilities sort

33771 arithmetic value. An empty digit string shall be treated as zero. Leading zeros and 33772 signs on zeros shall not affect ordering. 33773 $-\mathbf{r}$ Reverse the sense of comparisons. 33774 The treatment of field separators can be altered using the options: -b Ignore leading
blank>s when determining the starting and ending positions of a 33775 restricted sort key. If the $-\mathbf{b}$ option is specified before the first $-\mathbf{k}$ option, it shall be 33776 33777 applied to all $-\mathbf{k}$ options. Otherwise, the $-\mathbf{b}$ option can be attached independently to each **-k** *field start* or *field end* option-argument (see below). 33778 33779 -t char Use *char* as the field separator character; *char* shall not be considered to be part of a field (although it can be included in a sort key). Each occurrence of char shall be 33780 significant (for example, <char><char> delimits an empty field). If -t is not 33781 specified, <blank>s shall be used as default field separators; each maximal non-33782 empty sequence of <blank>s that follows a non-
blank> shall be a field separator. 33783 Sort keys can be specified using the options: 33784 −**k** keydef The keydef argument is a restricted sort key field definition. The format of this 33785 definition is: 33786 field start[type][,field end[type]] 33787 where field_start and field_end define a key field restricted to a portion of the line 33788 (see the EXTENDED DESCRIPTION section), and type is a modifier from the list of 33789 characters 'b', 'd', 'f', 'i', 'n', 'r'. The 'b' modifier shall behave like the 33790 **−b** option, but shall apply only to the *field_start* or *field_end* to which it is attached. 33791 The other modifiers shall behave like the corresponding options, but shall apply 33792 only to the key field to which they are attached; they shall have this effect if 33793 specified with field_start, field_end, or both. If any modifier is attached to a 33794 field_start or to a field_end, no option shall apply to either. Implementations shall 33795 support at least nine occurrences of the -k option, which shall be significant in 33796 33797 command line order. If no -k option is specified, a default sort key of the entire line shall be used. 33798 When there are multiple key fields, later keys shall be compared only after all 33799 earlier keys compare equal. Except when the -u option is specified, lines that 33800 otherwise compare equal shall be ordered as if none of the options $-\mathbf{d}$, $-\mathbf{f}$, $-\mathbf{i}$, $-\mathbf{n}$, or 33801 -k were present (but with -r still in effect, if it was specified) and with all bytes in 33802 the lines significant to the comparison. The order in which lines that still compare 33803 33804 equal are written is unspecified. 33805 OPERANDS The following operand shall be supported: 33806 file A pathname of a file to be sorted, merged, or checked. If no file operands are 33807 specified, or if a *file* operand is '-', the standard input shall be used. 33808 33809 **STDIN** The standard input shall be used only if no *file* operands are specified, or if a *file* operand is '-'. 33810 See the INPUT FILES section. 33811 33812 INPUT FILES The input files shall be text files, except that the *sort* utility shall add a <newline> to the end of a 33813 file ending with an incomplete last line. 33814

sort Utilities

	33815 ENVIRONMENT VARIABLES			
33816		ng environment variables shall affect the execution of <i>sort</i> :		
33817 33818 33819 33820	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)		
33821 33822	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.		
33823 33824	LC_COLLAT	TE Determine the locale for ordering rules.		
33825 33826 33827 33828	LC_CTYPE	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) and the behavior of character classification for the $-\mathbf{b}$, $-\mathbf{d}$, $-\mathbf{f}$, $-\mathbf{i}$, and $-\mathbf{n}$ options.		
33829	LC_MESSA	GES		
33830 33831		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.		
33832	LC_NUMER	PIC		
33833 33834		Determine the locale for the definition of the radix character and thousands separator for the $-\mathbf{n}$ option.		
33835 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.		
33836 ASYNO 33837	CHRONOUS I Default.	EVENTS		
33838 STDOU	J T			
33839	Unless the $-\mathbf{o}$ or $-\mathbf{c}$ options are in effect, the standard output shall contain the sorted input.			
33840 STDER				
33841 33842	The standard error shall be used for diagnostic messages. A warning message about correcting an incomplete last line of an input file may be generated, but need not affect the final exit status.			
33843 OUTPU 33844		ion is in effect, the sorted input shall be written to the file <i>output</i> .		
33845 EXTEN 33846	DED DESCR The notation			
33847	-k field_	start[type][,field_end[type]]		
33848 33849 33850	falls beyond	a key field that begins at <i>field_start</i> and ends at <i>field_end</i> inclusive, unless <i>field_start</i> the end of the line or after <i>field_end</i> , in which case the key field is empty. A missing all mean the last character of the line.		
33851 33852		prises a maximal sequence of non-separating characters and, in the absence of option eding field separator.		
33853	The field_sta	The <i>field_start</i> portion of the <i>keydef</i> option-argument shall have the form:		
33854	field num	field number[.first character]		
33855 33856 33857	Fields and of first_characters	characters within fields shall be numbered starting with 1. The <i>field_number</i> and <i>er</i> pieces, interpreted as positive decimal integers, shall specify the first character to eart of a sort key. If <i>.first_character</i> is omitted, it shall refer to the first character of the		

Utilities sort

33858 field.

33859 The *field_end* portion of the *keydef* option-argument shall have the form:

33860 field number[.last character]

The *field_number* shall be as described above for *field_start*. The *last_character* piece, interpreted as a non-negative decimal integer, shall specify the last character to be used as part of the sort key. If *last_character* evaluates to zero or *.last_character* is omitted, it shall refer to the last character of the field specified by *field_number*.

If the **-b** option or **b** type modifier is in effect, characters within a field shall be counted from the first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

first non-

firs

33867 EXIT STATUS

33871 33872

33868 The following exit values shall be returned:

33869 0 All input files were output successfully, or **-c** was specified and the input file was correctly sorted.

1 Under the −c option, the file was not ordered as specified, or if the −c and −u options were both specified, two input lines were found with equal keys.

33873 >1 An error occurred.

33874 CONSEQUENCES OF ERRORS

33875 Default.

33876 APPLICATION USAGE

The default value for -t, <blank>, has different properties from, for example, -t"<space>". If a line contains:

33879 <space><space>foo

the following treatment would occur with default separation as opposed to specifically selecting a <space>:

33882	Field	Default	-t " <space>"</space>
33883	1	<space><space>foo</space></space>	empty
33884	2	empty	empty
33885	3	empty	foo

The leading field separator itself is included in a field when -t is not used. For example, this command returns an exit status of zero, meaning the input was already sorted:

```
33888 sort -c -k 2 <<eof
```

33889 y<tab>b

33890 x<space>a

33891 eo

33886 33887

33892 (assuming that a <tab> precedes the <space> in the current collating sequence). The field separator is not included in a field when it is explicitly set via -t. This is historical practice and allows usage such as:

```
33895 sort -t "|" -k 2n <<eof

33896 Atlanta|425022|Georgia

33897 Birmingham|284413|Alabama

33898 Columbia|100385|South Carolina
```

33899 eof

sort Utilities

where the second field can be correctly sorted numerically without regard to the non-numeric field separator.

The wording in the OPTIONS section clarifies that the $-\mathbf{b}$, $-\mathbf{d}$, $-\mathbf{f}$, $-\mathbf{n}$, and $-\mathbf{r}$ options have to come before the first sort key specified if they are intended to apply to all specified keys. The way it is described in this volume of IEEE Std 1003.1-2001 matches historical practice, not historical documentation. The results are unspecified if these options are specified after a $-\mathbf{k}$ option.

The **–f** option might not work as expected in locales where there is not a one-to-one mapping between an uppercase and a lowercase letter.

33909 EXAMPLES

1. The following command sorts the contents of **infile** with the second field as the sort key:

```
33911 sort -k 2,2 infile
```

2. The following command sorts, in reverse order, the contents of **infile1** and **infile2**, placing the output in **outfile** and using the second character of the second field as the sort key (assuming that the first character of the second field is the field separator):

```
sort -r -o outfile -k 2.2,2.2 infile1 infile2
```

```
sort -k 2.2b, 2.2b infile1 infile2
```

4. The following command prints the System V password file (user database) sorted by the numeric user ID (the third colon-separated field):

```
sort -t : -k 3,3n /etc/passwd
```

5. The following command prints the lines of the already sorted file **infile**, suppressing all but one occurrence of lines having the same third field:

```
33924 sort -um -k 3.1,3.0 infile
```

33925 RATIONALE

Examples in some historical documentation state that options —**um** with one input file keep the first in each set of lines with equal keys. This behavior was deemed to be an implementation artifact and was not standardized.

The -z option was omitted; it is not standard practice on most systems and is inconsistent with using *sort* to sort several files individually and then merge them together. The text concerning -z in historical documentation appeared to require implementations to determine the proper buffer length during the sort phase of operation, but not during the merge.

The -y option was omitted because of non-portability. The -M option, present in System V, was omitted because of non-portability in international usage.

An undocumented -T option exists in some implementations. It is used to specify a directory for intermediate files. Implementations are encouraged to support the use of the TMPDIR environment variable instead of adding an option to support this functionality.

The -k option was added to satisfy two objections. First, the zero-based counting used by *sort* is not consistent with other utility conventions. Second, it did not meet syntax guideline requirements.

Historical documentation indicates that "setting $-\mathbf{n}$ implies $-\mathbf{b}$ ". The description of $-\mathbf{n}$ already states that optional leading
 states that optional leading
 states that optional leading
 tolerated in doing the comparison. If $-\mathbf{b}$ is enabled,

Utilities sort

33943 rather than implied, by -n, this has unusual side effects. When a character offset is used in a 33944 column of numbers (for example, to sort modulo 100), that offset is measured relative to the most significant digit, not to the column. Based upon a recommendation from the author of the 33945 original sort utility, the -b implication has been omitted from this volume of 33946 IEEE Std 1003.1-2001, and an application wishing to achieve the previously mentioned side 33947 33948 effects has to code the $-\mathbf{b}$ flag explicitly. 33949 FUTURE DIRECTIONS None. 33950 33951 **SEE ALSO** 33952 comm, join, uniq, the System Interfaces volume of IEEE Std 1003.1-2001, toupper() 33953 CHANGE HISTORY 33954 First released in Issue 2. 33955 **Issue 6** 33956 IEEE PASC Interpretation 1003.2 #174 is applied, updating the DESCRIPTION of comparisons. 33957 IEEE PASC Interpretation 1003.2 #168 is applied.

split Utilities

```
33958 NAME
```

33959 split — split files into pieces

33960 SYNOPSIS

```
split [-l line_count] [-a suffix_length] [file[name]]
```

split -b n[k|m] [-a suffix length] [file[name]]

33963

33965

33966

33967

33968

33969

33970

33971

33972

33973

33974

33975

33976

33977

33978

33982

33984

33985 33986

33987

33988

33991

33992

33993

33994

34002

33964 **DESCRIPTION**

The *split* utility shall read an input file and write one or more output files. The default size of each output file shall be 1 000 lines. The size of the output files can be modified by specification of the $-\mathbf{b}$ or $-\mathbf{l}$ options. Each output file shall be created with a unique suffix. The suffix shall consist of exactly $suffix_length$ lowercase letters from the POSIX locale. The letters of the suffix shall be used as if they were a base-26 digit system, with the first suffix to be created consisting of all 'a' characters, the second with a 'b' replacing the last 'a', and so on, until a name of all 'z' characters is created. By default, the names of the output files shall be 'x', followed by a two-character suffix from the character set as described above, starting with "aa", "ab", "ac", and so on, and continuing until the suffix "zz", for a maximum of 676 files.

If the number of files required exceeds the maximum allowed by the suffix length provided, such that the last allowable file would be larger than the requested size, the *split* utility shall fail after creating the last file with a valid suffix; *split* shall not delete the files it created with valid suffixes. If the file limit is not exceeded, the last file created shall contain the remainder of the input file, and may be smaller than the requested size.

33979 OPTIONS

The *split* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported:

33983 —a suffix_length

Use *suffix_length* letters to form the suffix portion of the filenames of the split file. If **-a** is not specified, the default suffix length shall be two. If the sum of the *name* operand and the *suffix_length* option-argument would create a filename exceeding {NAME_MAX} bytes, an error shall result; *split* shall exit with a diagnostic message and no files shall be created.

33989 — $\mathbf{b} n$ Split a file into pieces n bytes in size.

33990 —**b** $n\mathbf{k}$ Split a file into pieces $n*1\,024$ bytes in size.

−b *n***m** Split a file into pieces *n**1 048 576 bytes in size.

-l line_count Specify the number of lines in each resulting file piece. The line_count argument is an unsigned decimal integer. The default is 1 000. If the input does not end with a <newline>, the partial line shall be included in the last output file.

33995 OPERANDS

33996 The following operands shall be supported:

33997 *file* The pathname of the ordinary file to be split. If no input file is given or *file* is '-', the standard input shall be used.

The prefix to be used for each of the files resulting from the split operation. If no name argument is given, 'x' shall be used as the prefix of the output files. The combined length of the basename of prefix and suffix_length cannot exceed

{NAME_MAX} bytes. See the OPTIONS section.

Utilities split

34003 **STDIN** 34004 See the INPUT FILES section. 34005 INPUT FILES Any file can be used as input. 34006 34007 ENVIRONMENT VARIABLES 34008 The following environment variables shall affect the execution of *split*: **LANG** 34009 Provide a default value for the internationalization variables that are unset or null. 34010 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables 34011 34012 used to determine the values of locale categories.) LC ALL If set to a non-empty string value, override the values of all the other 34013 internationalization variables. 34014 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 34015 characters (for example, single-byte as opposed to multi-byte characters in 34016 arguments and input files). 34017 LC_MESSAGES 34018 34019 Determine the locale that should be used to affect the format and contents of 34020 diagnostic messages written to standard error. 34021 XSI NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 34022 ASYNCHRONOUS EVENTS Default. 34023 34024 STDOUT Not used. 34025 34026 STDERR 34027 The standard error shall be used only for diagnostic messages. 34028 OUTPUT FILES The output files contain portions of the original input file; otherwise, unchanged. 34029 34030 EXTENDED DESCRIPTION 34031 None. 34032 EXIT STATUS The following exit values shall be returned: 34033 34034 Successful completion. 34035 >0 An error occurred. 34036 CONSEQUENCES OF ERRORS

Default.

34037

split Utilities

34038 APPLICATION USAGE

34039 None.

34040 EXAMPLES

In the following examples **foo** is a text file that contains 5 000 lines.

1. Create five files, xaa, xab, xac, xad, and xae:

```
34043 split foo
```

2. Create five files, but the suffixed portion of the created files consists of three letters, **xaaa**, **xaab**, **xaac**, **xaad**, and **xaae**:

```
34046 split -a 3 foo
```

3. Create three files with four-letter suffixes and a supplied prefix, **bar_aaaa**, **bar_aaab**, and **bar_aaaa**:

```
34049 split -a 4 -l 2000 foo bar_
```

4. Create as many files as are necessary to contain at most 20*1 024 bytes, each with the default prefix of **x** and a five-letter suffix:

```
split —a 5 —b 20k foo
```

34053 RATIONALE

34047

34048

34050

34051 34052

34054

34055

34056

34058

34059

34060

34061 34062

34064

34065

34066

The $-\mathbf{b}$ option was added to provide a mechanism for splitting files other than by lines. While most uses of the $-\mathbf{b}$ option are for transmitting files over networks, some believed it would have additional uses.

 $\frac{1}{2}$ The -a option was added to overcome the limitation of being able to create only 676 files.

Consideration was given to deleting this utility, using the rationale that the functionality provided by this utility is available via the *csplit* utility (see *csplit*). Upon reconsideration of the purpose of the User Portability Extension, it was decided to retain both this utility and the *csplit* utility because users use both utilities and have historical expectations of their behavior. Furthermore, the splitting on byte boundaries in *split* cannot be duplicated with the historical *csplit*

34063 *csplit*.

The text "split shall not delete the files it created with valid suffixes" would normally be assumed, but since the related utility, *csplit*, does delete files under some circumstances, the historical behavior of *split* is made explicit to avoid misinterpretation.

34067 FUTURE DIRECTIONS

34068 None.

34069 SEE ALSO

csplit 34070 *csplit*

34071 CHANGE HISTORY

First released in Issue 2.

34073 Issue 6

This utility is marked as part of the User Portability Utilities option.

34075 The APPLICATION USAGE section is added.

The obsolescent SYNOPSIS is removed.

strings **Utilities**

34077 **NAME**

34078 strings — find printable strings in files

34079 SYNOPSIS

strings [-a][-t format][-n number][file...] 34080 UP

34081

34082 DESCRIPTION

The strings utility shall look for printable strings in regular files and shall write those strings to 34083 standard output. A printable string is any sequence of four (by default) or more printable 34084 characters terminated by a <newline> or NUL character. Additional implementation-defined 34085 34086 strings may be written; see *localedef*.

34087 OPTIONS

34088 The strings utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines. 34089

The following options shall be supported: 34090

Scan files in their entirety. If -a is not specified, it is implementation-defined what 34091 34092 portion of each file is scanned for strings.

-n number Specify the minimum string length, where the *number* argument is a positive 34093 decimal integer. The default shall be 4. 34094

-t format Write each string preceded by its byte offset from the start of the file. The format 34095 34096 shall be dependent on the single character used as the *format* option-argument:

The offset shall be written in decimal. 34097 d

The offset shall be written in octal. 34098 0

34099 The offset shall be written in hexadecimal. v

34100 OPERANDS

The following operand shall be supported: 34101

file A pathname of a regular file to be used as input. If no *file* operand is specified, the 34102 34103 strings utility shall read from the standard input.

34104 **STDIN**

See the INPUT FILES section. 34105

34106 INPUT FILES

34107 The input files named by the utility arguments or the standard input shall be regular files of any 34108 format.

34109 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *strings*: 34110 LANG Provide a default value for the internationalization variables that are unset or null. 34111

(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 34119 Internationalization Variables for the precedence of internationalization variables 34113 34114 used to determine the values of locale categories.)

LC ALL If set to a non-empty string value, override the values of all the other 34115 34116 internationalization variables.

34117 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 34118 characters (for example, single-byte as opposed to multi-byte characters in 34119

arguments and input files) and to identify printable strings.

strings Utilities

34120	LC_MESSAGES	
34121 34122	Determine the locale that should be used t diagnostic messages written to standard error	
34123 XSI	NLSPATH Determine the location of message catalogs for	r the processing of <i>LC_MESSAGES</i> .
34124 ASYNC	CHRONOUS EVENTS	
34125	Default.	
34126 STDOU 34127	J T Strings found shall be written to the standard output, one p	er line
34128	When the –t option is not specified, the format of the output	
34129	"%s", <string></string>	t shan be.
34130	With the $-\mathbf{t}$ \mathbf{o} option, the format of the output shall be:	
34130	"%o %s", <byte offset="">, <string></string></byte>	
34132	With the – t x option, the format of the output shall be:	
34133	"%x %s", <byte offset="">, <string></string></byte>	
34134	With the -t d option, the format of the output shall be:	
34135	"%d %s", <byte offset="">, <string></string></byte>	
34136 STDER 34137	${f R}$ The standard error shall be used only for diagnostic messag	jes.
34138 OUTPU		
34139	None.	
34140 EXTEN 34141	DED DESCRIPTION None.	
34142 EXIT S	TATUS	
34143	The following exit values shall be returned:	
34144	0 Successful completion.	
34145	>0 An error occurred.	
34146 CONSE 34147	QUENCES OF ERRORS Default.	
34148 APPLIC	CATION USAGE	
34149 34150	By default the data area (as opposed to the text, "bss", or file is scanned. Implementations document which areas are	
34151 34152	Some historical implementations do not require NUL or permit those languages that do not use NUL as a string term	
34153 EXAMI 34154	PLES None.	
34155 RATIO	NALE	
34156 34157 34158	Apart from rationalizing the option syntax and slight dibinary files, <i>strings</i> is specified to match historical practice introduced to replace the non-conforming – and – <i>number</i> options.	closely. The -a and -n options were
34159 34160	The $-\mathbf{o}$ option historically means different things on different mean "offset in decimal", while others use it as "offset in occ	

Utilities strings

34161 way would be least objectionable, the -t option was added. It was originally named -O to mean 34162 "offset", but was changed to -t to be consistent with od. The ISO C standard function *isprint()* is restricted to a domain of **unsigned char**. This volume of 34163 IEEE Std 1003.1-2001 requires implementations to write strings as defined by the current locale. 34164 34165 FUTURE DIRECTIONS 34166 None. 34167 SEE ALSO 34168 localedef, nm 34169 CHANGE HISTORY First released in Issue 4. 34170 34171 Issue 6 34172 This utility is marked as part of the User Portability Utilities option. The obsolescent SYNOPSIS is removed. 34173

The normative text is reworded to avoid use of the term "must" for application requirements.

34174

strip Utilities

34175 **NAME** 34176 strip — remove unnecessary information from executable files (**DEVELOPMENT**) 34177 SYNOPSIS 34178 SD strip file... 34179 34180 DESCRIPTION The strip utility shall remove from executable files named by the file operands any information 34181 the implementor deems unnecessary for execution of those files. The nature of that information 34182 is unspecified. The effect of *strip* shall be similar to the use of the -s option to c99 or fort77. 34183 **34184 OPTIONS** None. 34185 34186 OPERANDS The following operand shall be supported: 34187 file 34188 A pathname referring to an executable file. 34189 **STDIN** Not used. 34190 34191 INPUT FILES The input files shall be in the form of executable files successfully produced by any compiler 34192 34193 defined by this volume of IEEE Std 1003.1-2001. 34194 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *strip*: 34195 LANG Provide a default value for the internationalization variables that are unset or null. 34196 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 34197 Internationalization Variables for the precedence of internationalization variables 34198 34199 used to determine the values of locale categories.) 34200 LC ALL If set to a non-empty string value, override the values of all the other internationalization variables. 34201 34202 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 34203 characters (for example, single-byte as opposed to multi-byte characters in arguments). 34204 LC MESSAGES 34205 Determine the locale that should be used to affect the format and contents of 34206 34207 diagnostic messages written to standard error. 34208 XSI **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 34209 ASYNCHRONOUS EVENTS Default. 34210

34211 **STDOUT**

34212 Not used.

34213 STDERR

The standard error shall be used only for diagnostic messages. 34214

Utilities strip

34215 OUTPUT FILES 34216 The *strip* utility shall produce executable files of unspecified format. 34217 EXTENDED DESCRIPTION 34218 None. 34219 EXIT STATUS 34220 The following exit values shall be returned: 34221 Successful completion. >0 An error occurred. 34222 **34223 CONSEQUENCES OF ERRORS** 34224 Default. 34225 APPLICATION USAGE None. 34226 34227 EXAMPLES 34228 None. 34229 RATIONALE 34230 Historically, this utility has been used to remove the symbol table from an executable file. It was 34231 included since it is known that the amount of symbolic information can amount to several megabytes; the ability to remove it in a portable manner was deemed important, especially for 34232 smaller systems. 34233 34234 The behavior of strip is said to be the same as the -s option to a compiler. While the end result is 34235 essentially the same, it is not required to be identical. 34236 FUTURE DIRECTIONS 34237 None. 34238 SEE ALSO ar, c99, fort77 34239 34240 CHANGE HISTORY

This utility is marked as part of the Software Development Utilities option.

First released in Issue 2.

34241

34243

34242 Issue 6

stty Utilities

34244 NAME

34250

34251

34252

34253

34254 34255

34256

34257 34258

34259

34260

34261

34262

34274

34275 34276

34277

34278

34280

34282

34283

34284

34285

34286

stty — set the options for a terminal

34246 SYNOPSIS

34247 stty [-a|-g]34248 stty operands

34249 **DESCRIPTION**

The *stty* utility shall set or report on terminal I/O characteristics for the device that is its standard input. Without options or operands specified, it shall report the settings of certain characteristics, usually those that differ from implementation-defined defaults. Otherwise, it shall modify the terminal state according to the specified operands. Detailed information about the modes listed in the first five groups below are described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface. Operands in the Combination Modes group (see **Combination Modes** (on page 889)) are implemented using operands in the previous groups. Some combinations of operands are mutually-exclusive on some terminal types; the results of using such combinations are unspecified.

Typical implementations of this utility require a communications line configured to use the **termios** interface defined in the System Interfaces volume of IEEE Std 1003.1-2001. On systems where none of these lines are available, and on lines not currently configured to support the **termios** interface, some of the operands need not affect terminal characteristics.

34263 OPTIONS

The *stty* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

34266 The following options shall be supported:

34267 —a Write to standard output all the current settings for the terminal.

Write to standard output all the current settings in an unspecified form that can be used as arguments to another invocation of the *stty* utility on the same system. The form used shall not contain any characters that would require quoting to avoid word expansion by the shell; see Section 2.6 (on page 36).

34272 **OPERANDS**

34273 The following operands shall be supported to set the terminal characteristics.

Control Modes

parenb (-parenb) Enable (disable) parity generation and detection. This shall have the effect of setting (not setting) PARENB in the termios c_cflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.

34279 parodd (-parodd)

Select odd (even) parity. This shall have the effect of setting (not setting) PARODD in the **termios** c_c cflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.

cs5 cs6 cs7 cs8 Select character size, if possible. This shall have the effect of setting CS5, CS6,

CS7, and CS8, respectively, in the **termios** *c_cflag* field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal

Interface.

34287 *number* Set terminal baud rate to the number given, if possible. If the baud rate is set to zero, the modem control lines shall no longer be asserted. This shall have

Utilities stty

34289 34290 34291		the effect of setting the input and output termios baud rate values as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34292 34293 34294 34295 34296	ispeed number	Set terminal input baud rate to the number given, if possible. If the input baud rate is set to zero, the input baud rate shall be specified by the value of the output baud rate. This shall have the effect of setting the input termios baud rate values as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34297 34298 34299 34300 34301	ospeed number	Set terminal output baud rate to the number given, if possible. If the output baud rate is set to zero, the modem control lines shall no longer be asserted. This shall have the effect of setting the output termios baud rate values as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34302 34303 34304 34305	hupcl (-hupcl)	Stop asserting modem control lines (do not stop asserting modem control lines) on last close. This shall have the effect of setting (not setting) HUPCL in the termios c_cflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34306	hup (-hup)	Equivalent to hupcl (- hupcl).
34307 34308 34309	cstopb (-cstopb)	Use two (one) stop bits per character. This shall have the effect of setting (not setting) CSTOPB in the termios c_cflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34310 34311 34312	cread (-cread)	Enable (disable) the receiver. This shall have the effect of setting (not setting) CREAD in the termios c_cflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34313 34314 34315 34316	clocal (-clocal)	Assume a line without (with) modem control. This shall have the effect of setting (not setting) CLOCAL in the termios c_cflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34317	It is unspecified v	whether <i>stty</i> shall report an error if an attempt to set a Control Mode fails.
34318	Input Modes	
34319 34320 34321	ignbrk (–ignbrk)	Ignore (do not ignore) break on input. This shall have the effect of setting (not setting) IGNBRK in the termios c_i field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34322 34323 34324	brkint (-brkint)	Signal (do not signal) INTR on break. This shall have the effect of setting (not setting) BRKINT in the termios c _iflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34325 34326 34327 34328	ignpar (–ignpar)	Ignore (do not ignore) bytes with parity errors. This shall have the effect of setting (not setting) IGNPAR in the termios c_{iflag} field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34329 34330 34331 34332 34333	parmrk (–parmrk	Mark (do not mark) parity errors. This shall have the effect of setting (not setting) PARMRK in the termios c _iflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.

stty Utilities

34334 34335 34336	inpck (–inpck)	Enable (disable) input parity checking. This shall have the effect of setting (not setting) INPCK in the termios c_i field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34337 34338 34339 34340	istrip (–istrip)	Strip (do not strip) input characters to seven bits. This shall have the effect of setting (not setting) ISTRIP in the termios c _ <i>iflag</i> field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34341 34342 34343	inlcr (-inlcr)	Map (do not map) NL to CR on input. This shall have the effect of setting (not setting) INLCR in the termios c_i field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34344 34345 34346	igncr (–igncr)	Ignore (do not ignore) CR on input. This shall have the effect of setting (not setting) IGNCR in the termios c_i field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34347 34348 34349	icrnl (-icrnl)	Map (do not map) CR to NL on input. This shall have the effect of setting (not setting) ICRNL in the termios c_i field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34350 34351 34352 34353 34354	ixon (–ixon)	Enable (disable) START/STOP output control. Output from the system is stopped when the system receives STOP and started when the system receives START. This shall have the effect of setting (not setting) IXON in the termios c_iflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34355 XSI 34356 34357	ixany (–ixany)	Allow any character to restart output. This shall have the effect of setting (not setting) IXANY in the termios c _iflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34358 34359 34360 34361 34362	ixoff (-ixoff)	Request that the system send (not send) STOP characters when the input queue is nearly full and START characters to resume data transmission. This shall have the effect of setting (not setting) IXOFF in the $termios\ c_iflag$ field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34363	Output Modes	
34364 34365 34366 34367	opost (-opost)	Post-process output (do not post-process output; ignore all other output modes). This shall have the effect of setting (not setting) OPOST in the termios c_oflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34368 XSI 34369 34370 34371	ocrnl (-ocrnl)	Map (do not map) CR to NL on output This shall have the effect of setting (not setting) OCRNL in the termios c_{-} of lag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34372 34373 34374	onocr (–onocr)	Do not (do) output CR at column zero. This shall have the effect of setting (not setting) ONOCR in the termios c _oflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34375 34376 34377 34378	onlret (-onlret)	The terminal newline key performs (does not perform) the CR function. This shall have the effect of setting (not setting) ONLRET in the termios c_{-} of lag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.

Utilities stty

34379 34380 34381 34382	ofill (-ofill)	Use fill characters (use timing) for delays. This shall have the effect of setting (not setting) OFILL in the termios c_oflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34383 34384 34385	ofdel (–ofdel)	Fill characters are DELs (NULs). This shall have the effect of setting (not setting) OFDEL in the termios c_oflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34386 34387 34388 34389	cr0 cr1 cr2 cr3	Select the style of delay for CRs. This shall have the effect of setting CRDLY to CR0, CR1, CR2, or CR3, respectively, in the termios c_oflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34390 34391 34392 34393	nl0 nl1	Select the style of delay for NL. This shall have the effect of setting NLDLY to NL0 or NL1, respectively, in the termios c_oflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34394	tab0 tab1 tab2 tal	03
34395 34396 34397 34398 34399		Select the style of delay for horizontal tabs. This shall have the effect of setting TABDLY to TAB0, TAB1, TAB2, or TAB3, respectively, in the termios c_oflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface. Note that TAB3 has the effect of expanding <tab>s to <space>s.</space></tab>
34400	tabs (–tabs)	Synonym for tab0 (tab3).
34401 34402 34403 34404	bs0 bs1	Select the style of delay for backspaces. This shall have the effect of setting BSDLY to BS0 or BS1, respectively, in the termios c_oflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34405 34406 34407 34408	ff0 ff1	Select the style of delay for form-feeds. This shall have the effect of setting FFDLY to FF0 or FF1, respectively, in the termios c_oflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34409 34410 34411 34412	vt0 vt1	Select the style of delay for vertical-tabs. This shall have the effect of setting VTDLY to VT0 or VT1, respectively, in the termios c_oflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34413	Local Modes	
34414 34415 34416 34417	isig (-isig)	Enable (disable) the checking of characters against the special control characters INTR, QUIT, and SUSP. This shall have the effect of setting (not setting) ISIG in the termios c_lflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34418 34419 34420 34421	icanon (–icanon)	Enable (disable) canonical input (ERASE and KILL processing). This shall have the effect of setting (not setting) ICANON in the termios c_lflag field, as defined in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
34422 34423 34424	iexten (–iexten)	Enable (disable) any implementation-defined special control characters not currently controlled by icanon , isig , ixon , or ixoff . This shall have the effect of setting (not setting) IEXTEN in the termios c_lflag field, as defined in the Base

stty Utilities

34425 34426		Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.	
34427	echo (-echo)	Echo back (do not echo back) every character typed. This shall have the effect	
34428	,	of setting (not setting) ECHO in the termios c_{-} <i>Iflag</i> field, as defined in the Base	
34429		Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal	
34430		Interface.	
34431	echoe (-echoe)	The ERASE character visually erases (does not erase) the last character in the	
34432		current line from the display, if possible. This shall have the effect of setting	
34433		(not setting) ECHOE in the termios c_{-} lflag field, as defined in the Base	
34434		Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal	
34435		Interface.	
34436	echok (-echok)	Echo (do not echo) NL after KILL character. This shall have the effect of	
34437		setting (not setting) ECHOK in the termios c_lflag field, as defined in the Base	
34438		Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal	
34439		Interface.	
34440	echonl (-echonl)	Echo (do not echo) NL, even if echo is disabled. This shall have the effect of	
34441		setting (not setting) ECHONL in the termios c_{l} field, as defined in the	
34442		Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General	
34443		Terminal Interface.	
34444	noflsh (-noflsh)	Disable (enable) flush after INTR, QUIT, SUSP. This shall have the effect of	
34445		setting (not setting) NOFLSH in the termios c_l field, as defined in the Base	
34446		Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal	
34447		Interface.	
34448	tostop (-tostop)	Send SIGTTOU for background output. This shall have the effect of setting	
34449	-	(not setting) TOSTOP in the termios c_l field, as defined in the Base	
34450		Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal	
34451		Interface.	
34452	Special Control C	Character Assignments	
34453	<control>-characte</control>	er string	
34454	Set <control>-character to string. If <control>-character is one of the character sequences in</control></control>		
34455	the first column of the following table, the corresponding Base Definitions volume of		
34456	IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface control character from the		
34457	second column shall be recognized. This has the effect of setting the corresponding element		
34458	of the termios $c_{-}cc$ array (see the Base Definitions volume of IEEE Std 1003.1-2001, Chapter		
34459	13, Headers,	<termios.h>).</termios.h>	

Utilities stty

Table 4-19	Control	Character	Names in	stty

Control Character	c_cc Subscript	Description
eof	VEOF	EOF character
eol	VEOL	EOL character
erase	VERASE	ERASE character
intr	VINTR	INTR character
kill	VKILL	KILL character
quit	VQUIT	QUIT character
susp	VSUSP	SUSP character
start	VSTART	START character
stop	VSTOP	STOP character

If *string* is a single character, the control character shall be set to that character. If *string* is the two-character sequence "^-" or the string *undef*, the control character shall be set to _POSIX_VDISABLE, if it is in effect for the device; if _POSIX_VDISABLE is not in effect for the device, it shall be treated as an error. In the POSIX locale, if *string* is a two-character sequence beginning with circumflex ('^'), and the second character is one of those listed in the "^c" column of the following table, the control character shall be set to the corresponding character value in the Value column of the table.

Table 4-20 Circumflex Control Characters in *stty*

^c	Value	^c	Value	^c	Value
a, A	<soh></soh>	1, L	<ff></ff>	w, W	<etb></etb>
b, B	<stx></stx>	m, M	<cr></cr>	x, X	<can></can>
c, C	<etx></etx>	n, N	<so></so>	у, Ү	
d, D	<eot></eot>	0, 0	<si></si>	z, Z	
e, E	<enq></enq>	p, P	<dle></dle>	[<esc></esc>
f, F	<ack></ack>	q, Q	<dc1></dc1>	\	<fs></fs>
g, G	<bel></bel>	r, R	<dc2></dc2>]	<gs></gs>
h, H	<bs></bs>	s, S	<dc3></dc3>	^	<rs></rs>
i, I	<ht></ht>	t, T	<dc4></dc4>	_	<us></us>
j, J	<lf></lf>	u, U	<nak></nak>	?	
k, K	<vt></vt>	v, V	<syn></syn>		

min number

Set the value of MIN to *number*. MIN is used in non-canonical mode input processing (icanon)

time number

Set the value of TIME to *number*. TIME is used in non-canonical mode input processing (**icanon**).

Combination Modes

34498 saved settings

Set the current terminal characteristics to the saved settings produced by the **-g** option.

evenp or **parity**

Enable **parenb** and **cs7**; disable **parodd**.

oddp

Enable parenb, cs7, and parodd.

stty Utilities

34504 34505		<pre>-parity, -evenp, or -oddp Disable parenb, and set cs8.</pre>		
34506 XSI 34507	•	raw (-raw or cooked) Enable (disable) raw input and output. Raw mode shall be equivalent to setting:		
34508 34509		cs8 erase ^- kill ^- intr ^- \ uit ^- eof ^- eol ^post -inpck		
34510 34511	nl (-nl) Disable	(enable) icrnl . In addition, – nl unsets inlcr and igncr .		
34512	ek Reset E	RASE and KILL characters back to system defaults.		
34513 34514	sane Reset al	ll modes to some reasonable, unspecified, values.		
34515 STDIN				
34516 34517		o input is read from standard input, standard input shall be used to get the current D characteristics and to set new terminal I/O characteristics.		
34518 INPUT 34519	FILES None.			
34520 ENVIR 34521	ONMENT VA	ARIABLES ng environment variables shall affect the execution of stty:		
34522 34523 34524 34525	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)		
34526 34527	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.		
34528 34529 34530	LC_CTYPE	This variable determines 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 which characters are in the class print .		
34531	LC_MESSA	GES		
34532 34533		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.		
34534 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .		
34535 ASYNC 34536	CHRONOUS Default.	EVENTS		
34537 STDOU 34538		are specified, no output shall be produced.		
34539 34540		If the $-\mathbf{g}$ option is specified, $stty$ shall write to standard output the current settings in a form that can be used as arguments to another instance of $stty$ on the same system.		
34541 34542 34543 34544	be written t <space>-sep</space>	If the –a option is specified, all of the information as described in the OPERANDS section shall be written to standard output. Unless otherwise specified, this information shall be written as <space>-separated tokens in an unspecified format, on one or more lines, with an unspecified number of tokens per line. Additional information may be written.</space>		
34545 34546	•	s or operands are specified, an unspecified subset of the information written for the hall be written.		

Utilities Stty

34547 If speed information is written as part of the default output, or if the -a option is specified and if 34548 the terminal input speed and output speed are the same, the speed information shall be written as follows: 34549 "speed %d baud;", < speed> 34550 34551 Otherwise, speeds shall be written as: "ispeed %d baud; ospeed %d baud; ", <ispeed>, <ospeed> 34552 In locales other than the POSIX locale, the word baud may be changed to something more 34553 34554 appropriate in those locales. If control characters are written as part of the default output, or if the -a option is specified, 34555 control characters shall be written as: 34556 "%s = %s;", <control-character name>, <value> 34557 where <value> is either the character, or some visual representation of the character if it is non-34558 34559 printable, or the string *undef* if the character is disabled. **34560 STDERR** 34561 The standard error shall be used only for diagnostic messages. 34562 OUTPUT FILES None. 34564 EXTENDED DESCRIPTION 34565 None. 34566 EXIT STATUS The following exit values shall be returned: 34567 0 The terminal options were read or set successfully. 34568 >0 An error occurred. 34569 34570 CONSEQUENCES OF ERRORS Default. 34571 34572 APPLICATION USAGE 34573 The $-\mathbf{g}$ flag is designed to facilitate the saving and restoring of terminal state from the shell level. 34574 For example, a program may:

```
34575 saveterm="$(stty -g)" # save terminal state
34576 stty (new settings) # set new state
34577 ... # ...
34578 stty $saveterm # restore terminal state
```

Since the format is unspecified, the saved value is not portable across systems.

Since the **–a** format is so loosely specified, scripts that save and restore terminal settings should

34581 use the $-\mathbf{g}$ option.

34582 EXAMPLES

34583 None.

34584 RATIONALE

The original *stty* description was taken directly from System V and reflected the System V terminal driver **termio**. It has been modified to correspond to the terminal driver **termios**.

Output modes are specified only for XSI-conformant systems. All implementations are expected to provide *stty* operands corresponding to all of the output modes they support.

stty Utilities

The *stty* utility is primarily used to tailor the user interface of the terminal, such as selecting the preferred ERASE and KILL characters. As an application programming utility, *stty* can be used within shell scripts to alter the terminal settings for the duration of the script.

The **termios** section states that individual disabling of control characters is possible through the option _POSIX_VDISABLE. If enabled, two conventions currently exist for specifying this:

System V uses "^-", and BSD uses *undef*. Both are accepted by *stty* in this volume of IEEE Std 1003.1-2001. The other BSD convention of using the letter 'u' was rejected because it conflicts with the actual letter 'u', which is an acceptable value for a control character.

Early proposals did not specify the mapping of "^c" to control characters because the control characters were not specified in the POSIX locale character set description file requirements. The control character set is now specified in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 3, Definitions so the historical mapping is specified. Note that although the mapping corresponds to control-character key assignments on many terminals that use the ISO/IEC 646: 1991 standard (or ASCII) character encodings, the mapping specified here is to the control characters, not their keyboard encodings.

Since **termios** supports separate speeds for input and output, two new options were added to specify each distinctly.

Some historical implementations use standard input to get and set terminal characteristics; others use standard output. Since input from a login TTY is usually restricted to the owner while output to a TTY is frequently open to anyone, using standard input provides fewer chances of accidentally (or maliciously) altering the terminal settings of other users. Using standard input also allows *stty* –**a** and *stty* –**g** output to be redirected for later use. Therefore, usage of standard input is required by this volume of IEEE Std 1003.1-2001.

34612 FUTURE DIRECTIONS

34613 None.

34614 SEE ALSO

34597 34598

34599 34600

34601

34602 34603

34606

34607

34608

34609 34610

34611

Chapter 2 (on page 29), the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface, <termios.h>

34617 CHANGE HISTORY

34618 First released in Issue 2.

34619 **Issue 5**

34620 The description of **tabs** is clarified.

34621 The FUTURE DIRECTIONS section is added.

34622 Issue 6

The legacy items iuclc(-iuclc), xcase, olcuc(-olcuc), lcase(-lcase), and LCASE(-LCASE) are removed.

34625 IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/37 is applied, applying IEEE PASC Interpretation 1003.2 #133, fixing an error in the OPERANDS section for the Combination Modes nl(-nl).

Utilities tabs

34628 NAME				
34629	tabs — set terminal tabs			
34630 SYNOI 34631 UP XSI	34630 SYNOPSIS 34631 UP XSI tabs $[-n -a -a2 -c -c2 -c3 -f -p -s -u][+m[n]]$ [-T type]			
	tabs [-T type] [+[n]] $n1[,n2,]$			
34632 34633	tabs [-1	Lype][+[n]] n1[,n2,]		
34634 DESCRIPTION				
34635	The tabs utility shall display a series of characters that first clears the hardware terminal tab			
34636 XSI 34637	settings and then initializes the tab stops at the specified positions and optionally adjusts the margin.			
34638 34639	The phrase "tab-stop position N " shall be taken to mean that, from the start of a line of output, tabbing to position N shall cause the next character output to be in the $(N+1)$ th column position			
34640	on that line. The maximum number of tab stops allowed is terminal-dependent.			
34641	It need not be possible to implement <i>tabs</i> on certain terminals. If the terminal type obtained from			
34642 34643	the <i>TERM</i> environment variable or -T option represents such a terminal, an appropriate diagnostic message shall be written to standard error and <i>tabs</i> shall exit with a status greater			
34644	than zero.			
34645 OPTIO	NS			
34646	The tabs utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section			
34647 XSI 34648	12.2, Utility Syntax Guidelines, except for various extensions: the options -a2 , -c2 , and -c3 are multi-character.			
34649	The following options shall be supported:			
34650	-n	Specify repetitive tab stops separated by a uniform number of column positions, n,		
34651 34652		where n is a single-digit decimal number. The default usage of <i>tabs</i> with no arguments shall be equivalent to tabs –8. When – 0 is used, the tab stops shall be		
34653		cleared and no new ones set.		
34654 XSI	-a	1,10,16,36,72		
34655	_	Assembler, applicable to some mainframes.		
34656 XSI 34657	-a2	1,10,16,40,72 Assembler, applicable to some mainframes.		
34658 XSI	-с	1,8,12,16,20,55		
34659		COBOL, normal format.		
34660 XSI 34661	-c2	1,6,10,14,49		
		COBOL, compact format (columns 1 to 6 omitted).		
34662 XSI 34663	-с3	1,6,10,14,18,22,26,30,34,38,42,46,50,54,58,62,67 COBOL compact format (columns 1 to 6 omitted), with more tabs than –c2 .		
34664 XSI	−f	1,7,11,15,19,23		
34665		FORTRAN		
34666 XSI 34667	-p	1,5,9,13,17,21,25,29,33,37,41,45,49,53,57,61 PL/1		
34668 XSI	-s	1,10,55		
34669		SNOBOL		
34670 XSI	–u	1,12,20,44		
34671		Assembler, applicable to some mainframes.		

tabs Utilities

34672 34673 34674	-T type	Indicate the type of terminal. If this option is not supplied and the <i>TERM</i> variable is unset or null, an unspecified default terminal type shall be used. The setting of <i>type</i> shall take precedence over the value in <i>TERM</i> .			
34675 OPERA	NDS				
34676	The following operand shall be supported:				
34677 34678 34679 34680 34681 34682	n1[,n2,]	A single command line argument that consists of tab-stop values separated using either commas or tab-s. The application shall ensure that the tab-stop values are positive decimal integers in strictly ascending order. If any number (except the first one) is preceded by a plus sign, it is taken as an increment to be added to the previous value. For example, the tab lists $1,10,20,30$ and $1,10,+10,+10$ are considered to be identical.			
34683 STDIN 34684	Not used.				
34685 INPUT	FILES				
34686	None.				
34687 ENVIRONMENT VARIABLES					
The following environment variables shall affect the execution of <i>tabs</i> :					
34689 34690 34691 34692	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)			
34693 34694	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.			
34695 34696 34697	LC_CTYPE	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).			
34698	LC_MESSAC	LC_MESSAGES			
34699 34700		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.			
34701 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .			
34702 34703	TERM	Determine the terminal type. If this variable is unset or null, and if the $-\mathbf{T}$ option is not specified, an unspecified default terminal type shall be used.			
34704 ASYNCHRONOUS EVENTS 34705 Default.					
34706 STDOU	T				
34707		output is a terminal, the appropriate sequence to clear and set the tab stops may be			
34708 34709	written to standard output in an unspecified format. If standard output is not a terminal, undefined results occur.				
34710 STDERR					
34711					
34712 OUTPUT FILES					

34713

None.

Utilities tabs

34714 EXTENDED DESCRIPTION

34715 None.

34716 EXIT STATUS

34717 The following exit values shall be returned:

34718 0 Successful completion.

34719 >0 An error occurred.

34720 CONSEQUENCES OF ERRORS

34721 Default.

34722 APPLICATION USAGE

This utility makes use of the terminal's hardware tabs and the *stty tabs* option.

This utility is not recommended for application use.

Some integrated display units might not have escape sequences to set tab stops, but may be set by internal system calls. On these terminals, *tabs* works if standard output is directed to the terminal; if output is directed to another file, however, *tabs* fails.

34728 EXAMPLES

34729 None.

34730 RATIONALE

34731

34732

34733

34734

34735

34739

34740

34741

34742

34743

34744

34745

34746

34748

34749

Consideration was given to having the *tput* utility handle all of the functions described in *tabs*. However, the separate *tabs* utility was retained because it seems more intuitive to use a command named *tabs* than *tput* with a new option. The *tput* utility does not support setting or clearing tabs, and no known historical version of *tabs* supports the capability of setting arbitrary tab stops.

The System V *tabs* interface is very complex; the version in this volume of IEEE Std 1003.1-2001 has a reduced feature list, but many of the features omitted were restored as XSI extensions even though the supported languages and coding styles are primarily historical.

There was considerable sentiment for specifying only a means of resetting the tabs back to a known state—presumably the "standard" of tabs every eight positions. The following features were omitted:

• Setting tab stops via the first line in a file, using -- file. Since even the SVID has no complete explanation of this feature, it is doubtful that it is in widespread use.

In an early proposal, a -t *tablist* option was added for consistency with *expand*; this was later removed when inconsistencies with the historical list of tabs were identified.

Consideration was given to adding a $-\mathbf{p}$ option that would output the current tab settings so that they could be saved and then later restored. This was not accepted because querying the tab stops of the terminal is not a capability in historical *terminfo* or *termcap* facilities and might not be supported on a wide range of terminals.

34750 FUTURE DIRECTIONS

34751 None.

34752 SEE ALSO

34753 expand, stty, tput, unexpand

tabs Utilities

34754 CHANGE HISTORY

First released in Issue 2.

34756 **Issue 6**

34757 This utility is marked as part of the User Portability Utilities option.

34758 The normative text is reworded to avoid use of the term "must" for application requirements.

Utilities tail

34759 **NAME** 34760 tail — copy the last part of a file 34761 SYNOPSIS 34762 tail [-f] [-c number | -n number] [file] 34763 DESCRIPTION 34764 The tail utility shall copy its input file to the standard output beginning at a designated place. Copying shall begin at the point in the file indicated by the -c number or -n number options. The 34765 34766 option-argument *number* shall be counted in units of lines or bytes, according to the options -n and -c. Both line and byte counts start from 1. 34767 Tails relative to the end of the file may be saved in an internal buffer, and thus may be limited in 34768 length. Such a buffer, if any, shall be no smaller than {LINE MAX}*10 bytes. 34769 34770 OPTIONS The tail utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 34771 12.2, Utility Syntax Guidelines. 34772 The following options shall be supported: 34773 -c number The application shall ensure that the number option-argument is a decimal integer 34774 whose sign affects the location in the file, measured in bytes, to begin the copying: 34775 **Copying Starts** 34776 Sign Relative to the beginning of the file. 34777 + Relative to the end of the file. 34778 Relative to the end of the file. none 34779 The origin for counting shall be 1; that is, -c + 1 represents the first byte of the file, 34780 -c -1 the last. 34781 $-\mathbf{f}$ If the input file is a regular file or if the file operand specifies a FIFO, do not 34782 terminate after the last line of the input file has been copied, but read and copy 34783 further bytes from the input file when they become available. If no file operand is 34784 34785 specified and standard input is a pipe, the -f option shall be ignored. If the input file is not a FIFO, pipe, or regular file, it is unspecified whether or not the -f option 34786 shall be ignored. 34787 -**n** number This option shall be equivalent to -c *number*, except the starting location in the file 34788 shall be measured in lines instead of bytes. The origin for counting shall be 1; that 34789 is, $-\mathbf{n} + 1$ represents the first line of the file, $-\mathbf{n} - 1$ the last. 34790 If neither $-\mathbf{c}$ nor $-\mathbf{n}$ is specified, $-\mathbf{n}$ 10 shall be assumed. 34791 34792 **OPERANDS** The following operand shall be supported: 34793 A pathname of an input file. If no file operands are specified, the standard input file 34794 shall be used. 34795 34796 **STDIN** The standard input shall be used only if no file operands are specified. See the INPUT FILES 34797 section. 34798

tail **Utilities**

34799 INPUT FILES 34800 If the -c option is specified, the input file can contain arbitrary data; otherwise, the input file shall be a text file. 34801 34802 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *tail*: 34803 Provide a default value for the internationalization variables that are unset or null. 34804 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 34805 Internationalization Variables for the precedence of internationalization variables 34806 used to determine the values of locale categories.) 34807 LC_ALL If set to a non-empty string value, override the values of all the other 34808 internationalization variables. 34809 Determine the locale for the interpretation of sequences of bytes of text data as 34810 LC_CTYPE 34811 characters (for example, single-byte as opposed to multi-byte characters in 34812 arguments and input files). LC_MESSAGES 34813 Determine the locale that should be used to affect the format and contents of 34814 diagnostic messages written to standard error. 34815 34816 XSI NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 34817 ASYNCHRONOUS EVENTS 34818 Default. **34819 STDOUT** 34820 The designated portion of the input file shall be written to standard output. **34821 STDERR** 34822 The standard error shall be used only for diagnostic messages. 34823 OUTPUT FILES 34824 None. 34825 EXTENDED DESCRIPTION 34826 None. 34827 EXIT STATUS 34828 The following exit values shall be returned: Successful completion. 34829 >0 An error occurred. 34830

34833 APPLICATION USAGE

34831 CONSEQUENCES OF ERRORS Default.

The -c option should be used with caution when the input is a text file containing multi-byte 34834 characters; it may produce output that does not start on a character boundary. 34835

Although the input file to tail can be any type, the results might not be what would be expected on some character special device files or on file types not described by the System Interfaces volume of IEEE Std 1003.1-2001. Since this volume of IEEE Std 1003.1-2001 does not specify the block size used when doing input, tail need not read all of the data from devices that only perform block transfers.

34832

34836

34837

34838

34839

34840

Utilities tail

34841 EXAMPLES

The –f option can be used to monitor the growth of a file that is being written by some other process. For example, the command:

34844 tail -f fred

prints the last ten lines of the file **fred**, followed by any lines that are appended to **fred** between the time *tail* is initiated and killed. As another example, the command:

34847 tail -f -c 15 fred

prints the last 15 bytes of the file **fred**, followed by any bytes that are appended to **fred** between the time *tail* is initiated and killed.

34850 RATIONALE

This version of *tail* was created to allow conformance to the Utility Syntax Guidelines. The historical **-b** option was omitted because of the general non-portability of block-sized units of text. The **-c** option historically meant "characters", but this volume of IEEE Std 1003.1-2001 indicates that it means "bytes". This was selected to allow reasonable implementations when multi-byte characters are possible; it was not named **-b** to avoid confusion with the historical **-b**.

34856 —**D.**

The origin of counting both lines and bytes is 1, matching all widespread historical implementations.

The restriction on the internal buffer is a compromise between the historical System V implementation of 4 096 bytes and the BSD 32 768 bytes.

The –f option has been implemented as a loop that sleeps for 1 second and copies any bytes that are available. This is sufficient, but if more efficient methods of determining when new data are available are developed, implementations are encouraged to use them.

34864 Historical documentation indicates that *tail* ignores the -f option if the input file is a pipe (pipe and FIFO on systems that support FIFOs). On BSD-based systems, this has been true; on System 34865 34866 V-based systems, this was true when input was taken from standard input, but it did not ignore the -f flag if a FIFO was named as the file operand. Since the -f option is not useful on pipes and 34867 34868 all historical implementations ignore -f if no file operand is specified and standard input is a pipe, this volume of IEEE Std 1003.1-2001 requires this behavior. However, since the -f option is 34869 useful on a FIFO, this volume of IEEE Std 1003.1-2001 also requires that if standard input is a 34870 FIFO or a FIFO is named, the -f option shall not be ignored. Although historical behavior does 34871 not ignore the -f option for other file types, this is unspecified so that implementations are 34872 allowed to ignore the –**f** option if it is known that the file cannot be extended. 34873

This was changed to the current form based on comments noting that **-c** was almost never used without specifying a number and that there was no need to specify **-l** if **-n** *number* was given.

34876 FUTURE DIRECTIONS

34877 None.

34878 **SEE ALSO**

34879 *head*

34880 CHANGE HISTORY

34881 First released in Issue 2.

34882 **Issue 6**

34883 The obsolescent SYNOPSIS lines and associated text are removed.

The normative text is reworded to avoid use of the term "must" for application requirements.

talk Utilities

```
34885 NAME
34886 talk — talk to another user

34887 SYNOPSIS
34888 UP talk address [terminal]
34889

34890 DESCRIPTION
34891 The talk utility is a two-way, screen-oriented communication program.
```

34892 When first invoked, *talk* shall send a message similar to:

```
34893 Message from <unspecified string>
34894 talk: connection requested by your_address
34895 talk: respond with: talk your_address
```

to the specified *address*. At this point, the recipient of the message can reply by typing:

```
34897 talk your address
```

34900 34901

34902 34903

34904

34905 34906

34907

34908 34909

34910

34911

34912

34913

34914 34915

34916 34917

34918

34919

34920 34921

34922

Once communication is established, the two parties can type simultaneously, with their output displayed in separate regions of the screen. Characters shall be processed as follows:

- Typing the alert character shall alert the recipient's terminal.
- Typing <control>-L shall cause the sender's screen regions to be refreshed.
- Typing the erase and kill characters shall affect the sender's terminal in the manner described by the **termios** interface in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
- Typing the interrupt or end-of-file characters shall terminate the local *talk* utility. Once the *talk* session has been terminated on one side, the other side of the *talk* session shall be notified that the *talk* session has been terminated and shall be able to do nothing except exit.
- Typing characters from *LC_CTYPE* classifications **print** or **space** shall cause those characters to be sent to the recipient's terminal.
- When and only when the stty iexten local mode is enabled, the existence and processing of additional special control characters and multi-byte or single-byte functions shall be implementation-defined.
- Typing other non-printable characters shall cause implementation-defined sequences of printable characters to be sent to the recipient's terminal.

Permission to be a recipient of a *talk* message can be denied or granted by use of the *mesg* utility. However, a user's privilege may further constrain the domain of accessibility of other users' terminals. The *talk* utility shall fail when the user lacks the appropriate privileges to perform the requested action.

Certain block-mode terminals do not have all the capabilities necessary to support the simultaneous exchange of messages required for *talk*. When this type of exchange cannot be supported on such terminals, the implementation may support an exchange with reduced levels of simultaneous interaction or it may report an error describing the terminal-related deficiency.

34923 OPTIONS

34924 None.

Utilities talk

34925 OPERANDS 34926 The following operands shall be supported:							
34927 34928 34929	address	The recipient of the <i>talk</i> session. One form of <i>address</i> is the <i><user name=""></user></i> , as returned by the <i>who</i> utility. Other address formats and how they are handled are unspecified.					
34930 34931 34932 34933	terminal	If the recipient is logged in more than once, the <i>terminal</i> argument can be used to indicate the appropriate terminal name. If <i>terminal</i> is not specified, the <i>talk</i> message shall be displayed on one or more accessible terminals in use by the recipient. The format of <i>terminal</i> shall be the same as that returned by the <i>who</i> utility.					
34934 STDIN 34935 34936 34937	Characters read from standard input shall be copied to the recipient's terminal in an unspecified manner. If standard input is not a terminal, talk shall write a diagnostic message and exit with a non-zero status.						
34938 INPUT 34939	FILES None.						
34940 ENVIR 34941	ONMENT VA	ARIABLES ng environment variables shall affect the execution of <i>talk</i> :					
34942 34943 34944 34945	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)					
34946 34947	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.					
34948 34949 34950 34951	LC_CTYPE	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). If the recipient's locale does not use an <i>LC_CTYPE</i> equivalent to the sender's, the results are undefined.					
34952 34953 34954 34955	LC_MESSAGES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.						
34956 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.					
34957 34958	TERM	Determine the name of the invoker's terminal type. If this variable is unset or null, an unspecified default terminal type shall be used.					
34959 ASYNO 34960 34961		EVENTS alk utility receives a SIGINT signal, the utility shall terminate and exit with a zero all take the standard action for all other signals.					
34962 STDOU 34963		output is a terminal, characters copied from the recipient's standard input may be					

If standard output is a terminal, characters copied from the recipient's standard input may be written to standard output. Standard output also may be used for diagnostic messages. If standard output is not a terminal, *talk* shall exit with a non-zero status.

34966 STDERR

34967 None.

talk Utilities

34968 OUTPUT FILES

34969 None.

34970 EXTENDED DESCRIPTION

34971 None.

34972 EXIT STATUS

34973 The following exit values shall be returned:

34974 0 Successful completion.

34975 >0 An error occurred or *talk* was invoked on a terminal incapable of supporting it.

34976 CONSEQUENCES OF ERRORS

34977 Default.

34978 APPLICATION USAGE

Because the handling of non-printable, non-<space>s is tied to the *stty* description of **iexten**, implementation extensions within the terminal driver can be accessed. For example, some implementations provide line editing functions with certain control character sequences.

34982 EXAMPLES

34983 None.

34984 RATIONALE

34985

34986 34987

34988

34989

34990

34991

34992 34993

34994

34995

34996

34997

34998 34999

35000

35001 35002

35003

35004

35005

35006

35007

35008

35009 35010 The *write* utility was included in this volume of IEEE Std 1003.1-2001 since it can be implemented on all terminal types. The *talk* utility, which cannot be implemented on certain terminals, was considered to be a "better" communications interface. Both of these programs are in widespread use on historical implementations. Therefore, both utilities have been specified.

All references to networking abilities (*talk*ing to a user on another system) were removed as being outside the scope of this volume of IEEE Std 1003.1-2001.

Historical BSD and System V versions of *talk* terminate both of the conversations when either user breaks out of the session. This can lead to adverse consequences if a user unwittingly continues to enter text that is interpreted by the shell when the other terminates the session. Therefore, the version of *talk* specified by this volume of IEEE Std 1003.1-2001 requires both users to terminate their end of the session explicitly.

Only messages sent to the terminal of the invoking user can be internationalized in any way:

- The original "Message from *<unspecified string>* ..." message sent to the terminal of the recipient cannot be internationalized because the environment of the recipient is as yet inaccessible to the *talk* utility. The environment of the invoking party is irrelevant.
- Subsequent communication between the two parties cannot be internationalized because the two parties may specify different languages in their environment (and non-portable characters cannot be mapped from one language to another).
- Neither party can be required to communicate in a language other than C and/or the one specified by their environment because unavailable terminal hardware support (for example, fonts) may be required.

The text in the STDOUT section reflects the usage of the verb "display" in this section; some *talk* implementations actually use standard output to write to the terminal, but this volume of IEEE Std 1003.1-2001 does not require that to be the case.

The format of the terminal name is unspecified, but the descriptions of *ps, talk, who*, and *write* require that they all use or accept the same format.

Utilities talk

35011	The handling of non-printable characters is partially implementation-defined because the details
35012	of mapping them to printable sequences is not needed by the user. Historical implementations,
35013	for security reasons, disallow the transmission of non-printable characters that may send
35014	commands to the other terminal.
35015 FUTUR	E DIRECTIONS
35016	None.
35017 SEE AL	
35018	mesg, stty, who, write, the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General
35019	Terminal Interface
35020 CHANO	GE HISTORY
35021	First released in Issue 4.
35022 Issue 6	
35023	This utility is marked as part of the User Portability Utilities option.

tee Utilities

35024 **NAME** 35025 tee — duplicate standard input 35026 SYNOPSIS 35027 tee [-ai][file...] 35028 DESCRIPTION The tee utility shall copy standard input to standard output, making a copy in zero or more files. 35029 The tee utility shall not buffer output. 35030 35031 If the -a option is not specified, output files shall be written (see Section 1.7.1.4 (on page 4). 35032 OPTIONS The tee utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, 35033 35034 Utility Syntax Guidelines. The following options shall be supported: 35035 35036 Append the output to the files. -a −i Ignore the SIGINT signal. 35037 35038 OPERANDS The following operands shall be supported: 35039 file A pathname of an output file. Processing of at least 13 file operands shall be 35040 35041 supported. 35042 **STDIN** The standard input can be of any type. 35043 35044 INPUT FILES None. 35045 35046 ENVIRONMENT VARIABLES 35047 The following environment variables shall affect the execution of *tee*: LANG 35048 Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 35049 35050 Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 35051 LC ALL 35052 If set to a non-empty string value, override the values of all the other internationalization variables. 35053 35054 LC_CTYPE 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 35055 arguments). 35056 LC_MESSAGES 35057 Determine the locale that should be used to affect the format and contents of 35058 diagnostic messages written to standard error. 35059 **NLSPATH** 35060 XSI Determine the location of message catalogs for the processing of *LC_MESSAGES*.

35061 ASYNCHRONOUS EVENTS

35062 Default, except that if the -i option was specified, SIGINT shall be ignored.

Utilities tee

35063 **STDOUT** 35064 The standard output shall be a copy of the standard input. 35065 STDERR The standard error shall be used only for diagnostic messages. 35066 35067 OUTPUT FILES If any file operands are specified, the standard input shall be copied to each named file. 35068 35069 EXTENDED DESCRIPTION None. 35070 35071 EXIT STATUS The following exit values shall be returned: 35072 35073 The standard input was successfully copied to all output files. 35074 >0 An error occurred. 35075 CONSEQUENCES OF ERRORS If a write to any successfully opened file operand fails, writes to other successfully opened file 35076 35077 operands and standard output shall continue, but the exit status shall be non-zero. Otherwise, the default actions specified in Section 1.11 (on page 20) apply. 35078 35079 APPLICATION USAGE The tee utility is usually used in a pipeline, to make a copy of the output of some utility. 35080 35081 The *file* operand is technically optional, but *tee* is no more useful than *cat* when none is specified. 35082 EXAMPLES 35083 Save an unsorted intermediate form of the data in a pipeline: 35084 ... | tee unsorted | sort > sorted 35085 RATIONALE The buffering requirement means that tee is not allowed to use ISO C standard fully buffered or 35086 35087 line-buffered writes. It does not mean that *tee* has to do 1-byte reads followed by 1-byte writes. It should be noted that early versions of BSD ignore any invalid options and accept a single '-' 35088 35089 as an alternative to -i. They also print a message if unable to open a file: 35090 "tee: cannot access %s\n", <pathname> Historical implementations ignore write errors. This is explicitly not permitted by this volume of 35091 IEEE Std 1003.1-2001. 35092

Some historical implementations use O_APPEND when providing append mode; others use the 35093 lseek() function to seek to the end-of-file after opening the file without O_APPEND. This volume 35094 of IEEE Std 1003.1-2001 requires functionality equivalent to using O_APPEND; see Section 35095 35096 1.7.1.4 (on page 4).

35097 FUTURE DIRECTIONS

None. 35098

35099 SEE ALSO

Chapter 1 (on page 1), cat, the System Interfaces volume of IEEE Std 1003.1-2001, lseek() 35100

35101 CHANGE HISTORY

First released in Issue 2. 35102

tee Utilities

35103 **Issue 6**

35104 IEEE PASC Interpretation 1003.2 #168 is applied.

test **Utilities**

35105 NAME							
35106	106 test — evaluate expression						
35107 SYNOP 35108	rest [expression]						
35109	[[expres	sion]]					
35110 DESCR 35111 35112 35113	RIPTION The <i>test</i> utility shall evaluate the <i>expression</i> and indicate the result of the evaluation by its exit status. An exit status of zero indicates that the expression evaluated as true and an exit status of 1 indicates that the expression evaluated as false.						
35114 35115		Index of the utility, which uses "[]" rather than <i>test</i> , the application shall ensure are brackets are separate arguments.					
35116 OPTIO	NS						
35117 35118		ity shall not recognize the " $$ " argument in the manner specified by guideline 10 in finitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.					
35119	No options :	shall be supported.					
35120 OPERA							
35121 35122		tion shall ensure that all operators and elements of primaries are presented as suments to the <i>test</i> utility.					
35123	The following	ng primaries can be used to construct expression:					
35124	− b file	True if <i>file</i> exists and is a block special file.					
35125	−c file	True if <i>file</i> exists and is a character special file.					
35126	− d file	True if <i>file</i> exists and is a directory.					
35127	− e file	True if <i>file</i> exists.					
35128	− f file	True if <i>file</i> exists and is a regular file.					
35129	− g file	True if <i>file</i> exists and its set-group-ID flag is set.					
35130	-h file	True if <i>file</i> exists and is a symbolic link.					
35131	−L file	True if <i>file</i> exists and is a symbolic link.					
35132	–n string	True if the length of <i>string</i> is non-zero.					
35133	− p file	True if <i>file</i> is a FIFO.					
35134 35135	− r file	True if <i>file</i> exists and is readable. True shall indicate that permission to read from <i>file</i> will be granted, as defined in Section 1.7.1.4 (on page 4).					
35136	−S file	True if <i>file</i> exists and is a socket.					
35137	−s file	True if <i>file</i> exists and has a size greater than zero.					
35138 35139 35140	-t file_descriptor True if the file whose file descriptor number is file_descriptor is open and is associated with a terminal.						
35141	−u file	True if <i>file</i> exists and its set-user-ID flag is set.					
35142 35143	-w file True if file exists and is writable. True shall indicate that permission to write from file will be granted, as defined in Section 1.7.1.4 (on page 4).						

test Utilities

35144 35145 35146	− x file	True if <i>file</i> exists and is executable. True shall indicate that permission to execute <i>file</i> will be granted, as defined in Section 1.7.1.4 (on page 4). If <i>file</i> is a directory, true shall indicate that permission to search <i>file</i> will be granted.					
35147	− z string	Ing True if the length of string string is zero.					
35148	string	True if the string string is not the null string.					
35149	s1 = s2	True if the strings $s1$ and $s2$ are identical.					
35150	s1 != s2	True if the strings $s1$ and $s2$ are not identical.					
35151	n1 – eq n2	True if the integers $n1$ and $n2$ are algebraically equal.					
35152	n1 – ne n2	True if the integers $n1$ and $n2$ are not algebraically equal.					
35153	n1 – gt n2	True if the integer $n1$ is algebraically greater than the integer $n2$.					
35154	n1 – ge n2	True if the integer $n1$ is algebraically greater than or equal to the integer $n2$.					
35155	n1 – lt n2	True if the integer $n1$ is algebraically less than the integer $n2$.					
35156	n1 – le n2	True if the integer $n1$ is algebraically less than or equal to the integer $n2$.					
35157 XSI 35158 35159	expression1 –a expression2 True if both expression1 and expression2 are true. The –a binary primary is left associative. It has a higher precedence than –o.						
35160 XSI 35161 35162	expression1 – o expression2 True if either expression1 or expression2 is true. The – o binary primary is left associative.						
35163 35164 35165	With the exception of the -h <i>file</i> and -L <i>file</i> primaries, if a <i>file</i> argument is a symbolic link, <i>test</i> shall evaluate the expression by resolving the symbolic link and using the file referenced by the link.						
35166	These prima	ries can be combined with the following operators:					
35167	! expression	True if <i>expression</i> is false.					
35168 XSI 35169	(expression) True if expression is true. The parentheses can be used to alter the normal precedence and associativity.						
35170	The primario	es with two elements of the form:					
35171	-primary_	operator primary_operand					
35172	are known as unary primaries. The primaries with three elements in either of the two forms:						
35173	primary_o	perand -primary_operator primary_operand					
35174	primary_operand primary_operator primary_operand						
35175 35176 35177		as <i>binary primaries</i> . Additional implementation-defined operators and <i>rators</i> may be provided by implementations. They shall be of the form <i>-operator</i> set character of <i>operator</i> is not a digit.					
35178 35179 35180	The algorithm for determining the precedence of the operators and the return value that shall be generated is based on the number of arguments presented to <i>test</i> . (However, when using the "[]" form, the right-bracket final argument shall not be counted in this algorithm.)						
35181	In the follow	ring list, \$1, \$2, \$3, and \$4 represent the arguments presented to <i>test</i> :					
35182	0 arguments	Exit false (1).					

Utilities test

35183	1 argument: Exit true (0) if \$1 is not null; otherwise, exit false.					
35184	2 arguments:	• If \$1 is '!', exit true if \$2 is null, false if \$2 is not null.				
35185 35186		• If \$1 is a unary primary, exit true if the unary test is true, false if the unary test is false.				
35187		 Otherwise, produce unspecified results. 				
35188	3 arguments:	• If \$2 is a binary primary, perform the binary test of \$1 and \$3.				
35189		• If $\$1$ is $'$! $'$, negate the two-argument test of $\$2$ and $\$3$.				
35190 XSI		• If 1 is ' (' and 3 is ')', perform the unary test of 2 .				
35191		 Otherwise, produce unspecified results. 				
35192	4 arguments:	• If \$1 is '!', negate the three-argument test of \$2, \$3, and \$4.				
35193 XSI		\bullet If \$1 is \prime (\prime and \$4 is \prime) \prime , perform the two-argument test of \$2 and \$3.				
35194		Otherwise, the results are unspecified.				
35195	>4 arguments	s: The results are unspecified.				
35196 XSI 35197 35198 35199		On XSI-conformant systems, combinations of primaries and operators shall be evaluated using the precedence and associativity rules described previously. In addition, the string comparison binary primaries '=' and "!=" shall have a higher precedence than any unary primary.				
35200 STDIN						
35201	Not used.					
35202 INPUT 35203	None.					
35204 ENVIR 35205	ONMENT VAI The following	RIABLES g environment variables shall affect the execution of <i>test</i> :				
35206 35207 35208 35209		Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)				
35210 35211		If set to a non-empty string value, override the values of all the other internationalization variables.				
35212 35213 35214		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).				
35215 35216 35217	LC_MESSAGES Determine the locale that should be used to affect the format and contents diagnostic messages written to standard error.					
35218 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.				
35219 ASYNO 35220	CHRONOUS E Default.	VENTS				

test Utilities

```
35221 STDOUT
35222
             Not used.
35223 STDERR
             The standard error shall be used only for diagnostic messages.
35224
35225 OUTPUT FILES
35226
             None.
35227 EXTENDED DESCRIPTION
             None.
35228
35229 EXIT STATUS
             The following exit values shall be returned:
35230
35231
                 expression evaluated to true.
                  expression evaluated to false or expression was missing.
35232
             >1 An error occurred.
35233
35234 CONSEQUENCES OF ERRORS
             Default.
35235
35236 APPLICATION USAGE
             Scripts should be careful when dealing with user-supplied input that could be confused with
35237
             primaries and operators. Unless the application writer knows all the cases that produce input to
35238
35239
             the script, invocations like:
              test "$1" -a "$2"
35240
             should be written as:
35241
35242
             test "$1" && test "$2"
35243
             to avoid problems if a user supplied values such as $1 set to '!' and $2 set to the null string.
35244
             That is, in cases where maximal portability is of concern, replace:
              test expr1 -a expr2
35245
35246
             with:
35247
              test expr1 && test expr2
             and replace:
35248
35249
              test expr1 -o expr2
             with:
35250
35251
             test expr1 | test expr2
             but note that, in test, -a has higher precedence than -o while "&&" and "||" have equal
35252
35253
             precedence in the shell.
             Parentheses or braces can be used in the shell command language to effect grouping.
35254
             Parentheses must be escaped when using sh; for example:
35255
              test \( expr1 -a expr2 \) -o expr3
35256
             This command is not always portable outside XSI-conformant systems. The following form can
35257
             be used instead:
35258
```

Utilities test

```
35259
             ( test expr1 && test expr2 ) || test expr3
             The two commands:
35260
             test "$1"
35261
35262
             test ! "$1"
             could not be used reliably on some historical systems. Unexpected results would occur if such a
35263
             string expression were used and $1 expanded to '!', '(', or a known unary primary. Better
35264
35265
             constructs are:
             test -n "$1"
35266
35267
             test -z "$1"
35268
             respectively.
35269
             Historical systems have also been unreliable given the common construct:
             test "$response" = "expected string"
35270
35271
             One of the following is a more reliable form:
35272
             test "X$response" = "Xexpected string"
35273
             test "expected string" = "$response"
             Note that the second form assumes that expected string could not be confused with any unary
35274
             primary. If expected string starts with '-', '(', '!', or even '=', the first form should be used
35275
             instead. Using the preceding rules without the XSI marked extensions, any of the three
35276
             comparison forms is reliable, given any input. (However, note that the strings are quoted in all
35277
35278
             cases.)
             Because the string comparison binary primaries, '=' and "!=", have a higher precedence than
35279
35280
             any unary primary in the greater than 4 argument case, unexpected results can occur if
35281
             arguments are not properly prepared. For example, in:
             test -d $1 -o -d $2
35282
             If $1 evaluates to a possible directory name of '=', the first three arguments are considered a
35283
35284
             string comparison, which shall cause a syntax error when the second -\mathbf{d} is encountered. One of
             the following forms prevents this; the second is preferred:
35285
             test \( -d "$1" \) -o \( -d "$2" \)
35286
             test -d "$1" || test -d "$2"
35287
35288
             Also in the greater than 4 argument case:
             test "$1" = "bat" -a "$2" = "ball"
35289
35290
             syntax errors occur if $1 evaluates to '(' or '!'. One of the following forms prevents this; the
35291
             third is preferred:
             test "X$1" = "Xbat" -a "X$2" = "Xball"
35292
             test "$1" = "bat" && test "$2" = "ball"
35293
             test "X$1" = "Xbat" && test "X$2" = "Xball"
35294
35295 EXAMPLES
               1. Exit if there are not two or three arguments (two variations):
35296
35297
                  if [ $# -ne 2 -a $# -ne 3 ]; then exit 1; fi
                  if [ $# -lt 2 -o $# -gt 3 ]; then exit 1; fi
35298
```

test Utilities

```
35299
                2. Perform a mkdir if a directory does not exist:
                    test ! -d tempdir && mkdir tempdir
35300
                3. Wait for a file to become non-readable:
35301
                   while test -r thefile
35302
35303
                   do
                         sleep 30
35304
35305
                    done
                   echo '"thefile" is no longer readable'
35306
35307
                4. Perform a command if the argument is one of three strings (two variations):
                    if [ "$1" = "pear" ] || [ "$1" = "grape" ] || [ "$1" = "apple" ]
35308
35309
                    then
35310
                         command
                    fi
35311
                    case "$1" in
35312
35313
                         pear | grape | apple) command ;;
35314
                    esac
35315 RATIONALE
35316
              The KornShell-derived conditional command (double bracket [[]]) was removed from the shell
              command language description in an early proposal. Objections were raised that the real
35317
35318
              problem is misuse of the test command ([), and putting it into the shell is the wrong way to fix
              the problem. Instead, proper documentation and a new shell reserved word (!) are sufficient.
35319
35320
              Tests that require multiple test operations can be done at the shell level using individual
              invocations of the test command and shell logicals, rather than using the error-prone –o flag of
35321
35322
              test.
35323
              XSI-conformant systems support more than four arguments.
              XSI-conformant systems support the combining of primaries with the following constructs:
35324
35325
              expression1 -a expression2
35326
                  True if both expression1 and expression2 are true.
35327
              expression1 -o expression2
35328
                  True if at least one of expression1 and expression2 are true.
35329
              (expression)
35330
                  True if expression is true.
              In evaluating these more complex combined expressions, the following precedence rules are
35331
              used:
35332

    The unary primaries have higher precedence than the algebraic binary primaries.

35333
35334

    The unary primaries have lower precedence than the string binary primaries.

    The unary and binary primaries have higher precedence than the unary string primary.

35335
               • The ! operator has higher precedence than the -a operator, and the -a operator has higher
35336
                 precedence than the -\mathbf{o} operator.
35337
```

The parentheses can be used to alter the normal precedence and associativity.

The -a and -o operators are left associative.

35338

35339

Utilities test

35340 The BSD and System V versions of –f are not the same. The BSD definition was: −**f** file True if *file* exists and is not a directory. 35341 The SVID version (true if the file exists and is a regular file) was chosen for this volume of 35342 35343 IEEE Std 1003.1-2001 because its use is consistent with the $-\mathbf{b}$, $-\mathbf{c}$, $-\mathbf{d}$, and $-\mathbf{p}$ operands (file exists 35344 and is a specific file type). 35345 The -e primary, possessing similar functionality to that provided by the C shell, was added 35346 because it provides the only way for a shell script to find out if a file exists without trying to open the file. Since implementations are allowed to add additional file types, a portable script 35347 cannot use: 35348 test -b foo -o -c foo -o -d foo -o -f foo -o -p foo 35349 to find out if foo is an existing file. On historical BSD systems, the existence of a file could be 35350 determined by: 35351 35352 test -f foo -o -d foo but there was no easy way to determine that an existing file was a regular file. An early proposal 35353 35354 used the KornShell -a primary (with the same meaning), but this was changed to -e because 35355 there were concerns about the high probability of humans confusing the -a primary with the -a binary operator. 35356 The following options were not included in this volume of IEEE Std 1003.1-2001, although they 35357 are provided by some implementations. These operands should not be used by new 35358 35359 implementations for other purposes: -k file True if *file* exists and its sticky bit is set. 35360 −C file True if *file* is a contiguous file. 35361 −V file 35362 True if *file* is a version file. The following option was not included because it was undocumented in most implementations, 35363 35364 has been removed from some implementations (including System V), and the functionality is provided by the shell (see Section 2.6.2 (on page 37). 35365 -l string The length of the string *string*. 35366 The $-\mathbf{b}$, $-\mathbf{c}$, $-\mathbf{g}$, $-\mathbf{p}$, $-\mathbf{u}$, and $-\mathbf{x}$ operands are derived from the SVID; historical BSD does not 35367 35368 provide them. The **-k** operand is derived from System V; historical BSD does not provide it. 35369 On historical BSD systems, test -w directory always returned false because test tried to open the 35370 directory for writing, which always fails. Some additional primaries newly invented or from the KornShell appeared in an early proposal 35371 as part of the conditional command ([[]]): s1 > s2, s1 < s2, str = pattern, str! = pattern, 35372 -ot f2, and f1 -ef f2. They were not carried forward into the test utility when the conditional 35373 command was removed from the shell because they have not been included in the test utility 35374 35375 built into historical implementations of the *sh* utility. The -t file descriptor primary is shown with a mandatory argument because the grammar is 35376 ambiguous if it can be omitted. Historical implementations have allowed it to be omitted, 35377 35378 providing a default of 1. 35379 FUTURE DIRECTIONS

35380

None.

test Utilities

35381 S I	EE ALSO
35382	Section 1.7.1.4 (on page 4), find
35383 C]	HANGE HISTORY First released in Issue 2.
35385 Is 35386	The FUTURE DIRECTIONS section is added.
35387 Is 35388 35389	The -h operand is added for symbolic links, and access permission requirements are clarified for the -r, -w, and -x operands to align with the IEEE P1003.2b draft standard.
35390	The normative text is reworded to avoid use of the term "must" for application requirements.
35391	The $-\mathbf{L}$ and $-\mathbf{S}$ operands are added for symbolic links and sockets.
35392 35393 35394	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/38 is applied, adding XSI margin marking and shading to a line in the OPERANDS section referring to the use of parentheses as arguments to the <i>test</i> utility.

Utilities time

35395 **NAME**

35401

35402 35403

35404 35405

35406

35407 35408

35409

35410

35411

35412

35413 35414

35415

35416

35396 time — time a simple command

35397 SYNOPSIS

```
35398 UP time [-p] utility [argument...]
35399
```

35400 **DESCRIPTION**

The *time* utility shall invoke the utility named by the *utility* operand with arguments supplied as the *argument* operands and write a message to standard error that lists timing statistics for the utility. The message shall include the following information:

- The elapsed (real) time between invocation of utility and its termination.
- The User CPU time, equivalent to the sum of the *tms_utime* and *tms_cutime* fields returned by the *times*() function defined in the System Interfaces volume of IEEE Std 1003.1-2001 for the process in which *utility* is executed.
- The System CPU time, equivalent to the sum of the *tms_stime* and *tms_cstime* fields returned by the *times*() function for the process in which *utility* is executed.

The precision of the timing shall be no less than the granularity defined for the size of the clock tick unit on the system, but the results shall be reported in terms of standard time units (for example, 0.02 seconds, 00:00:00.02, 1m33.75s, 365.21 seconds), not numbers of clock ticks.

When *time* is used as part of a pipeline, the times reported are unspecified, except when it is the sole command within a grouping command (see Section 2.9.4.1 (on page 52)) in that pipeline. For example, the commands on the left are unspecified; those on the right report on utilities **a** and **c**, respectively:

35419 OPTIONS

The *time* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

35422 The following option shall be supported:

35423 — **p** Write the timing output to standard error in the format shown in the STDERR section.

35425 **OPERANDS**

35426 The following operands shall be supported:

35427 *utility* The name of a utility that is to be invoked. If the *utility* operand names any of the special built-in utilities in Section 2.14 (on page 64), the results are undefined.

35429 argument Any string to be supplied as an argument when invoking the utility named by the utility operand.

35431 **STDIN**

35432 Not used.

35433 INPUT FILES

35434 None.

time Utilities

	85435 ENVIR 85436	55 ENVIRONMENT VARIABLES 66 The following environment variables shall affect the execution of <i>time</i> :						
3	35437 35438 35439 35440	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)					
	35441 35442	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.					
3	35443 35444 35445	LC_CTYPE	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).					
3	35446 35447 35448	LC_MESSA	GES Determine the locale that should be used to affect the format and contents of diagnostic and informative messages written to standard error.					
	35449	LC_NUMER						
3	35450		Determine the locale for numeric formatting.					
3	35451 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .					
3	85452 85453 85454	PATH	Determine the search path that shall be used to locate the utility to be invoked; se the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environmen Variables.					
	85455 ASYNC 85456	CHRONOUS I Default.	EVENTS					
	85457 STDOU 85458	J T Not used.						
3	35459 STDER	R						
	35460 35461	The standard error shall be used to write the timing statistics. If $-\mathbf{p}$ is specified, the following format shall be used in the POSIX locale:						
	35462 35463	<pre>"real %f\nuser %f\nsys %f\n", <real seconds="">, <user seconds="">,</user></real></pre>						
3 3 3	85464 85465 85466 85467 85468 85469	where each floating-point number shall be expressed in seconds. The precision used may be less than the default six digits of %f, but shall be sufficiently precise to accommodate the size of the clock tick on the system (for example, if there were 60 clock ticks per second, at least two digits shall follow the radix character). The number of digits following the radix character shall be no less than one, even if this always results in a trailing zero. The implementation may append white space and additional information following the format shown here.						
3	35470 OUTPU							
3	35471 None.							
	35472 EXTENDED DESCRIPTION 35473 None.							
3	85474 EXIT S 85475 85476	STATUS If the <i>utility</i> utility is invoked, the exit status of <i>time</i> shall be the exit status of <i>utility</i> ; otherwise, the <i>time</i> utility shall exit with one of the following values:						
		4 405	Let all the collections of the collection of the					

An error occurred in the *time* utility.

35477

1-125

Utilities time

- 35478 126 The utility specified by *utility* was found but could not be invoked.
- 35479 127 The utility specified by *utility* could not be found.

35480 CONSEQUENCES OF ERRORS

35481 Default.

35482 APPLICATION USAGE

The *command, env, nice, nohup, time,* 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.

35493 EXAMPLES

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

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

35498 time sh -c 'complex-command-line'

35499 RATIONALE

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

- The underlying CPU definitions from the System Interfaces volume of IEEE Std 1003.1-2001 are vague, so the numeric output could not be compared accurately between systems or even between invocations.
- The creation of portable benchmark programs was outside the scope this volume of IEEE Std 1003.1-2001.

However, *time* does fit in the scope of user portability. Human judgement can be applied to the analysis of the output, and it could be very useful in hands-on debugging of applications or in providing subjective measures of system performance. Hence it has been included in this volume of IEEE Std 1003.1-2001.

The default output format has been left unspecified because historical implementations differ greatly in their style of depicting this numeric output. The $-\mathbf{p}$ option was invented to provide scripts with a common means of obtaining this information.

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

The term *utility* is used, rather than *command*, to highlight the fact that shell compound commands, pipelines, special built-ins, and so on, cannot be used directly. However, *utility*

time Utilities

includes user application programs and shell scripts, not just the standard utilities.

FUTURE DIRECTIONS
None.

SEE ALSO
Chapter 2 (on page 29), sh, the System Interfaces volume of IEEE Std 1003.1-2001, times()

CHANGE HISTORY
First released in Issue 2.

Issue 6

This utility is marked as part of the User Portability Utilities option.

touch **Utilities**

35533 **NAME** 35534 touch — change file access and modification times 35535 SYNOPSIS touch [-acm] [-r ref file | -t time] file... 35536 35537 **DESCRIPTION** The touch utility shall change the modification times, access times, or both of files. The 35538 modification time shall be equivalent to the value of the *st_mtime* member of the **stat** structure 35539 for a file, as described in the System Interfaces volume of IEEE Std 1003.1-2001; the access time 35540 35541 shall be equivalent to the value of *st_atime*. The time used can be specified by the -t time option-argument, the corresponding time fields of 35542 the file referenced by the -r ref_file option-argument, or the date_time operand, as specified in the 35543 35544 following sections. If none of these are specified, touch shall use the current time (the value returned by the equivalent of the time() function defined in the System Interfaces volume of 35545 IEEE Std 1003.1-2001). 35546 For each file operand, touch shall perform actions equivalent to the following functions defined 35547 in the System Interfaces volume of IEEE Std 1003.1-2001: 35548 1. If file does not exist, a creat() function call is made with the file operand used as the path 35549 argument and the value of the bitwise-inclusive OR of S_IRUSR, S_IWUSR, S_IRGRP, 35550 S_IWGRP, S_IROTH, and S_IWOTH used as the *mode* argument. 35551 2. The *utime*() function is called with the following arguments: 35552 The *file* operand is used as the *path* argument. 35553 The **utimbuf** structure members actime and modtime are determined as described in 35554 the OPTIONS section. 35555 35556 OPTIONS

35557

35558

The touch utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported: 35559

35560 35561	- a	Change the access time of <i>file</i> . Do not change the modification time unless $-\mathbf{m}$ is also specified.					
35562 35563	- c	Do not create a specified <i>file</i> if it does not exist. Do not write any diagnostic messages concerning this condition.					
35564 35565	- m	Change the modification time of <i>file</i> . Do not change the access time unless $-\mathbf{a}$ is also specified.					
35566 35567	- r ref_file	Use the corresponding time of the file named by the pathname <i>ref_file</i> instead of the current time.					
35568 35569	−t time	Use the specified <i>time</i> instead of the current time. The option-argument shall be a decimal number of the form:					
35570		[[CC]YY]MMDDhhmm[.SS]					
35571		where each two digits represents the following:					
35572		MM The month of the year [01,12].					
35573		DD The day of the month [01,31].					

touch Utilities

The minute of the hour 00,59 .	35574		hh	The hour of the day [00,23].						
S5577 YY The second two digits of the year.	35575		mm	The minute of the hour [00,59].						
SS	35576		CC	The first two digits of the year (the century).						
Both CC and YY shall be optional. If neither is given, the current year shall be assumed. If YY is specified, but CC is not, CC shall be derived as follows: FYY is: CC becomes:	35577		YY							
assumed. If YY is specified, but CC is not, CC shall be derived as follows: If YY is: CC becomes: [69.99] 19 19 19 19 19 19 19	35578		SS	The second of the minute [00,60].						
assumed. If YY is specified, but CC is not, CC shall be derived as follows: If YY is: CC becomes: [69.99] 19 19 19 19 19 19 19	35579		Both Co	C and YY shall be optional. If neither is given, the current year shall be						
Record R	35580									
Note:	35581			If YY is: CC becomes:						
Note: It is expected that in a future version of IEEE Std 1003.1-2001 the default century inferred from a 2-digit year will change. (This would apply to all commands accepting a 2-digit year as input.) The resulting time shall be affected by the value of the TZ environment variable. If the resulting time walue precedes the Epoch, touch shall exit immediately with an assistance of valid times past the Epoch is implementation-defined, but it shall extend to at least the time 0 hours, 0 minutes, 0 seconds, January 1, 2038, Coordinated Universal Time. Some implementations may not be able to on the state of the transport of the transport of valid times past the Epoch is implementation-defined, but it shall extend to at least the time 0 hours, 0 minutes, 0 seconds, January 1, 2038, Coordinated Universal Time. Some implementations may not be able to one present dates beyond January 18, 2038, because they use signed int as a time holder. The range for SS is [00,60] rather than [00,59] because of leap seconds. If SS is 60, and the resulting time, as affected by the TZ environment variable, does not refer to a leap second, the resulting time shall be one second after a time where SS is 59. If SS is not given a value, it is assumed to be zero. If neither the ¬a nor ¬m options were specified, touch shall behave as if both the ¬a and ¬m options were specified. S600 OPERANDS The following operands shall be supported: A pathname of a file whose times shall be modified. S601 INPUT FILES S602 FIVIRONMENT VARIABLES The following environment variables shall affect the execution of touch: (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization variables for the precedence of internationalization variables used to determine the values of locale categories.) LANG Provide a default value for the internationalization variables used to determine the values of locale categories.)	35582									
century inferred from a 2-digit year will change. (This would apply to all commands accepting a 2-digit year as input.) The resulting time shall be affected by the value of the TZ environment variable. If the resulting time value precedes the Epoch, touch shall exit immediately with an error status. The range of valid times past the Epoch is implementation-defined, but it shall extend to at least the time 0 hours, 0 minutes, 0 seconds, January 1, 2038, Coordinated Universal Time. Some implementations may not be able to represent dates beyond January 18, 2038, because they use signed int as a time holder. The range for SS is [00,60] rather than [00,59] because of leap seconds. If SS is 60, and the resulting time, as affected by the TZ environment variable, does not refer to a leap second, the resulting time shall be one second after a time where SS is 59. If SS is not given a value, it is assumed to be zero. If neither the —a nor —m options were specified, touch shall behave as if both the —a and —m options were specified. S600 OPERANDS S600 OPERANDS S600 The following operands shall be supported: S600 INPUT FILES S600 None. S600 None. S600 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of touch: S600 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) S611 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables.	35583			[00,68] 20						
the resulting time value precedes the Epoch, touch shall exit immediately with an error status. The range of valid times past the Epoch is implementation-defined, but it shall extend to at least the time 0 hours, 0 minutes, 0 seconds, January 1, 2038, Coordinated Universal Time. Some implementations may not be able to represent dates beyond January 18, 2038, because they use signed int as a time holder. The range for SS is [00,60] rather than [00,59] because of leap seconds. If SS is 60, and the resulting time, as affected by the TZ environment variable, does not refer to a leap second, the resulting time shall be one second after a time where SS is 59. If SS is not given a value, it is assumed to be zero. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. S600 OPERANDS The following operands shall be supported: The following operands shall be supported: The following operands shall be supported: The following environment variables shall affect the execution of touch: The following environment variables shall affect the execution of touch: LANG FEVIRONMENT VARIABLES The following environment variables shall affect the execution of touch: LANG Forvide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables.	35585		Note:	century inferred from a 2-digit year will change. (This would apply to all						
the resulting time value precedes the Epoch, touch shall exit immediately with an error status. The range of valid times past the Epoch is implementation-defined, but it shall extend to at least the time 0 hours, 0 minutes, 0 seconds, January 1, 2038, Coordinated Universal Time. Some implementations may not be able to represent dates beyond January 18, 2038, because they use signed int as a time holder. The range for SS is [00,60] rather than [00,59] because of leap seconds. If SS is 60, and the resulting time, as affected by the TZ environment variable, does not refer to a leap second, the resulting time shall be one second after a time where SS is 59. If SS is not given a value, it is assumed to be zero. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. S600 OPERANDS The following operands shall be supported: The following operands shall be supported: The following operands shall be supported: The following environment variables shall affect the execution of touch: The following environment variables shall affect the execution of touch: LANG FEVIRONMENT VARIABLES The following environment variables shall affect the execution of touch: LANG Forvide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables.	35587		The resi	ulting time shall be affected by the value of the TZ environment variable. If						
but it shall extend to at least the time 0 hours, 0 minutes, 0 seconds, January 1, 35591 2038, Coordinated Universal Time. Some implementations may not be able to represent dates beyond January 18, 2038, because they use signed int as a time holder. 35593 holder. 35594 The range for SS is [00,60] rather than [00,59] because of leap seconds. If SS is 60, 35595 and the resulting time, as affected by the TZ environment variable, does not refer to a leap second, the resulting time shall be one second after a time where SS is 59. If SS is not given a value, it is assumed to be zero. 35598 If neither the —a nor —m options were specified, touch shall behave as if both the —a and —m options were specified. 35600 OPERANDS 35601 The following operands shall be supported: 35602 file A pathname of a file whose times shall be modified. 35603 STDIN 35604 Not used. 35605 INPUT FILES 35606 None. 35607 ENVIRONMENT VARIABLES 35608 The following environment variables shall affect the execution of touch: 35609 LANG Provide a default value for the internationalization variables that are unset or null. 35609 See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 35613 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables.										
2038, Coordinated Universal Time. Some implementations may not be able to represent dates beyond January 18, 2038, because they use signed int as a time holder. 35593 holder. 35594 The range for SS is [00,60] rather than [00,59] because of leap seconds. If SS is 60, 35595 and the resulting time, as affected by the TZ environment variable, does not refer to a leap second, the resulting time shall be one second after a time where SS is 59. If SS is not given a value, it is assumed to be zero. 35598 If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. 35600 OPERANDS 35601 The following operands shall be supported: 35602 file A pathname of a file whose times shall be modified. 35603 STDIN 35604 Not used. 35605 INPUT FILES 35606 None. 35607 ENVIRONMENT VARIABLES 35608 The following environment variables shall affect the execution of touch: 35609 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 11 Internationalization Variables of the precedence of internationalization variables used to determine the values of locale categories.) 35613 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 35615 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as										
represent dates beyond January 18, 2038, because they use signed int as a time holder. The range for SS is [00,60] rather than [00,59] because of leap seconds. If SS is 60, and the resulting time, as affected by the TZ environment variable, does not refer to a leap second, the resulting time shall be one second after a time where SS is 59. If SS is not given a value, it is assumed to be zero. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. The following operands shall be supported: The following operands shall be supported: A pathname of a file whose times shall be modified. The following operands shall be supported: The following operands shall be one second after a time where SS is 59. The following operands operands shall be one second after a time where SS is 59. The following operands operands shall be one second after a time where SS is 59. The following operands operands shall be one secon				· ·						
The range for SS is [00,60] rather than [00,59] because of leap seconds. If SS is 60, and the resulting time, as affected by the TZ environment variable, does not refer to a leap second, the resulting time shall be one second after a time where SS is 59. If SS is not given a value, it is assumed to be zero. S5598 If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. S5600 OPERANDS S5601 The following operands shall be supported: S5602 file A pathname of a file whose times shall be modified. S5603 STDIN S5604 Not used. S5605 INPUT FILES S5606 None. S5607 ENVIRONMENT VARIABLES S5608 The following environment variables shall affect the execution of touch: S5609 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 11 Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) S5613 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. S5615 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as	35592		represe							
and the resulting time, as affected by the TZ environment variable, does not refer to a leap second, the resulting time shall be one second after a time where SS is 59. If SS is not given a value, it is assumed to be zero. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified. I	35593		holder.							
to a leap second, the resulting time shall be one second after a time where SS is 59. If SS is not given a value, it is assumed to be zero. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. The following operands shall be supported: The following operands shall be supported: The following operands shall be supported: The following operands shall be modified. The following operands shall affect the execution of touch: The following operands shall affect the execution of touch: The following operands shall affect the execution of touch: The following operands shall affect the execution of touch: The following operands shall affect the execution of touch: The following operands shall affect the execution of touch: The following operands shall be supported: The following oper										
If SS is not given a value, it is assumed to be zero. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. If neither the -a nor -m options were specified, touch shall behave as if both the -a and -m options were specified. It following operands shall be supported: It following operands shall be supported: It following operands shall be supported: It following operands shall be were specified, touch shall behave as if both the -a and -m options were specified, touch shall behave as if both the -a and -m options were specified, touch shall behave as if both the -a and -m options were specified, touch shall behave as if both the -a and -m options were specified, touch shall behave as if both the -a and -m options were specified, touch shall behave as if both the -a and -m options were specified. It following operands shall be supported: It followi										
options were specified. SF600 OPERANDS SF601 The following operands shall be supported: SF602 file A pathname of a file whose times shall be modified. SF603 STDIN SF604 Not used. SF605 INPUT FILES SF606 None. SF607 ENVIRONMENT VARIABLES SF608 The following environment variables shall affect the execution of touch: SF609 LANG Provide a default value for the internationalization variables that are unset or null. SF601 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) SF613 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. SF615 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as										
The following operands shall be supported: The following operands shall be supported: File A pathname of a file whose times shall be modified.										
35602 file A pathname of a file whose times shall be modified. 35603 STDIN 35604 Not used. 35605 INPUT FILES 35606 None. 35607 ENVIRONMENT VARIABLES 35608 The following environment variables shall affect the execution of touch: 35609 LANG Provide a default value for the internationalization variables that are unset or null. 35610 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 35611 Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 35613 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 35615 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as	35600 OPERA	NDS								
35603 STDIN 35604 Not used. 35605 INPUT FILES 35606 None. 35607 ENVIRONMENT VARIABLES 35608 The following environment variables shall affect the execution of touch: 35609 LANG Provide a default value for the internationalization variables that are unset or null. 35610 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 35611 Internationalization Variables for the precedence of internationalization variables 35612 used to determine the values of locale categories.) 35613 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 35615 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as	35601	The following	ng operan	ds shall be supported:						
Not used. 35605 INPUT FILES 35606 None. 35607 ENVIRONMENT VARIABLES 35608 The following environment variables shall affect the execution of touch: 35609 LANG Provide a default value for the internationalization variables that are unset or null. 35610 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 35611 Internationalization Variables for the precedence of internationalization variables 35612 used to determine the values of locale categories.) 35613 LC_ALL If set to a non-empty string value, override the values of all the other 35614 internationalization variables. 35615 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as	35602	file	A pathr	ame of a file whose times shall be modified.						
35605 INPUT FILES 35606 None. 35607 ENVIRONMENT VARIABLES 35608 The following environment variables shall affect the execution of touch: 35609 LANG Provide a default value for the internationalization variables that are unset or null. 35610 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 35611 Internationalization Variables for the precedence of internationalization variables 35612 used to determine the values of locale categories.) 35613 LC_ALL If set to a non-empty string value, override the values of all the other 35614 internationalization variables. 35615 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as		NT . I								
None. 35607 ENVIRONMENT VARIABLES 35608 The following environment variables shall affect the execution of touch: 35609 LANG Provide a default value for the internationalization variables that are unset or null. 35610 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 35611 Internationalization Variables for the precedence of internationalization variables 35612 used to determine the values of locale categories.) 35613 LC_ALL If set to a non-empty string value, override the values of all the other 35614 internationalization variables. 35615 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as										
The following environment variables shall affect the execution of <i>touch</i> : LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as										
(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as										
internationalization variables. LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as	35610 35611	LANG	(See tl Internat	ne Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, ionalization Variables for the precedence of internationalization variables						
		LC_ALL								
		LC_CTYPE								

Utilities touch

35617 arguments). LC_MESSAGES 35618 35619 Determine the locale that should be used to affect the format and contents of 35620 diagnostic messages written to standard error. 35621 XSI NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 35622 TZDetermine the timezone to be used for interpreting the *time* option-argument. If TZ 35623 is unset or null, an unspecified default timezone shall be used. 35624 ASYNCHRONOUS EVENTS 35625 Default. 35626 STDOUT Not used. 35627 35628 STDERR 35629 The standard error shall be used only for diagnostic messages. 35630 OUTPUT FILES None. 35631 35632 EXTENDED DESCRIPTION None. 35633 35634 EXIT STATUS 35635 The following exit values shall be returned: The utility executed successfully and all requested changes were made. 35636 An error occurred. 35637 35638 CONSEQUENCES OF ERRORS 35639 Default. 35640 APPLICATION USAGE 35641 The interpretation of time is taken to be seconds since the Epoch (see the Base Definitions volume 35642 of IEEE Std 1003.1-2001, Section 4.14, Seconds Since the Epoch). It should be noted that 35643 implementations conforming to the System Interfaces volume of IEEE Std 1003.1-2001 do not take leap seconds into account when computing seconds since the Epoch. When SS=60 is used, 35644 the resulting time always refers to 1 plus *seconds since the Epoch* for a time when *SS*=59. 35645 Although the -t time option-argument specifies values in 1969, the access time and modification 35646 time fields are defined in terms of seconds since the Epoch (00:00:00 on 1 January 1970 UTC). 35647 Therefore, depending on the value of TZ when touch is run, there is never more than a few valid 35648 hours in 1969 and there need not be any valid times in 1969. 35649 One ambiguous situation occurs if -t time is not specified, -r ref_file is not specified, and the first 35650 operand is an eight or ten-digit decimal number. A portable script can avoid this problem by 35651 35652 using: touch -- file 35653 35654 or. touch ./file 35655

in this case.

35656

touch Utilities

35657 EXAMPLES

35658 None.

35659 RATIONALE

35660

35661

35662

35663

35664

35665 35666

35667

35668

35669

35670

35671

35672

35673

35674

35675

35676 35677

35678 35679

35680

The functionality of *touch* is described almost entirely through references to functions in the System Interfaces volume of IEEE Std 1003.1-2001. In this way, there is no duplication of effort required for describing such side effects as the relationship of user IDs to the user database, permissions, and so on.

There are some significant differences between the *touch* utility in this volume of IEEE Std 1003.1-2001 and those in System V and BSD systems. They are upwards-compatible for historical applications from both implementations:

- 1. In System V, an ambiguity exists when a pathname that is a decimal number leads the operands; it is treated as a time value. In BSD, no *time* value is allowed; files may only be *touch*ed to the current time. The -t *time* construct solves these problems for future conforming applications (note that the -t option is not historical practice).
- 2. The inclusion of the century digits, *CC*, is also new. Note that a ten-digit *time* value is treated as if *YY*, and not *CC*, were specified. The caveat about the range of dates following the Epoch was included as recognition that some implementations are not able to represent dates beyond 18 January 2038 because they use **signed int** as a time holder.

The **–r** option was added because several comments requested this capability. This option was named **–f** in an early proposal, but was changed because the **–f** option is used in the BSD version of *touch* with a different meaning.

At least one historical implementation of *touch* incremented the exit code if –c was specified and the file did not exist. This volume of IEEE Std 1003.1-2001 requires exit status zero if no errors occur.

35681 FUTURE DIRECTIONS

35682 Applications should use the $-\mathbf{r}$ or $-\mathbf{t}$ options.

35683 **SEE ALSO**

date, the System Interfaces volume of IEEE Std 1003.1-2001, creat(), time(), utime(), the Base Definitions volume of IEEE Std 1003.1-2001, <sys/stat.h>

35686 CHANGE HISTORY

35687 First released in Issue 2.

35688 Issue 6

35689 The obsolescent *date_time* operand is removed.

The Open Group Corrigendum U027/1 is applied. This extends the range of valid time past the Epoch to at least the time 0 hours, 0 minutes, 0 seconds, January 1, 2038, Coordinated Universal Time. This is a new requirement on POSIX implementations.

The range for seconds is changed from [00,61] to [00,60] to align with the ISO/IEC 9899:1999 standard, and to allow for positive leap seconds.

Utilities tput

35695 **NAME** 35696 tput — change terminal characteristics 35697 SYNOPSIS tput [-T type] operand... 35698 UP 35699 35700 **DESCRIPTION** The tput utility shall display terminal-dependent information. The manner in which this 35701 information is retrieved is unspecified. The information displayed shall clear the terminal screen, 35702 initialize the user's terminal, or reset the user's terminal, depending on the operand given. The 35703 35704 exact consequences of displaying this information are unspecified. 35705 OPTIONS 35706 The tput utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines. 35707 The following option shall be supported: 35708 -T type Indicate the type of terminal. If this option is not supplied and the TERM variable 35709 35710 is unset or null, an unspecified default terminal type shall be used. The setting of 35711 *type* shall take precedence over the value in *TERM*. 35712 OPERANDS The following strings shall be supported as operands by the implementation in the POSIX locale: 35713 35714 clear Display the clear-screen sequence. init Display the sequence that initializes the user's terminal in an implementation-35715 35716 defined manner. 35717 reset Display the sequence that resets the user's terminal in an implementation-defined 35718 manner. If a terminal does not support any of the operations described by these operands, this shall not 35719 35720 be considered an error condition. 35721 **STDIN** 35799 Not used. 35723 INPUT FILES 35724 None. 35725 ENVIRONMENT VARIABLES 35726 The following environment variables shall affect the execution of *tput*: LANG Provide a default value for the internationalization variables that are unset or null. 35727 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 35728 Internationalization Variables for the precedence of internationalization variables 35729 used to determine the values of locale categories.) 35730 LC_ALL If set to a non-empty string value, override the values of all the other 35731 internationalization variables. 35732 Determine the locale for the interpretation of sequences of bytes of text data as 35733 LC_CTYPE characters (for example, single-byte as opposed to multi-byte characters in 35734 arguments). 35735

diagnostic messages written to standard error.

Determine the locale that should be used to affect the format and contents of

LC_MESSAGES

35736

35737 35738 **tput** Utilities

35739 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .							
35740 35741	TERM Determine the terminal type. If this variable is unset or null, and if the -T option is not specified, an unspecified default terminal type shall be used.								
35742 ASYNCHRONOUS EVENTS 35743 Default.									
35744 STDOU	35744 STDOUT								
35745 35746	If standard output is a terminal device, it may be used for writing the appropriate sequence to clear the screen or reset or initialize the terminal. If standard output is not a terminal device,								
35747	undefined results occur.								
35748 STDER									
35749		rd error shall be used only for diagnostic messages.							
35750 OUTPU 35751	None.								
35752 EXTEN	DED DESCI	RIPTION							
35753	None.								
35754 EXIT S		to a soft and and all he make one de							
35755		ing exit values shall be returned:							
35756		quested string was written successfully.							
35757	1 Unspe								
35758	2 Usage error.								
35759	3 No information is available about the specified terminal type.								
35760	4 The specified operand is invalid.								
35761		or occurred.							
35762 CONSI 35763	•	of ERRORS e operands is not available for the terminal, <i>tput</i> continues processing the remaining							
35764	operands.								
35765 APPLIC									
35766 35767		ence between resetting and initializing a terminal is left unspecified, as they vary ed on hardware types. In general, resetting is a more severe action.							
35768		inals use control characters to perform the stated functions, and on such terminals it e sense to use <i>tput</i> to store the initialization strings in a file or environment variable							
35769 35770		e. However, because other terminals might rely on system calls to do this work, the							
35771		utput cannot be used in a portable manner, such as the following non-portable							
35772	constructs:	Neural along							
35773 35774		=`tput clear` et mailx -s "Wake Up" ddg							
35775 EXAMI	PLES								
35776 35777		lize the terminal according to the type of terminal in the environmental variable <i>f</i> . This command can be included in a .profile file.							
35778	tput	init							
35779	2. Reset	a 450 terminal.							

Utilities tput

35780 tput -T 450 reset 35781 RATIONALE 35782 The list of operands was reduced to a minimum for the following reasons: 35783 The only features chosen were those that were likely to be used by human users interacting 35784 with a terminal. • Specifying the full terminfo set was not considered desirable, but the standard developers did 35785 35786 not want to select among operands. • This volume of IEEE Std 1003.1-2001 does not attempt to provide applications with 35787 35788 sophisticated terminal handling capabilities, as that falls outside of its assigned scope and intersects with the responsibilities of other standards bodies. 35789 35790 The difference between resetting and initializing a terminal is left unspecified as this varies greatly based on hardware types. In general, resetting is a more severe action. 35791 35792 The exit status of 1 is historically reserved for finding out if a Boolean operand is not set. 35793 Although the operands were reduced to a minimum, the exit status of 1 should still be reserved for the Boolean operands, for those sites that wish to support them. 35794 35795 FUTURE DIRECTIONS 35796 None. 35797 SEE ALSO 35798 stty, tabs 35799 CHANGE HISTORY First released in Issue 4. 35800 35801 **Issue 6** 35802 This utility is marked as part of the User Portability Utilities option.

tr Utilities

```
35803 NAME
35804
              tr — translate characters
35805 SYNOPSIS
              tr [-c | -C][-s] string1 string2
35806
35807
              tr -s [-c | -C] string1
              tr -d [-c | -C] string1
35808
              tr -ds [-c | -C] string1 string2
35809
35810 DESCRIPTION
              The tr utility shall copy the standard input to the standard output with substitution or deletion
35811
35812
              of selected characters. The options specified and the string1 and string2 operands shall control
              translations that occur while copying characters and single-character collating elements.
35813
35814 OPTIONS
35815
              The tr utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2,
35816
              Utility Syntax Guidelines.
35817
              The following options shall be supported:
                           Complement the set of values specified by string1. See the EXTENDED
35818
              -с
35819
                           DESCRIPTION section.
              -\mathbf{C}
                           Complement the set of characters specified by string1. See the EXTENDED
35820
                           DESCRIPTION section.
35821
              -\mathbf{d}
                           Delete all occurrences of input characters that are specified by string1.
35822
                           Replace instances of repeated characters with a single character, as described in the
35823
              -s
                           EXTENDED DESCRIPTION section.
35824
35825 OPERANDS
              The following operands shall be supported:
35826
              string1, string2
35827
35828
                           Translation control strings. Each string shall represent a set of characters to be
                           converted into an array of characters used for the translation. For a detailed
35829
                           description of how the strings are interpreted, see the EXTENDED DESCRIPTION
35830
                           section.
35831
35832 STDIN
35833
              The standard input can be any type of file.
35834 INPUT FILES
35835
              None
35836 ENVIRONMENT VARIABLES
              The following environment variables shall affect the execution of tr:
35837
              LANG
                           Provide a default value for the internationalization variables that are unset or null.
35838
                           (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
35839
                           Internationalization Variables for the precedence of internationalization variables
35840
                           used to determine the values of locale categories.)
35841
35842
              LC ALL
                           If set to a non-empty string value, override the values of all the other
                           internationalization variables.
35843
35844
              LC_COLLATE
35845
                           Determine the locale for the behavior of range expressions and equivalence classes.
```

Utilities tr

35853 ASYNCHRONOUS EVENTS

35854 Default.

35855 **STDOUT**

35856 The *tr* output shall be identical to the input, with the exception of the specified transformations.

35857 STDERR

35862

35863

35864 35865

35873

35874

35875

35876

35877

35878

35879

35880

35881

35882

35883

35884

35885

35886

35887

35888

35889 35890

35858 The standard error shall be used only for diagnostic messages.

35859 OUTPUT FILES

35860 None.

35861 EXTENDED DESCRIPTION

C-C

The operands *string1* and *string2* (if specified) define two arrays of characters. The constructs in the following list can be used to specify characters or single-character collating elements. If any of the constructs result in multi-character collating elements, *tr* shall exclude, without a diagnostic, those multi-character elements from the resulting array.

Any character not described by one of the conventions below shall represent itself. 35866 character \octal Octal sequences can be used to represent characters with specific coded values. An 35867 octal sequence shall consist of a backslash followed by the longest sequence of one, 35868 35869 two, or three-octal-digit characters (01234567). The sequence shall cause the value whose encoding is represented by the one, two, or three-digit octal integer to be 35870 placed into the array. If the size of a byte on the system is greater than nine bits, the 35871 valid escape sequence used to represent a byte is implementation-defined. Multi-35872

including the leading '\' for each byte.

\character The backslash-escape sequences in the Base Definitions volume of IEEE Std 1003.1-2001, Table 5-1, Escape Sequences and Associated Actions ('\\', '\a', '\b', '\f', '\r', '\t', '\v') shall be supported. The results of using any other character, other than an octal digit, following the backslash are

unspecified.

In the POSIX locale, this construct shall represent the range of collating elements between the range endpoints (as long as neither endpoint is an octal sequence of the form \octal), inclusive, as defined by the collation sequence. The characters or collating elements in the range shall be placed in the array in ascending collation sequence. If the second endpoint precedes the starting endpoint in the collation sequence, it is unspecified whether the range of collating elements is empty, or this construct is treated as invalid. In locales other than the POSIX locale, this construct

byte characters require multiple, concatenated escape sequences of this type,

has unspecified behavior.

If either or both of the range endpoints are octal sequences of the form \octal, this shall represent the range of specific coded values between the two range

endpoints, inclusive.

tr Utilities

35891 35892 35893	[:class:]	Represents all characters belonging to the defined character class, as defined by the current setting of the LC_CTYPE locale category. The following character class names shall be accepted when specified in $string1$:								
35894 35895		alnum blank digit lower punct upper alpha cntrl graph print space xdigit								
35896 XSI 35897 35898		In addition, character class expressions of the form [:name:] shall be recognized those locales where the name keyword has been given a charclass definition <i>LC_CTYPE</i> category.								
35899 35900 35901 35902 35903 35904 35905 35906 35907 35908 35909 35910		When both the -d and -s options are specified, any of the character class shall be accepted in <i>string2</i> . Otherwise, only character class names lower or are valid in <i>string2</i> and then only if the corresponding character class (uppe lower , respectively) is specified in the same relative position in <i>string1</i> . Suppersification shall be interpreted as a request for case conversion. When [suppers in <i>string1</i> and [suppers] appears in <i>string2</i> , the arrays shall contain characters from the toupper mapping in the <i>LC_CTYPE</i> category of the colocale. When [suppers] appears in <i>string1</i> and [slowers] appears in <i>string2</i> , the shall contain the characters from the tolower mapping in the <i>LC_CTYPE</i> category of the current locale. The first character from each mapping pair shall be array for <i>string1</i> and the second character from each mapping pair shall be array for <i>string2</i> in the same relative position.								
35911 35912		Except for case conversion, the characters specified by a character class expression shall be placed in the array in an unspecified order.								
35913 35914		If the name specified for <i>class</i> does not define a valid character class in the current locale, the behavior is undefined.								
35915 35916 35917 35918 35919 35920	[=equiv=] Represents all characters or collating elements belonging to the same equivaclass as equiv, as defined by the current setting of the LC_COLLATE leading category. An equivalence class expression shall be allowed only in string1, string2 when it is being used by the combined -d and -s options. The charabelonging to the equivalence class shall be placed in the array in an unspectorder.									
35921 35922 35923 35924 35925 35926	[<i>x</i> * <i>n</i>]	Represents <i>n</i> repeated occurrences of the character <i>x</i> . Because this expression is used to map multiple characters to one, it is only valid when it occurs in <i>string2</i> . If <i>n</i> is omitted or is zero, it shall be interpreted as large enough to extend the <i>string2</i> -based sequence to the length of the <i>string1</i> -based sequence. If <i>n</i> has a leading zero, it shall be interpreted as an octal value. Otherwise, it shall be interpreted as a decimal value.								
35927	When the –	- d option is not specified:								
35928 35929 35930	in the sa	• Each input character found in the array specified by <i>string1</i> shall be replaced by the character in the same relative position in the array specified by <i>string2</i> . When the array specified by <i>string2</i> is shorter that the one specified by <i>string1</i> , the results are unspecified.								
35931 35932							ters specified by <i>string1</i> (the set of the current setting of <i>LC_CTYPE</i> ,			

- If the –C option is specified, the complements of the characters specified by *string1* (the set of all characters in the current character set, as defined by the current setting of *LC_CTYPE*, except for those actually specified in the *string1* operand) shall be placed in the array in ascending collation sequence, as defined by the current setting of *LC_COLLATE*.
- If the –c option is specified, the complement of the values specified by *string1* shall be placed in the array in ascending order by binary value.

35933

35934 35935

35936

Utilities tr

Because the order in which characters specified by character class expressions or equivalence
 class expressions is undefined, such expressions should only be used if the intent is to map
 several characters into one. An exception is case conversion, as described previously.

When the $-\mathbf{d}$ option is specified:

- Input characters found in the array specified by *string1* shall be deleted.
- When the –C option is specified with –d, all characters except those specified by *string1* shall be deleted. The contents of *string2* are ignored, unless the –s option is also specified.
- When the -c option is specified with -d, all values except those specified by *string1* shall be deleted. The contents of *string2* shall be ignored, unless the -s option is also specified.
- The same string cannot be used for both the **-d** and the **-s** option; when both options are specified, both *string1* (used for deletion) and *string2* (used for squeezing) shall be required.

When the -s option is specified, after any deletions or translations have taken place, repeated sequences of the same character shall be replaced by one occurrence of the same character, if the character is found in the array specified by the last operand. If the last operand contains a character class, such as the following example:

```
35952 tr -s '[:space:]'
```

the last operand's array shall contain all of the characters in that character class. However, in a case conversion, as described previously, such as:

```
35955 tr -s '[:upper:]' '[:lower:]'
```

the last operand's array shall contain only those characters defined as the second characters in each of the **toupper** or **tolower** character pairs, as appropriate.

35958 An empty string used for *string1* or *string2* produces undefined results.

35959 EXIT STATUS

35940

35941

35942 35943

35944

35945

35946

35947

35948

35949

35950

35951

35953 35954

35966

35969

35970 35971

35973

35974

35975

35976

35960 The following exit values shall be returned:

35961 0 All input was processed successfully.

35962 >0 An error occurred.

35963 CONSEQUENCES OF ERRORS

35964 Default.

35965 APPLICATION USAGE

If necessary, *string1* and *string2* can be quoted to avoid pattern matching by the shell.

If an ordinary digit (representing itself) is to follow an octal sequence, the octal sequence must use the full three digits to avoid ambiguity.

When *string2* is shorter than *string1*, a difference results between historical System V and BSD systems. A BSD system pads *string2* with the last character found in *string2*. Thus, it is possible to do the following:

35972 tr 0123456789 d

which would translate all digits to the letter 'd'. Since this area is specifically unspecified in this volume of IEEE Std 1003.1-2001, both the BSD and System V behaviors are allowed, but a conforming application cannot rely on the BSD behavior. It would have to code the example in the following way:

35977 tr 0123456789 '[d*]'

tr Utilities

It should be noted that, despite similarities in appearance, the string operands used by *tr* are not regular expressions.

Unlike some historical implementations, this definition of the *tr* utility correctly processes NUL characters in its input stream. NUL characters can be stripped by using:

```
35982 tr -d '\000'
```

35983 EXAMPLES

1. The following example creates a list of all words in **file1** one per line in **file2**, where a word is taken to be a maximal string of letters.

```
tr -cs "[:alpha:]" "[\n*]" <file1 >file2
```

2. The next example translates all lowercase characters in **file1** to uppercase and writes the results to standard output.

```
tr "[:lower:]" "[:upper:]" <file1</pre>
```

3. This example uses an equivalence class to identify accented variants of the base character 'e' in **file1**, which are stripped of diacritical marks and written to **file2**.

```
tr "[=e=]" e <file1 >file2
```

35993 RATIONALE

In some early proposals, an explicit option $-\mathbf{n}$ was added to disable the historical behavior of stripping NUL characters from the input. It was considered that automatically stripping NUL characters from the input was not correct functionality. However, the removal of $-\mathbf{n}$ in a later proposal does not remove the requirement that tr correctly process NUL characters in its input stream. NUL characters can be stripped by using tr $-\mathbf{d}$ $^{\prime}$ \000'.

Historical implementations of *tr* differ widely in syntax and behavior. For example, the BSD version has not needed the bracket characters for the repetition sequence. The *tr* utility syntax is based more closely on the System V and XPG3 model while attempting to accommodate historical BSD implementations. In the case of the short *string2* padding, the decision was to unspecify the behavior and preserve System V and XPG3 scripts, which might find difficulty with the BSD method. The assumption was made that BSD users of *tr* have to make accommodations to meet the syntax defined here. Since it is possible to use the repetition sequence to duplicate the desired behavior, whereas there is no simple way to achieve the System V method, this was the correct, if not desirable, approach.

The use of octal values to specify control characters, while having historical precedents, is not portable. The introduction of escape sequences for control characters should provide the necessary portability. It is recognized that this may cause some historical scripts to break.

An early proposal included support for multi-character collating elements. It was pointed out that, while tr does employ some syntactical elements from REs, the aim of tr is quite different; ranges, for example, do not have a similar meaning ("any of the chars in the range matches", versus "translate each character in the range to the output counterpart"). As a result, the previously included support for multi-character collating elements has been removed. What remains are ranges in current collation order (to support, for example, accented characters), character classes, and equivalence classes.

In XPG3 the [:class:] and [=equiv=] conventions are shown with double brackets, as in RE syntax. However, tr does not implement RE principles; it just borrows part of the syntax. Consequently, [:class:] and [=equiv=] should be regarded as syntactical elements on a par with [x*n], which is not an RE bracket expression.

Utilities tr

36022 The standard developers will consider changes to tr that allow it to translate characters between 36023 different character encodings, or they will consider providing a new utility to accomplish this. 36024 On historical System V systems, a range expression requires enclosing square-brackets, such as: tr '[a-z]' '[A-Z]' 36025 However, BSD-based systems did not require the brackets, and this convention is used here to 36026 avoid breaking large numbers of BSD scripts: 36027 tr a-z A-Z 36028 36029 The preceding System V script will continue to work because the brackets, treated as regular characters, are translated to themselves. However, any System V script that relied on "a-z" 36030 representing the three characters 'a', '-', and 'z' have to be rewritten as "az-". 36031 The ISO POSIX-2: 1993 standard had a -c option that behaved similarly to the -C option, but did 36032 not supply functionality equivalent to the -c option specified in IEEE Std 1003.1-2001. This 36033 36034 meant that historical practice of being able to specify $tr - d \ge 00 - 377$ (which would delete all 36035 bytes with the top bit set) would have no effect because, in the C locale, bytes with the values octal 200 to octal 377 are not characters. 36036 The earlier version also said that octal sequences referred to collating elements and could be 36037 placed adjacent to each other to specify multi-byte characters. However, it was noted that this 36038 caused ambiguities because tr would not be able to tell whether adjacent octal sequences were 36039 intending to specify multi-byte characters or multiple single byte characters. 36040 36041 IEEE Std 1003.1-2001 specifies that octal sequences always refer to single byte binary values. **36042 FUTURE DIRECTIONS** 36043 None. 36044 SEE ALSO 36045 sed 36046 CHANGE HISTORY

First released in Issue 2.

36048 **Issue 6**

The -C operand is added, and the description of the -c operand is changed to align with the IEEE P1003.2b draft standard.

The normative text is reworded to avoid use of the term "must" for application requirements.

true Utilities

36052 **NAME** 36053 true — return true value 36054 SYNOPSIS 36055 true 36056 **DESCRIPTION** 36057 The true utility shall return with exit code zero. 36058 OPTIONS 36059 None. 36060 OPERANDS 36061 None. 36062 STDIN Not used. 36064 INPUT FILES 36065 None. **36066 ENVIRONMENT VARIABLES** 36067 None. **36068 ASYNCHRONOUS EVENTS** 36069 Default. 36070 STDOUT 36071 Not used. 36072 STDERR Not used. 36073 **36074 OUTPUT FILES** 36075 None. 36076 EXTENDED DESCRIPTION 36077 None. 36078 EXIT STATUS 36079 Zero. 36080 CONSEQUENCES OF ERRORS 36081 None. **36082 APPLICATION USAGE** 36083 This utility is typically used in shell scripts, as shown in the EXAMPLES section. The special built-in utility: is sometimes more efficient than true. 36084 36085 EXAMPLES This command is executed forever: 36086 36087 while true 36088 36089 command

done

36090

Utilities true

36091 RATIONALE 36092 The true utility has been retained in this volume of IEEE Std 1003.1-2001, even though the shell special built-in: provides similar functionality, because true is widely used in historical scripts 36093 and is less cryptic to novice script readers. 36094 **36095 FUTURE DIRECTIONS** None. 36096 36097 **SEE ALSO** false, Section 2.9 (on page 47) 36098 **36099 CHANGE HISTORY** First released in Issue 2. 36100 36101 **Issue 6** IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/39 is applied, replacing the terms "None" 36102 and "Default" from the STDERR and EXIT STATUS sections, respectively, with terms as defined 36103 36104 in Section 1.11 (on page 20).

tsort Utilities

36105 **NAME** 36106 tsort — topological sort 36107 SYNOPSIS 36108 XSI tsort [file] 36109 36110 DESCRIPTION The tsort utility shall write to standard output a totally ordered list of items consistent with a 36111 partial ordering of items contained in the input. 36112 The application shall ensure that the input consists of pairs of items (non-empty strings) 36113 36114 separated by <blank>s. Pairs of different items indicate ordering. Pairs of identical items indicate presence, but not ordering. 36115 36116 OPTIONS 36117 None. 36118 OPERANDS The following operand shall be supported: 36119 file A pathname of a text file to order. If no file operand is given, the standard input 36120 shall be used. 36121 36122 **STDIN** 36123 The standard input shall be a text file that is used if no *file* operand is given. 36124 INPUT FILES The input file named by the *file* operand is a text file. 36125 36126 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *tsort*: 36127 LANG 36128 Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 36129 36130 Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 36131 36132 LC_ALL If set to a non-empty string value, override the values of all the other 36133 internationalization variables. 36134 LC_CTYPE 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 36135 arguments and input files). 36136 LC_MESSAGES 36137 Determine the locale that should be used to affect the format and contents of 36138 diagnostic messages written to standard error. 36139 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 36140 36141 ASYNCHRONOUS EVENTS Default. 36142 36143 **STDOUT**

ordered input.

36144 36145 The standard output shall be a text file consisting of the order list produced from the partially

Utilities tsort

```
36146 STDERR
36147
             The standard error shall be used only for diagnostic messages.
36148 OUTPUT FILES
36149
             None.
36150 EXTENDED DESCRIPTION
36151
             None.
36152 EXIT STATUS
36153
             The following exit values shall be returned:
36154
                 Successful completion.
36155
             >0 An error occurred.
36156 CONSEQUENCES OF ERRORS
             Default.
36157
36158 APPLICATION USAGE
             The LC_COLLATE variable need not affect the actions of tsort. The output ordering is not
36159
36160
             lexicographic, but depends on the pairs of items given as input.
36161 EXAMPLES
             The command:
36162
36163
             tsort <<EOF
36164
             abccde
36165
             g g
             fgef
36166
             h h
36167
             EOF
36168
36169
             produces the output:
36170
             a
36171
             b
36172
             С
36173
             d
36174
             е
             f
36175
36176
             g
36177
             h
36178 RATIONALE
             None.
36180 FUTURE DIRECTIONS
             None.
36181
36182 SEE ALSO
             None.
36183
36184 CHANGE HISTORY
             First released in Issue 2.
36185
```

The normative text is reworded to avoid use of the term "must" for application requirements.

36186 Issue 6

tty Utilities

36188 NAME 36189 tty — return user's terminal name 36190 SYNOPSIS 36191 tty 36192 **DESCRIPTION** The tty utility shall write to the standard output the name of the terminal that is open as 36193 standard input. The name that is used shall be equivalent to the string that would be returned by 36194 the *ttyname*() function defined in the System Interfaces volume of IEEE Std 1003.1-2001. 36195 36196 OPTIONS 36197 The tty utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines. 36198 36199 OPERANDS 36200 None. 36201 STDIN While no input is read from standard input, standard input shall be examined to determine 36202 whether or not it is a terminal, and, if so, to determine the name of the terminal. 36203 36204 INPUT FILES 36205 None. 36206 ENVIRONMENT VARIABLES 36207 The following environment variables shall affect the execution of *tty*: LANG Provide a default value for the internationalization variables that are unset or null. 36208 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 36209 Internationalization Variables for the precedence of internationalization variables 36210 used to determine the values of locale categories.) 36211 36212 LC_ALL If set to a non-empty string value, override the values of all the other internationalization variables. 36213 Determine the locale for the interpretation of sequences of bytes of text data as LC_CTYPE 36214 characters (for example, single-byte as opposed to multi-byte characters in 36215 arguments). 36216 LC_MESSAGES 36217 Determine the locale that should be used to affect the format and contents of 36218 diagnostic messages written to standard error and informative messages written to 36219 36220 standard output. NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 36221 XSI **36222 ASYNCHRONOUS EVENTS** Default. 36223 **36224 STDOUT** If standard input is a terminal device, a pathname of the terminal as specified by the *ttyname()* 36225 function defined in the System Interfaces volume of IEEE Std 1003.1-2001 shall be written in the 36226 following format: 36227 "%s\n", <terminal name> 36228

36229 36230 Otherwise, a message shall be written indicating that standard input is not connected to a

terminal. In the POSIX locale, the *tty* utility shall use the format:

Utilities tty

36231 "not a tty\n" 36232 STDERR 36233 The standard error shall be used only for diagnostic messages. 36234 OUTPUT FILES 36235 None. 36236 EXTENDED DESCRIPTION 36237 None. 36238 EXIT STATUS 36239 The following exit values shall be returned: 36240 Standard input is a terminal. Standard input is not a terminal. 36241 >1 An error occurred. 36242 **36243 CONSEQUENCES OF ERRORS** 36244 Default. **36245 APPLICATION USAGE** 36246 This utility checks the status of the file open as standard input against that of an implementation-defined set of files. It is possible that no match can be found, or that the match 36247 36248 found need not be the same file as that which was opened for standard input (although they are the same device). 36249 36250 EXAMPLES 36251 None. 36252 RATIONALE 36253 None. 36254 FUTURE DIRECTIONS 36255 None. 36256 SEE ALSO 36257 The System Interfaces volume of IEEE Std 1003.1-2001, isatty(), ttyname() 36258 CHANGE HISTORY First released in Issue 2. 36259 36260 Issue 5 36261 The SYNOPSIS is changed to indicate two forms of the command, with the second form marked as obsolete. This is a clarification and does not change the functionality published in previous 36262 issues. 36263

The obsolescent –**s** option is removed.

36264 Issue 6

type Utilities

36266 **NAME** 36267 type — write a description of command type 36268 SYNOPSIS 36269 XSI type name... 36270 36271 **DESCRIPTION** The type utility shall indicate how each argument would be interpreted if used as a command 36272 36273 36274 OPTIONS 36275 None. 36276 OPERANDS 36277 The following operand shall be supported: A name to be interpreted. 36278 name 36279 **STDIN** Not used. 36280 36281 INPUT FILES 36282 None. 36283 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *type*: 36284 LANG Provide a default value for the internationalization variables that are unset or null. 36285 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 36286 Internationalization Variables for the precedence of internationalization variables 36287 used to determine the values of locale categories.) 36288 LC ALL If set to a non-empty string value, override the values of all the other 36289 internationalization variables. 36290 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 36291 36292 characters (for example, single-byte as opposed to multi-byte characters in 36293 arguments). 36294 LC_MESSAGES 36295 Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. 36296 36297 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. **PATH** Determine the location of *name*, as described in the Base Definitions volume of 36298 IEEE Std 1003.1-2001, Chapter 8, Environment Variables. 36299 36300 ASYNCHRONOUS EVENTS 36301 Default. 36302 **STDOUT** 36303 The standard output of *type* contains information about each operand in an unspecified format.

36304

36305

The information provided typically identifies the operand as a shell built-in, function, alias, or

keyword, and where applicable, may display the operand's pathname.

Utilities type

36306 STDERR 36307 The standard error shall be used only for diagnostic messages. **36308 OUTPUT FILES** 36309 None. 36310 EXTENDED DESCRIPTION 36311 None. 36312 EXIT STATUS 36313 The following exit values shall be returned: 36314 Successful completion. 36315 >0 An error occurred. 36316 CONSEQUENCES OF ERRORS Default. 36317 36318 APPLICATION USAGE 36319 Since type must be aware of the contents of the current shell execution environment (such as the lists of commands, functions, and built-ins processed by hash), it is always provided as a shell 36320 regular built-in. If it is called in a separate utility execution environment, such as one of the 36321 following: 36322 36323 nohup type writer find . -type f | xargs type 36324 36325 it might not produce accurate results. 36326 EXAMPLES None. 36327 36328 RATIONALE 36329 None. **36330 FUTURE DIRECTIONS** 36331 None.

36332 **SEE ALSO**

36334 CHANGE HISTORY

command, hash

First released in Issue 2.

36333

ulimit Utilities

36336 **NAME** 36337 ulimit — set or report file size limit 36338 SYNOPSIS 36339 XSI ulimit [-f][blocks] 36340 36341 **DESCRIPTION** The *ulimit* utility shall set or report the file-size writing limit imposed on files written by the 36342 shell and its child processes (files of any size may be read). Only a process with appropriate 36343 privileges can increase the limit. 36344 36345 OPTIONS The *ulimit* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 36346 12.2, Utility Syntax Guidelines. 36347 The following option shall be supported: 36348 -f Set (or report, if no *blocks* operand is present), the file size limit in blocks. The -f 36349 option shall also be the default case. 36350 36351 OPERANDS The following operand shall be supported: 36352 blocks The number of 512-byte blocks to use as the new file size limit. 36353 36354 **STDIN** Not used. 36355 36356 INPUT FILES None. 36357 36358 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *ulimit*: 36359 LANG Provide a default value for the internationalization variables that are unset or null. 36360 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 36361 Internationalization Variables for the precedence of internationalization variables 36362 used to determine the values of locale categories.) 36363 If set to a non-empty string value, override the values of all the other LC_ALL 36364 internationalization variables. 36365 Determine the locale for the interpretation of sequences of bytes of text data as 36366 LC_CTYPE 36367 characters (for example, single-byte as opposed to multi-byte characters in arguments). 36368 LC_MESSAGES 36369 Determine the locale that should be used to affect the format and contents of 36370 diagnostic messages written to standard error. 36371 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 36372

36373 ASYNCHRONOUS EVENTS

36374 Default.

36375 **STDOUT**

The standard output shall be used when no *blocks* operand is present. If the current number of blocks is limited, the number of blocks in the current limit shall be written in the following format:

Utilities **ulimit**

36379 "%d\n", <number of 512-byte blocks> If there is no current limit on the number of blocks, in the POSIX locale the following format 36380 36381 shall be used: "unlimited\n" 36382 **36383 STDERR** The standard error shall be used only for diagnostic messages. 36384 **36385 OUTPUT FILES** None. 36386 36387 EXTENDED DESCRIPTION None. 36388 36389 EXIT STATUS The following exit values shall be returned: 36390 Successful completion. 36391 >0 A request for a higher limit was rejected or an error occurred. 36392 **36393 CONSEQUENCES OF ERRORS** Default. 36394 36395 APPLICATION USAGE 36396 Since *ulimit* affects the current shell execution environment, it is always provided as a shell regular built-in. If it is called in a separate utility execution environment, such as one of the 36397 36398 following: nohup ulimit -f 10000 36399 env ulimit 10000 36400 it does not affect the file size limit of the caller's environment. 36401 Once a limit has been decreased by a process, it cannot be increased (unless appropriate 36402 privileges are involved), even back to the original system limit. 36403 36404 EXAMPLES Set the file size limit to 51 200 bytes: 36405 ulimit -f 100 36406 36407 RATIONALE None. 36408 36409 FUTURE DIRECTIONS None. 36410 **36411 SEE ALSO** The System Interfaces volume of IEEE Std 1003.1-2001, *ulimit()* 36412 36413 CHANGE HISTORY

First released in Issue 2.

umask Utilities

```
36415 NAME
36416
              umask — get or set the file mode creation mask
36417 SYNOPSIS
36418
              umask [-S] [mask]
36419 DESCRIPTION
              The umask utility shall set the file mode creation mask of the current shell execution
36420
              environment (see Section 2.12 (on page 61)) to the value specified by the mask operand. This
36421
              mask shall affect the initial value of the file permission bits of subsequently created files. If umask
36422
              is called in a subshell or separate utility execution environment, such as one of the following:
36423
              (umask 002)
36424
              nohup umask ...
36425
36426
              find . -exec umask ... \;
              it shall not affect the file mode creation mask of the caller's environment.
36427
              If the mask operand is not specified, the umask utility shall write to standard output the value of
36428
              the invoking process' file mode creation mask.
36429
36430 OPTIONS
              The umask utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
36431
              12.2, Utility Syntax Guidelines.
36432
              The following option shall be supported:
36433
              -S
                            Produce symbolic output.
36434
              The default output style is unspecified, but shall be recognized on a subsequent invocation of
36435
              umask on the same system as a mask operand to restore the previous file mode creation mask.
36436
36437 OPERANDS
              The following operand shall be supported:
36438
                            A string specifying the new file mode creation mask. The string is treated in the
36439
              mask
                            same way as the mode operand described in the EXTENDED DESCRIPTION
36440
                            section for chmod.
36441
                            For a symbolic_mode value, the new value of the file mode creation mask shall be
36442
36443
                            the logical complement of the file permission bits portion of the file mode specified
36444
                            by the symbolic_mode string.
                            In a symbolic_mode value, the permissions op characters '+' and '-' shall be
36445
36446
                            interpreted relative to the current file mode creation mask; '+' shall cause the bits
                            for the indicated permissions to be cleared in the mask; '-' shall cause the bits for
36447
                            the indicated permissions to be set in the mask.
36448
                            The interpretation of mode values that specify file mode bits other than the file
36449
                            permission bits is unspecified.
36450
                            In the octal integer form of mode, the specified bits are set in the file mode creation
36451
36452
                            The file mode creation mask shall be set to the resulting numeric value.
36453
```

operand also shall be recognized as a *mask* operand.

36454

36455

The default output of a prior invocation of *umask* on the same system with no

Utilities umask

36456 STDIN				
36457 36458 INPUT	Not used.			
36459	None.			
36460 ENVIR 36461	36460 ENVIRONMENT VARIABLES			
36462 36463 36464 36465	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)		
36466 36467	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.		
36468 36469 36470	LC_CTYPE	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).		
36471 36472 36473	LC_MESSA	Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.		
36474 XSI	NLSPATH	Determine the location of message catalogs for the processing of $LC_MESSAGES$.		
36475 ASYNCHRONOUS EVENTS 36476 Default.				
36477 STDOU				
36478 36479	When the <i>mask</i> operand is not specified, the <i>umask</i> utility shall write a message to standard output that can later be used as a <i>umask mask</i> operand.			
36480	If $-S$ is specified, the message shall be in the following format:			
36481 36482	"u=%s,g=%s,o=%s\n", <owner permissions="">, <group permissions="">, <other permissions=""></other></group></owner>			
36483 36484	where the three values shall be combinations of letters from the set $\{r, w, x\}$; the presence of a letter shall indicate that the corresponding bit is clear in the file mode creation mask.			
36485	If a <i>mask</i> operand is specified, there shall be no output written to standard output.			
36486 STDERR 36487 The standard error shall be used only for diagnostic messages.				
36488 OUTPUT FILES 36489 None.				
36490 EXTENDED DESCRIPTION 36491 None.				
36492 EXIT S				
36493		ng exit values shall be returned:		
36494	0 The file	mode creation mask was successfully changed, or no <i>mask</i> operand was supplied.		

>0 An error occurred.

umask Utilities

36496 CONSEQUENCES OF ERRORS

36497 Default.

36498 APPLICATION USAGE

Since umask affects the current shell execution environment, it is generally provided as a shell 36499

36500 regular built-in.

In contrast to the negative permission logic provided by the file mode creation mask and the 36501 octal number form of the *mask* argument, the symbolic form of the *mask* argument specifies those 36502

permissions that are left alone.

36504 EXAMPLES

36503

36510

Either of the commands: 36505

36506 umask a=rx,ug+w

36507 umask 002

36508 sets the mode mask so that subsequently created files have their S_IWOTH bit cleared.

After setting the mode mask with either of the above commands, the umask command can be 36509

used to write out the current value of the mode mask:

\$ umask 36511

0002 36512

(The output format is unspecified, but historical implementations use the octal integer mode 36513

36514 format.)

36515 \$ umask -S

36516 u=rwx,g=rwx,o=rx

36517 Either of these outputs can be used as the mask operand to a subsequent invocation of the *umask*

36518 utility.

Assuming the mode mask is set as above, the command: 36519

36520 umask q-w

sets the mode mask so that subsequently created files have their S_IWGRP and S_IWOTH bits 36521

cleared. 36522

The command: 36523

umask -- -w 36524

36525 sets the mode mask so that subsequently created files have all their write bits cleared. Note that 36526

mask operands $-\mathbf{r}$, $-\mathbf{w}$, $-\mathbf{x}$ or anything beginning with a hyphen, must be preceded by "--" to

keep it from being interpreted as an option.

36528 RATIONALE

36527

36531

Since *umask* affects the current shell execution environment, it is generally provided as a shell 36529 36530 regular built-in. If it is called in a subshell or separate utility execution environment, such as one

of the following:

(umask 002) 36532 36533 nohup umask ... find . -exec umask ... \; 36534

it does not affect the file mode creation mask of the environment of the caller. 36535

The description of the historical utility was modified to allow it to use the symbolic modes of 36536 36537 *chmod.* The -s option used in early proposals was changed to -S because -s could be confused

Utilities umask

36538	with a <i>symbolic_mode</i> form of mask referring to the S_ISUID and S_ISGID bits.
36539 36540 36541 36542 36543	The default output style is implementation-defined to permit implementors to provide migration to the new symbolic style at the time most appropriate to their users. A $-\mathbf{o}$ flag to force octal mode output was omitted because the octal mode may not be sufficient to specify all of the information that may be present in the file mode creation mask when more secure file access permission checks are implemented.
36544 36545 36546 36547 36548 36549	It has been suggested that trusted systems developers might appreciate ameliorating the requirement that the mode mask "affects" the file access permissions, since it seems access control lists might replace the mode mask to some degree. The wording has been changed to say that it affects the file permission bits, and it leaves the details of the behavior of how they affect the file access permissions to the description in the System Interfaces volume of IEEE Std 1003.1-2001.
36550 FUTUI	REDIRECTIONS
36551	None.
36552 SEE Al 36553	LSO Chapter 2 (on page 29), <i>chmod</i> , the System Interfaces volume of IEEE Std 1003.1-2001, <i>umask</i> ()
36554 CHAN 36555	GE HISTORY First released in Issue 2.
36556 Issue 6 36557 36558	The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:
36559	The octal mode is supported.

unalias Utilities

36560 **NAME** 36561 unalias — remove alias definitions 36562 SYNOPSIS unalias alias-name.. 36563 UP 36564 unalias -a 36565 36566 **DESCRIPTION** The unalias utility shall remove the definition for each alias name specified. See Section 2.3.1 (on 36567 page 32). The aliases shall be removed from the current shell execution environment; see Section 36568 2.12 (on page 61). 36569 36570 OPTIONS The unalias utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 36571 12.2, Utility Syntax Guidelines. 36572 The following option shall be supported: 36573 Remove all alias definitions from the current shell execution environment. 36574 36575 OPERANDS The following operand shall be supported: 36576 alias-name The name of an alias to be removed. 36577 36578 STDIN 36579 Not used. 36580 INPUT FILES None. 36581 36582 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *unalias*: 36583 LANG 36584 Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 36585 Internationalization Variables for the precedence of internationalization variables 36586 used to determine the values of locale categories.) 36587 LC_ALL If set to a non-empty string value, override the values of all the other 36588 internationalization variables. 36589 Determine the locale for the interpretation of sequences of bytes of text data as 36590 LC_CTYPE characters (for example, single-byte as opposed to multi-byte characters in 36591 arguments). 36592 LC_MESSAGES 36593 Determine the locale that should be used to affect the format and contents of 36594 36595 diagnostic messages written to standard error. NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 36596 XSI 36597 ASYNCHRONOUS EVENTS Default. 36598

36599 STDOUT

36600

Not used.

Utilities unalias

36601 STDERR

36602 The standard error shall be used only for diagnostic messages.

36603 OUTPUT FILES

36604 None.

36605 EXTENDED DESCRIPTION

36606 None.

36607 EXIT STATUS

36608 The following exit values shall be returned:

36609 0 Successful completion.

>0 One of the *alias-name* operands specified did not represent a valid alias definition, or an

36611 error occurred.

36612 CONSEQUENCES OF ERRORS

36613 Default.

36614 APPLICATION USAGE

Since *unalias* affects the current shell execution environment, it is generally provided as a shell

36616 regular built-in.

36617 EXAMPLES

36618 None.

36619 RATIONALE

The *unalias* description is based on that from historical KornShell implementations. Known differences exist between that and the C shell. The KornShell version was adopted to be consistent with all the other KornShell features in this volume of IEEE Std 1003.1-2001, such as

36623 command line editing.

The -a option is the equivalent of the *unalias* * form of the C shell and is provided to address security concerns about unknown aliases entering the environment of a user (or application) through the allowable implementation-defined predefined alias route or as a result of an *ENV* file. (Although *unalias* could be used to simplify the "secure" shell script shown in the *command* rationale, it does not obviate the need to quote all command names. An initial call to *unalias* -a would have to be quoted in case there was an alias for *unalias*.)

36630 FUTURE DIRECTIONS

36631 None.

36632 **SEE ALSO**

36633 Chapter 2 (on page 29), alias

36634 CHANGE HISTORY

36635 First released in Issue 4.

36636 Issue 6

36637 This utility is marked as part of the User Portability Utilities option.

Utilities uname

36638 NAME

36639 uname — return system name

36640 SYNOPSIS

36641 uname [-snrvma]

36642 DESCRIPTION

By default, the *uname* utility shall write the operating system name to standard output. When 36643 options are specified, symbols representing one or more system characteristics shall be written 36644 to the standard output. The format and contents of the symbols are implementation-defined. On 36645 systems conforming to the System Interfaces volume of IEEE Std 1003.1-2001, the symbols 36646 36647 written shall be those supported by the *uname()* function as defined in the System Interfaces volume of IEEE Std 1003.1-2001. 36648

36649 **OPTIONS**

The *uname* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 36650 12.2, Utility Syntax Guidelines. 36651

The following options shall be supported: 36652

36653 −a Behave as though all of the options –**mnrsv** were specified.

Write the name of the hardware type on which the system is running to standard 36654 -m output. 36655

Write the name of this node within an implementation-defined communications 36656 -n

36657 network.

Write the current release level of the operating system implementation. 36658 $-\mathbf{r}$

Write the name of the implementation of the operating system. 36659 -s

_v Write the current version level of this release of the operating system 36660 36661

implementation.

If no options are specified, the *uname* utility shall write the operating system name, as if the -s 36662 option had been specified. 36663

36664 OPERANDS

None. 36665

36666 STDIN

Not used.

36668 INPUT FILES

None 36669

36670 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *uname*: 36671

LANG Provide a default value for the internationalization variables that are unset or null. 36672 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 36673 Internationalization Variables for the precedence of internationalization variables 36674 36675 used to determine the values of locale categories.)

LC ALL If set to a non-empty string value, override the values of all the other 36676 36677

internationalization variables.

36678 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 36679

characters (for example, single-byte as opposed to multi-byte characters in

arguments).

Utilities uname

36681 LC_MESSAGES 36682 Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. 36683 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 36684 XSI 36685 ASYNCHRONOUS EVENTS Default. 36686 **36687 STDOUT** By default, the output shall be a single line of the following form: 36688 36689 "%s\n", <sysname> If the -a option is specified, the output shall be a single line of the following form: 36690 "%s %s %s %s %s\n", <sysname>, <nodename>, <release>, 36691 <version>, <machine> 36692 Additional implementation-defined symbols may be written; all such symbols shall be written at 36693 the end of the line of output before the <newline>. 36694 If options are specified to select different combinations of the symbols, only those symbols shall 36695 be written, in the order shown above for the -a option. If a symbol is not selected for writing, its 36696 corresponding trailing <blank>s also shall not be written. 36697 **36698 STDERR** 36699 The standard error shall be used only for diagnostic messages. 36700 OUTPUT FILES 36701 None. 36702 EXTENDED DESCRIPTION 36703 None. 36704 EXIT STATUS 36705 The following exit values shall be returned: 36706 The requested information was successfully written. An error occurred. 36707 36708 CONSEQUENCES OF ERRORS Default. 36710 APPLICATION USAGE 36711 Note that any of the symbols could include embedded <space>s, which may affect parsing 36712 algorithms if multiple options are selected for output. The node name is typically a name that the system uses to identify itself for inter-system 36713 communication addressing. 36714 36715 EXAMPLES The following command: 36716 36717 uname -sr 36718 writes the operating system name and release level, separated by one or more

blank>s.

uname Utilities

36719 RATIONALE

It was suggested that this utility cannot be used portably since the format of the symbols is implementation-defined. The POSIX.1 working group could not achieve consensus on defining these formats in the underlying *uname()* function, and there was no expectation that this volume of IEEE Std 1003.1-2001 would be any more successful. Some applications may still find this historical utility of value. For example, the symbols could be used for system log entries or for comparison with operator or user input.

36726 FUTURE DIRECTIONS

36727 None.

36728 SEE ALSO

The System Interfaces volume of IEEE Std 1003.1-2001, uname()

36730 CHANGE HISTORY

First released in Issue 2.

Utilities uncompress

36732 **NAME** 36733 uncompress — expand compressed data 36734 SYNOPSIS uncompress [-cfv][file...] 36735 XSI 36736 36737 **DESCRIPTION** The uncompress utility shall restore files to their original state after they have been compressed 36738 using the *compress* utility. If no files are specified, the standard input shall be uncompressed to 36739 the standard output. If the invoking process has appropriate privileges, the ownership, modes, 36740 access time, and modification time of the original file shall be preserved. 36741 This utility shall support the uncompressing of any files produced by the *compress* utility on the 36742 36743 same implementation. For files produced by compress on other systems, uncompress supports 9 to 14-bit compression (see *compress*, -**b**); it is implementation-defined whether values of -**b** greater 36744 than 14 are supported. 36745 **36746 OPTIONS** The uncompress utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, 36747 Section 12.2, Utility Syntax Guidelines. 36748 The following options shall be supported: 36749 36750 **-с** Write to standard output; no files are changed. $-\mathbf{f}$ 36751 Do not prompt for overwriting files. Except when run in the background, if -f is not given the user shall be prompted as to whether an existing file should be 36752 overwritten. If the standard input is not a terminal and -f is not given, uncompress 36753 shall write a diagnostic message to standard error and exit with a status greater 36754 than zero. 36755 Write messages to standard error concerning the expansion of each file. 36756 $-\mathbf{v}$ 36757 OPERANDS The following operand shall be supported: 36758 file A pathname of a file. If file already has the .Z suffix specified, it shall be used as the 36759 input file and the output file shall be named file with the .Z suffix removed. 36760 Otherwise, file shall be used as the name of the output file and file with the .Z 36761 36762 suffix appended shall be used as the input file. 36763 **STDIN** 36764 The standard input shall be used only if no file operands are specified, or if a file operand is '-'. 36765 INPUT FILES

36766

36768

Input files shall be in the format produced by the *compress* utility.

36767 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *uncompress*:

36769	LANG	Provide a default value for the internationalization variables that are unset or null.
36770		(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
36771		Internationalization Variables for the precedence of internationalization variables
36772		used to determine the values of locale categories.)
36773	LC_ALL	If set to a non-empty string value, override the values of all the other

internationalization variables. 36774

uncompressUtilities

36775 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 36776 characters (for example, single-byte as opposed to multi-byte characters in arguments). 36777 LC MESSAGES 36778 Determine the locale that should be used to affect the format and contents of 36779 diagnostic messages written to standard error. 36780 **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 36781 36782 ASYNCHRONOUS EVENTS 36783 Default. 36784 **STDOUT** 36785 When there are no file operands or the -c option is specified, the uncompressed output is written to standard output. 36786 36787 STDERR Prompts shall be written to the standard error output under the conditions specified in the 36788 DESCRIPTION and OPTIONS sections. The prompts shall contain the *file* pathname, but their 36789 format is otherwise unspecified. Otherwise, the standard error output shall be used only for 36790 36791 diagnostic messages. 36792 OUTPUT FILES Output files are the same as the respective input files to *compress*. 36793 36794 EXTENDED DESCRIPTION None. 36795 36796 EXIT STATUS The following exit values shall be returned: 36797 36798 Successful completion. >0 An error occurred. 36799 36800 CONSEQUENCES OF ERRORS 36801 The input file remains unmodified. **36802 APPLICATION USAGE** The limit of 14 on the compress -b bits argument is to achieve portability to all systems (within 36803 36804 the restrictions imposed by the lack of an explicit published file format). Some implementations based on 16-bit architectures cannot support 15 or 16-bit uncompression. 36805 36806 EXAMPLES None. 36807 36808 RATIONALE None. 36809 36810 FUTURE DIRECTIONS 36811 None. 36812 SEE ALSO 36813 compress, zcat

36815

36814 CHANGE HISTORY

First released in Issue 4.

Utilities uncompress

36816 **Issue 6**

36817

The normative text is reworded to avoid use of the term "must" for application requirements.

unexpand Utilities

NAME

36819 unexpand — convert spaces to tabs

36820 SYNOPSIS

36821 UP unexpand [-a | -t tablist][file...]
36822

DESCRIPTION

36830 OPTIONS

The *unexpand* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported:

In addition to translating
sequences of two or more
blank>s at the beginning of each line, translate all sequences of two or more
blank>s immediately preceding a tab stop to the maximum number of <tab>s followed by the minimum number of <space>s
needed to fill the same column positions originally filled by the translated
>blank>s.

Specify the tab stops. The application shall ensure that the *tablist* option-argument is a single argument consisting of a single positive decimal integer or multiple positive decimal integers, separated by

blank>s or commas, in ascending order. If a single number is given, tabs shall be set *tablist* column positions apart instead of

the default 8. If multiple numbers are given, the tabs shall be set at those specific column positions.

The application shall ensure that each tab-stop position N is an integer value greater than zero, and the list shall be in strictly ascending order. This is taken to mean that, from the start of a line of output, tabbing to position N shall cause the next character output to be in the (N+1)th column position on that line. When the $-\mathbf{t}$ option is not specified, the default shall be the equivalent of specifying $-\mathbf{t}$ 8 (except for the interaction with $-\mathbf{a}$, described below).

No <space>-to-<tab> conversions shall occur for characters at positions beyond the last of those specified in a multiple tab-stop list.

When -t is specified, the presence or absence of the -a option shall be ignored; conversion shall not be limited to the processing of leading <blank>s.

OPERANDS

36856 The following operand shall be supported:

file A pathname of a text file to be used as input.

STDIN

36859 See the INPUT FILES section.

Utilities unexpand

36860 INPUT FILES

36861 The input files shall be text files.

36862 ENVIRONMENT VARIABLES

The following environment variables shall affect the execution of *unexpand*: 36863

LANG 36864 Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 36865 Internationalization Variables for the precedence of internationalization variables 36866 36867

used to determine the values of locale categories.)

LC ALL If set to a non-empty string value, override the values of all the other 36868

internationalization variables. 36869

Determine the locale for the interpretation of sequences of bytes of text data as 36870 LC_CTYPE characters (for example, single-byte as opposed to multi-byte characters in 36871 arguments and input files), the processing of <tab>s and <space>s, and for the 36872 36873 determination of the width in column positions each character would occupy on

an output device.

LC MESSAGES 36875

Determine the locale that should be used to affect the format and contents of 36876 36877

diagnostic messages written to standard error.

NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 36878 XSI

36879 ASYNCHRONOUS EVENTS

Default. 36880

36881 STDOUT

36874

The standard output shall be equivalent to the input files with the specified <space>-to-<tab> 36882 conversions.

36883

36884 STDERR

36885

The standard error shall be used only for diagnostic messages.

36886 OUTPUT FILES

None. 36887

36888 EXTENDED DESCRIPTION

None. 36889

36890 EXIT STATUS

The following exit values shall be returned: 36891

Successful completion. 36892

36893 >0 An error occurred.

36894 CONSEQUENCES OF ERRORS

36895 Default. unexpand **Utilities**

36896 APPLICATION USAGE

36897 One non-intuitive aspect of *unexpand* is its restriction to leading spaces when neither -a nor -t is specified. Users who always want to convert all spaces in a file can easily alias unexpand to use 36898

the $-\mathbf{a}$ or $-\mathbf{t}$ 8 option. 36899

36900 EXAMPLES

None. 36901

36902 RATIONALE

On several occasions, consideration was given to adding a -t option to the unexpand utility to 36903 complement the -t in expand (see expand). The historical intent of unexpand was to translate 36904 36905 multiple <blank>s into tab stops, where tab stops were a multiple of eight column positions on most UNIX systems. An early proposal omitted -t because it seemed outside the scope of the 36906 User Portability Utilities option; it was not described in any of the base documents. However, 36907 hard-coding tab stops every eight columns was not suitable for the international community and 36908 broke historical precedents for some vendors in the FORTRAN community, so -t was restored 36909 36910 in conjunction with the list of valid extension categories considered by the standard developers. 36911 Thus, *unexpand* is now the logical converse of *expand*.

36912 FUTURE DIRECTIONS

None. 36913

36914 SEE ALSO

36915 expand, tabs

36916 CHANGE HISTORY

First released in Issue 4. 36917

36918 Issue 6

This utility is marked as part of the User Portability Utilities option. 36919

36920 The definition of the LC_CTYPE environment variable is changed to align with the 36921

IEEE P1003.2b draft standard.

36922 The normative text is reworded to avoid use of the term "must" for application requirements. **Utilities** unget

36923 **NAME** 36924 unget — undo a previous get of an SCCS file (**DEVELOPMENT**) 36925 SYNOPSIS unget [-ns] [-r SID] file... 36926 XSI 36927 36928 **DESCRIPTION** The *unget* utility shall reverse the effect of a *get* –e done prior to creating the intended new delta. 36929 36930 OPTIONS 36931 The *unget* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines. 36932 36933 The following options shall be supported: -r SID Uniquely identify which delta is no longer intended. (This would have been 36934 specified by get as the new delta.) The use of this option is necessary only if two or 36935 more outstanding get commands for editing on the same SCCS file were done by 36936 the same person (login name). 36937 Suppress the writing to standard output of the intended delta's SID. 36938 -sRetain the file that was obtained by get, which would normally be removed from 36939 -n the current directory. 36940 36941 **OPERANDS** 36942 The following operands shall be supported: file A pathname of an existing SCCS file or a directory. If file is a directory, the unget 36943 utility shall behave as though each file in the directory were specified as a named 36944 file, except that non-SCCS files (last component of the pathname does not begin 36945 36946 with **s.**) and unreadable files shall be silently ignored. If exactly one *file* operand appears, and it is '-', the standard input shall be read; 36947 36948 each line of the standard input shall be taken to be the name of an SCCS file to be processed. Non-SCCS files and unreadable files shall be silently ignored. 36949 36950 **STDIN** The standard input shall be a text file used only when the file operand is specified as '-'. Each 36951 line of the text file shall be interpreted as an SCCS pathname. 36952 36953 INPUT FILES Any SCCS files processed shall be files of an unspecified format. 36954

36955 ENVIRONMENT VARIABLES

36956 The following environment variables shall affect the execution of *unget*:

36957 36958 36959 36960	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
36961 36962	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
36963 36964 36965	LC_CTYPE	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).

unget Utilities

36966 LC_MESSAGES 36967 Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error. 36968 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 36969 36970 ASYNCHRONOUS EVENTS 36971 Default. 36972 **STDOUT** The standard output shall consist of a line for each file, in the following format: 36973 36974 "%s\n", <SID removed from file> 36975 If there is more than one named file or if a directory or standard input is named, each pathname shall be written before each of the preceding lines: 36976 "\n%s:\n", <pathname> 36977 **36978 STDERR** The standard error shall be used only for diagnostic messages. 36979 36980 OUTPUT FILES Any SCCS files updated shall be files of an unspecified format. During processing of a file, a 36981 locking z-file, as described in get, and a q-file (a working copy of the p-file), may be created and 36982 deleted. The *p-file* and *g-file*, as described in *get*, shall be deleted. 36983 36984 EXTENDED DESCRIPTION None. 36985 36986 EXIT STATUS The following exit values shall be returned: 36987 36988 Successful completion. >0 An error occurred. 36989 36990 CONSEQUENCES OF ERRORS Default. 36991 36992 APPLICATION USAGE 36993 None. 36994 EXAMPLES 36995 None. 36996 RATIONALE None. 36998 FUTURE DIRECTIONS 36999 None. 37000 SEE ALSO 37001 delta, get, sact 37002 CHANGE HISTORY 37003 First released in Issue 2. 37004 **Issue 6**

37005

The normative text is reworded to avoid use of the term "must" for application requirements.

Utilities uniq

37006 NAME 37007 uniq — report or filter out repeated lines in a file 37008 SYNOPSIS uniq [-c|-d|-u] [-f fields] [-s char] [input file [output file]] 37009 37010 **DESCRIPTION** The *uniq* utility shall read an input file comparing adjacent lines, and write one copy of each 37011 input line on the output. The second and succeeding copies of repeated adjacent input lines shall 37012 37013 37014 Repeated lines in the input shall not be detected if they are not adjacent. 37015 OPTIONS The uniq utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 37016 12.2, Utility Syntax Guidelines. 37017 The following options shall be supported: 37018 37019 Precede each output line with a count of the number of times the line occurred in -c the input. 37020 $-\mathbf{d}$ Suppress the writing of lines that are not repeated in the input. 37021 -f fields Ignore the first *fields* fields on each input line when doing comparisons, where 37022 fields is a positive decimal integer. A field is the maximal string matched by the 37023 basic regular expression: 37024 [[:blank:]] * [^[:blank:]] * 37025 37026 If the *fields* option-argument specifies more fields than appear on an input line, a null string shall be used for comparison. 37027 Ignore the first chars characters when doing comparisons, where chars shall be a 37028 -s chars 37029 positive decimal integer. If specified in conjunction with the -f option, the first chars characters after the first fields fields shall be ignored. If the chars option-37030 argument specifies more characters than remain on an input line, a null string shall 37031 be used for comparison. 37032 Suppress the writing of lines that are repeated in the input. 37033 -u 37034 OPERANDS The following operands shall be supported: 37035 A pathname of the input file. If the *input file* operand is not specified, or if the 37036 input file $input_file$ is '-', the standard input shall be used. 37037 A pathname of the output file. If the output_file operand is not specified, the 37038 output_file standard output shall be used. The results are unspecified if the file named by 37039 37040 output file is the file named by input file. 37041 **STDIN** 37042 The standard input shall be used only if no *input_file* operand is specified or if *input_file* is '-'. See the INPUT FILES section. 37043 37044 INPUT FILES 37045 The input file shall be a text file.

uniq **Utilities**

37046 ENVIRONMENT VARIABLES				
37047 37048 37049 37050 37051	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)		
37052 37053	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.		
37054 37055	LC_COLLATE Determine the locale for ordering rules.			
37056 37057 37058 37059	LC_CTYPE	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) and which characters constitute a blank> in the current locale.		
37060 37061 37062	LC_MESSA	GES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.		
37063 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .		
37064 ASYNC 37065	CHRONOUS Default.	EVENTS		
37066 STDOUT 37067 The standard output shall be used only if no <i>output_file</i> operand is specified. See the OUTPUT 37068 FILES section.				
37069 STDERR 37070 The standard error shall be used only for diagnostic messages.				
37071 OUTPU				
37072	If the $-c$ option is specified, the output file shall be empty or each line shall be of the form:			
37073	"%d %s", <number duplicates="" of="">, <line></line></number>			
37074	otherwise, the output file shall be empty or each line shall be of the form:			
37075	"%s", <11			
37076 EXTENDED DESCRIPTION 37077 None.				
37078 EXIT S 7 37079		ng exit values shall be returned:		
37080	0 The uti	lity executed successfully.		
37081	>0 An erro	or occurred.		
37082 CONSE 37083	E QUENCES (Default.	OF ERRORS		

Utilities uniq

37084 APPLICATION USAGE

37085 The *sort* utility can be used to cause repeated lines to be adjacent in the input file.

37086 EXAMPLES

The following input file data (but flushed left) was used for a test series on *uniq*:

```
37088#01 foo0 bar0 foo1 bar137089#02 bar0 foo1 bar1 foo137090#03 foo0 bar0 foo1 bar137091#0437092#05 foo0 bar0 foo1 bar137093#06 foo0 bar0 foo1 bar137094#07 bar0 foo1 bar1 foo0
```

What follows is a series of test invocations of the *uniq* utility that use a mixture of *uniq* options against the input file data. These tests verify the meaning of *adjacent*. The *uniq* utility views the input data as a sequence of strings delimited by ' \n' . Accordingly, for the *fields*th member of the sequence, *uniq* interprets unique or repeated adjacent lines strictly relative to the *fields*+1th member.

1. This first example tests the line counting option, comparing each line of the input file data starting from the second field:

```
uniq -c -f 1 uniq_0I.t
    1 #01 foo0 bar0 foo1 bar1
1 #02 bar0 foo1 bar1 foo0
1 #03 foo0 bar0 foo1 bar1
1 #04
2 #05 foo0 bar0 foo1 bar1
1 #07 bar0 foo1 bar1 foo0
```

The number '2', prefixing the fifth line of output, signifies that the *uniq* utility detected a pair of repeated lines. Given the input data, this can only be true when *uniq* is run using the –**f 1** option (which shall cause *uniq* to ignore the first field on each input line).

2. The second example tests the option to suppress unique lines, comparing each line of the input file data starting from the second field:

```
uniq -d -f 1 uniq_0I.t
#05 foo0 bar0 foo1 bar1
```

3. This test suppresses repeated lines, comparing each line of the input file data starting from the second field:

```
uniq -u -f 1 uniq_0I.t
#01 foo0 bar0 foo1 bar1
#02 bar0 foo1 bar1 foo1
#03 foo0 bar0 foo1 bar1
#04
#07 bar0 foo1 bar1 foo0
```

4. This suppresses unique lines, comparing each line of the input file data starting from the third character:

```
uniq -d -s 2 uniq 0I.t
```

In the last example, the *uniq* utility found no input matching the above criteria.

uniq Utilities

37128 RATIONALE 37129 Some historical implementations have limited lines to be 1080 bytes in length, which does not $meet\ the\ implied\ \{LINE_MAX\}\ limit.$ 37130 37131 FUTURE DIRECTIONS 37132 None. 37133 **SEE ALSO** 37134 comm, sort 37135 CHANGE HISTORY First released in Issue 2. 37136 37137 **Issue 6** The obsolescent SYNOPSIS and associated text are removed. 37138 37139 The normative text is reworded to avoid use of the term "must" for application requirements. IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/40 is applied, adding LC_COLLATE to the 37140 37141 ENVIRONMENT VARIABLES section, and changing "the application shall ensure that" in the **OUTPUT FILES section.** 37142

Utilities unlink

```
37143 NAME
37144
             unlink — call the unlink() function
37145 SYNOPSIS
37146 XSI
             unlink file
37147
37148 DESCRIPTION
             The unlink utility shall perform the function call:
37149
37150
             unlink(file);
37151
             A user may need appropriate privilege to invoke the unlink utility.
37152 OPTIONS
             None.
37153
37154 OPERANDS
37155
             The following operands shall be supported:
                           The pathname of an existing file.
37156
37157 STDIN
             Not used.
37158
37159 INPUT FILES
37160
             Not used.
37161 ENVIRONMENT VARIABLES
             The following environment variables shall affect the execution of unlink:
37162
             LANG
                           Provide a default value for the internationalization variables that are unset or null.
37163
                           (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
37164
                           Internationalization Variables for the precedence of internationalization variables
37165
37166
                           used to determine the values of locale categories.)
37167
             LC ALL
                           If set to a non-empty string value, override the values of all the other
                           internationalization variables.
37168
37169
             LC_CTYPE
                           Determine the locale for the interpretation of sequences of bytes of text data as
37170
                           characters (for example, single-byte as opposed to multi-byte characters in
                           arguments).
37171
             LC MESSAGES
37172
                           Determine the locale that should be used to affect the format and contents of
37173
37174
                           diagnostic messages written to standard error.
             NLSPATH
37175
                           Determine the location of message catalogs for the processing of LC_MESSAGES.
37176 ASYNCHRONOUS EVENTS
             Default.
37177
37178 STDOUT
37179
             None.
37180 STDERR
```

The standard error shall be used only for diagnostic messages.

unlink Utilities

37182 OUTPUT FILES 37183 None. 37184 EXTENDED DESCRIPTION 37185 None. 37186 EXIT STATUS 37187 The following exit values shall be returned: 0 Successful completion. 37188 >0 An error occurred. 37189 37190 CONSEQUENCES OF ERRORS Default. 37192 APPLICATION USAGE 37193 None. 37194 EXAMPLES None. 37195 37196 RATIONALE 37197 None.

37198 FUTURE DIRECTIONS

37199 None.37200 SEE ALSO

link, rm, the System Interfaces volume of IEEE Std 1003.1-2001, unlink()

37202 **CHANGE HISTORY** 37203 First released in Issue 5. *Utilities* **uucp**

37204 NAME	ı				
37205	uucp — system-to-system copy				
37206 SYNOI 37207 XSI 37208	1 - 3				
37209 DESCF					
37210 37211	The <i>uucp</i> utility shall copy files named by the <i>source-file</i> argument to the <i>destination-file</i> argument. The files named can be on local or remote systems.				
37212 37213 37214 37215 37216 37217 37218	example, tr filenames in circumstand standard In and that on	The <i>uucp</i> utility cannot guarantee support for all character encodings in all circumstances. For example, transmission data may be restricted to 7 bits by the underlying network, 8-bit data and filenames need not be portable to non-internationalized systems, and so on. Under these circumstances, it is recommended that only characters defined in the ISO/IEC 646:1991 standard International Reference Version (equivalent to ASCII) 7-bit range of characters be used, and that only characters defined in the portable filename character set be used for naming files. The protocol for transfer of files is unspecified by IEEE Std 1003.1-2001.			
37219 37220 37221 37222 37223	Typical implementations of this utility require a communications line configured to use the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface, but other communications means may be used. On systems where there are no available communications means (either temporarily or permanently), this utility shall write an error message describing the problem and exit with a non-zero exit status.				
37224 OPTIO					
37225 37226	The <i>uucp</i> utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.				
37227	The followi	The following options shall be supported:			
37228 37229	-с	Do not copy local file to the spool directory for transfer to the remote machine (default).			
37230	-C	Force the copy of local files to the spool directory for transfer.			
37231	$-\mathbf{d}$	Make all necessary directories for the file copy (default).			
37232	-f	Do not make intermediate directories for the file copy.			
37233 37234	- j	Write the job identification string to standard output. This job identification can be used by <i>uustat</i> to obtain the status or terminate a job.			
37235	-m	Send mail to the requester when the copy is completed.			
37236	-n user	Notify user on the remote system that a file was sent.			
37237	-r	Do not start the file transfer; just queue the job.			
37238 OPERANDS 37239 The following operands shall be supported:					
37240 37241 37242		file, source-file A pathname of a file to be copied to, or from, respectively. Either name can be a pathname on the local machine, or can have the form:			
37243		system-name!pathname			
37244 37245		where <i>system-name</i> is taken from a list of system names that <i>uucp</i> knows about. The destination <i>system-name</i> can also be a list of names such as:			

uucp Utilities

37246		sys	tem-name	e!system-name!!system-name!pathname
37247 37248 37249		dest	ination. C	, an attempt is made to send the file via the specified route to the are should be taken to ensure that intermediate nodes in the route forward information.
37250 37251				ern matching notation characters $'?'$, $'*'$, and "[]" appearing hall be expanded on the appropriate system.
37252		Path	names cai	n be one of:
37253		1.	An abso	lute pathname.
37254 37255 37256 37257		2.	system a login is	ame preceded by "user where user is a login name on the specified and is replaced by that user's login directory. Note that if an invalid specified, the default is to the public directory (called <i>PUBDIR</i> ; the exation of <i>PUBDIR</i> is implementation-defined).
37258 37259		3.	A pathr PUBDIR	name preceded by ~/destination where destination is appended to ?.
37260 37261 37262 37263 37264			Note:	This destination is treated as a filename unless more than one file is being transferred by this request or the destination is already a directory. To ensure that it is a directory, follow the destination with a '/'. For example, ~/dan/ as the destination makes the directory PUBDIR/dan if it does not exist and puts the requested files in that directory.
37265		4.	Anythin	g else shall be prefixed by the current directory.
37266 37267				an erroneous pathname for the remote system, the copy shall fail. If <i>-file</i> is a directory, the last part of the <i>source-file</i> name shall be used.
37268 37269		The defin		ite, and execute permissions given by uucp are implementation-
37270 STDIN 37271	Not used.			
37272 INPUT FILES 37273 The files to be copied are regular files.				
	ONMENT VA			vanishles shall affect the avecuation of your
37275				variables shall affect the execution of <i>uucp</i> :
37276 37277 37278 37279	LANG	(See Inter	the Ba	nult value for the internationalization variables that are unset or null. ase Definitions volume of IEEE Std 1003.1-2001, Section 8.2, zation Variables for the precedence of internationalization variables nine the values of locale categories.)
37280 37281	LC_ALL			non-empty string value, override the values of all the other zation variables.
37282	LC_COLLAT	Έ		
37283 37284				e locale for the behavior of ranges, equivalence classes, and multi- nting elements within bracketed filename patterns.
37285 37286 37287 37288	LC_CTYPE	char argu	acters (fo ments an	e locale for the interpretation of sequences of bytes of text data as or example, single-byte as opposed to multi-byte characters in d input files) and the behavior of character classes within bracketed erns (for example, "'[[:lower:]]*'").

Utilities uucp

37289 LC_MESSAGES 37290 Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error, and informative messages written 37291 to standard output. 37292 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 37293 37294 ASYNCHRONOUS EVENTS Default. 37295 37296 **STDOUT** Not used. 37297 **37298 STDERR** 37299 The standard error shall be used only for diagnostic messages. 37300 OUTPUT FILES The output files (which may be on other systems) are copies of the input files. 37301 37302 If **-m** is used, mail files are modified. 37303 EXTENDED DESCRIPTION None. 37304 37305 EXIT STATUS The following exit values shall be returned: 37306 37307 Successful completion. 37308 >0 An error occurred. 37309 CONSEQUENCES OF ERRORS Default. 37310 37311 APPLICATION USAGE 37312 The domain of remotely accessible files can (and for obvious security reasons usually should) be 37313 severely restricted. 37314 Note that the '!' character in addresses has to be escaped when using csh as a command 37315 interpreter because of its history substitution syntax. For ksh and sh the escape is not necessary, 37316 but may be used. 37317 As noted above, shell metacharacters appearing in pathnames are expanded on the appropriate system. On an internationalized system, this is done under the control of local settings of 37318 LC COLLATE and LC CTYPE. Thus, care should be taken when using bracketed filename 37319 37320 patterns, as collation and typing rules may vary from one system to another. Also be aware that 37321 certain types of expression (that is, equivalence classes, character classes, and collating symbols) 37322 need not be supported on non-internationalized systems. 37323 EXAMPLES 37324 None. 37325 RATIONALE 37326 None. 37327 FUTURE DIRECTIONS None. 37328

uucp Utilities

37329 SEE AL 37330	SO mailx, uuencode, uustat, uux
37331 CHAN 37332	GE HISTORY First released in Issue 2.
37333 Issue 6 37334	The LC_TIME and TZ entries are removed from the ENVIRONMENT VARIABLES section.
37335 37336	The UN margin codes and associated shading are removed from the $-\mathbf{C}$, $-\mathbf{j}$, $-\mathbf{n}$, and $-\mathbf{r}$ options in response to The Open Group Base Resolution bwg2001-003.

Utilities uudecode

37337 **NAME**

37338 uudecode — decode a binary file

37339 SYNOPSIS

37340 UP uudecode [-o outfile] [file]

37341

37352 37353

37354

37342 **DESCRIPTION**

The uudecode utility shall read a file, or standard input if no file is specified, that includes data 37343 created by the *uuencode* utility. The *uudecode* utility shall scan the input file, searching for data 37344 compatible with one of the formats specified in *uuencode*, and attempt to create or overwrite the 37345 37346 file described by the data (or overridden by the $-\mathbf{o}$ option). The pathname shall be contained in the data or specified by the $-\mathbf{o}$ option. The file access permission bits and contents for the file to 37347 be produced shall be contained in that data. The mode bits of the created file (other than 37348 37349 standard output) shall be set from the file access permission bits contained in the data; that is, other attributes of the mode, including the file mode creation mask (see umask), shall not affect 37350 37351 the file being produced.

If the pathname of the file to be produced exists, and the user does not have write permission on that file, *uudecode* shall terminate with an error. If the pathname of the file to be produced exists, and the user has write permission on that file, the existing file shall be overwritten.

If the input data was produced by *uuencode* on a system with a different number of bits per byte than on the target system, the results of *uudecode* are unspecified.

37357 OPTIONS

The *uudecode* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

37360 The following option shall be supported by the implementation:

37361 — o outfile A pathname of a file that shall be used instead of any pathname contained in the input data. Specifying an outfile option-argument of /dev/stdout shall indicate standard output.

37364 OPERANDS

37365 The following operand shall be supported:

37366 *file* The pathname of a file containing the output of *uuencode*.

37367 **STDIN**

37381

37368 See the INPUT FILES section.

37369 INPUT FILES

37370 The input files shall be files containing the output of *uuencode*.

37371 ENVIRONMENT VARIABLES

37372 The following environment variables shall affect the execution of *uudecode*:

37373	LANG	Provide a default value for the internationalization variables that are unset or null.
37374		(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
37375		Internationalization Variables for the precedence of internationalization variables
37376		used to determine the values of locale categories.)
37377	LC_ALL	If set to a non-empty string value, override the values of all the other
37378		internationalization variables.

37379 *LC_CTYPE* 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).

uudecodeUtilities

37382 LC_MESSAGES

Determine the locale that should be used to affect the format and contents of

37384 diagnostic messages written to standard error.

37385 XSI NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

37386 ASYNCHRONOUS EVENTS

37387 Default.

37388 STDOUT

If the file data header encoded by *uuencode* is – or /**dev/stdout**, or the **–o** /**dev/stdout** option overrides the file data, the standard output shall be in the same format as the file originally encoded by *uuencode*. Otherwise, the standard output shall not be used.

37392 STDERR

37393 The standard error shall be used only for diagnostic messages.

37394 OUTPUT FILES

The output file shall be in the same format as the file originally encoded by *uuencode*.

37396 EXTENDED DESCRIPTION

37397 None.

37398 EXIT STATUS

37399 The following exit values shall be returned:

37400 0 Successful completion.

37401 >0 An error occurred.

37402 CONSEQUENCES OF ERRORS

37403 Default.

37404 APPLICATION USAGE

The user who is invoking *uudecode* must have write permission on any file being created.

The output of *uuencode* is essentially an encoded bit stream that is not cognizant of byte boundaries. It is possible that a 9-bit byte target machine can process input from an 8-bit source, if it is aware of the requirement, but the reverse is unlikely to be satisfying. Of course, the only data that is meaningful for such a transfer between architectures is generally character data.

37410 EXAMPLES

37405

37411 None.

37412 RATIONALE

Input files are not necessarily text files, as stated by an early proposal. Although the *uuencode* output is a text file, that output could have been wrapped within another file or mail message that is not a text file.

37416 The **-o** option is not historical practice, but was added at the request of WG15 so that the user could override the target pathname without having to edit the input data itself.

In early drafts, the [**–o** outfile] option-argument allowed the use of – to mean standard output.

The symbol – has only been used previously in IEEE Std 1003.1-2001 as a standard input indicator. The developers of the standard did not wish to overload the meaning of – in this manner. The /**dev/stdout** concept exists on most modern systems. The /**dev/stdout** syntax does not refer to a new special file. It is just a magic cookie to specify standard output.

Utilities uudecode

37423 FUTURE DIRECTIONS

37424 None.

37425 **SEE ALSO**

37426 umask, uuencode

37427 CHANGE HISTORY

First released in Issue 4.

37429 **Issue 6**

37430 This utility is marked as part of the User Portability Utilities option.

37431 The **-o** *outfile* option is added, as specified in the IEEE P1003.2b draft standard.

The normative text is reworded to avoid use of the term "must" for application requirements.

uuencode Utilities

37433 **NAME** 37434 uuencode — encode a binary file 37435 SYNOPSIS 37436 UP uuencode [-m][file] decode pathname 37437 37438 **DESCRIPTION** The uuencode utility shall write an encoded version of the named input file, or standard input if 37439 no file is specified, to standard output. The output shall be encoded using one of the algorithms 37440 described in the STDOUT section and shall include the file access permission bits (in chmod octal 37441 37442 or symbolic notation) of the input file and the decode_pathname, for re-creation of the file on another system that conforms to this volume of IEEE Std 1003.1-2001. 37443 37444 OPTIONS The *uuencode* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, 37445 Section 12.2, Utility Syntax Guidelines. 37446 The following option shall be supported by the implementation: 37447 37448 -m Encode the output using the MIME Base64 algorithm described in STDOUT. If -m is not specified, the historical algorithm described in STDOUT shall be used. 37449 37450 OPERANDS The following operands shall be supported: 37451 37452 decode_pathname The pathname of the file into which the *uudecode* utility shall place the decoded 37453 file. Specifying a decode_pathname operand of /dev/stdout shall indicate that 37454 uudecode is to use standard output. If there are characters in decode_pathname that 37455 are not in the portable filename character set the results are unspecified. 37456 file A pathname of the file to be encoded. 37457 37458 **STDIN** See the INPUT FILES section. 37459 37460 INPUT FILES Input files can be files of any type. 37461 37462 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *uuencode*: 37463

37464 37465 37466 37467	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)
37468 37469	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.
37470 37471 37472	LC_CTYPE	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).

37473 *LC_MESSAGES*

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

37474

Utilities uuencode

37476 XSI NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

37477 ASYNCHRONOUS EVENTS

37478 Default.

37479 STDOUT

37480 uuencode Base64 Algorithm

The standard output shall be a text file (encoded in the character set of the current locale) that begins with the line:

"begin-base64 Δ %s Δ %sn", <mode>, <decode pathname>

37484 and ends with the line:

37485 "====\n"

In both cases, the lines shall have no preceding or trailing
blank>s.

The encoding process represents 24-bit groups of input bits as output strings of four encoded characters. Proceeding from left to right, a 24-bit input group shall be formed by concatenating three 8-bit input groups. Each 24-bit input group then shall be treated as four concatenated 6-bit groups, each of which shall be translated into a single digit in the Base64 alphabet. When encoding a bit stream via the Base64 encoding, the bit stream shall be presumed to be ordered with the most-significant bit first. That is, the first bit in the stream shall be the high-order bit in the first byte, and the eighth bit shall be the low-order bit in the first byte, and so on. Each 6-bit group is used as an index into an array of 64 printable characters, as shown in Table 4-21.

Table 4-21 uuencode Base64 Values

Value	Encoding	Value	Encoding	Value	Encoding	Value	Encoding
0	А	17	R	34	i	51	Z
1	В	18	S	35	j	52	0
2	C	19	Т	36	k	53	1
3	D	20	U	37	1	54	2
4	E	21	V	38	m	55	3
5	F	22	W	39	n	56	4
6	G	23	X	40	0	57	5
7	H	24	Y	41	р	58	6
8	I I	25	Z	42	q	59	7
9	J	26	a	43	r	60	8
10	K	27	b	44	s	61	9
11	L	28	С	45	t	62	+
12	M	29	d	46	u	63	/
13	N	30	е	47	v		
14	0	31	f	48	w	(pad)	=
15	P	32	g	49	х		
16	Q	33	h	50	У		

The character referenced by the index shall be placed in the output string.

The output stream (encoded bytes) shall be represented in lines of no more than 76 characters each. All line breaks or other characters not found in the table shall be ignored by decoding software (see *uudecode*).

Special processing shall be performed if fewer than 24 bits are available at the end of a message or encapsulated part of a message. A full encoding quantum shall always be completed at the

uuencode Utilities

end of a message. When fewer than 24 input bits are available in an input group, zero bits shall be added (on the right) to form an integral number of 6-bit groups. Output character positions that are not required to represent actual input data shall be set to the character '='. Since all Base64 input is an integral number of octets, only the following cases can arise:

- 1. The final quantum of encoding input is an integral multiple of 24 bits; here, the final unit of encoded output shall be an integral multiple of 4 characters with no '=' padding.
- 2. The final quantum of encoding input is exactly 16 bits; here, the final unit of encoded output shall be three characters followed by one '=' padding character.
- 3. The final quantum of encoding input is exactly 8 bits; here, the final unit of encoded output shall be two characters followed by two '=' padding characters.

A terminating "====" evaluates to nothing and denotes the end of the encoded data.

uuencode Historical Algorithm

The standard output shall be a text file (encoded in the character set of the current locale) that begins with the line:

```
"begin\Deltas\Deltas\n" < mode>, < decode_pathname>
```

and ends with the line:

```
37536 "end\n"
```

In both cases, the lines shall have no preceding or trailing
blank>s.

The algorithm that shall be used for lines in between **begin** and **end** takes three octets as input and writes four characters of output by splitting the input at six-bit intervals into four octets, containing data in the lower six bits only. These octets shall be converted to characters by adding a value of 0x20 to each octet, so that each octet is in the range [0x20,0x5f], and then it shall be assumed to represent a printable character in the ISO/IEC 646: 1991 standard encoded character set. It then shall be translated into the corresponding character codes for the codeset in use in the current locale. (For example, the octet 0x41, representing 'A', would be translated to 'A' in the current codeset, such as 0xc1 if it were EBCDIC.)

Where the bits of two octets are combined, the least significant bits of the first octet shall be shifted left and combined with the most significant bits of the second octet shifted right. Thus the three octets *A*, *B*, *C* shall be converted into the four octets:

These octets then shall be translated into the local character set.

Each encoded line contains a length character, equal to the number of characters to be decoded plus 0x20 translated to the local character set as described above, followed by the encoded characters. The maximum number of octets to be encoded on each line shall be 45.

37557 STDERR

37558 The standard error shall be used only for diagnostic messages.

37559 OUTPUT FILES

37560 None.

Utilities uuencode

37561 EXTENDED DESCRIPTION

37562 None.

37563 EXIT STATUS

37564 The following exit values shall be returned:

37565 0 Successful completion.

37566 >0 An error occurred.

37567 CONSEQUENCES OF ERRORS

37568 Default.

37569 APPLICATION USAGE

The file is expanded by 35 percent (each three octets become four, plus control information) causing it to take longer to transmit.

Since this utility is intended to create files to be used for data interchange between systems with possibly different codesets, and to represent binary data as a text file, the ISO/IEC 646:1991 standard was chosen for a midpoint in the algorithm as a known reference point. The output from *uuencode* is a text file on the local system. If the output were in the ISO/IEC 646:1991 standard codeset, it might not be a text file (at least because the <newline>s might not match), and the goal of creating a text file would be defeated. If this text file was then carried to another machine with the same codeset, it would be perfectly compatible with that system's *uudecode*. If it was transmitted over a mail system or sent to a machine with a different codeset, it is assumed that, as for every other text file, some translation mechanism would convert it (by the time it reached a user on the other system) into an appropriate codeset. This translation only makes sense from the local codeset, not if the file has been put into a ISO/IEC 646:1991 standard representation first. Similarly, files processed by *uuencode* can be placed in *pax* archives, intermixed with other text files in the same codeset.

37585 EXAMPLES

37586 None.

37587 RATIONALE

A new algorithm was added at the request of the international community to parallel work in RFC 2045 (MIME). As with the historical *uuencode* format, the Base64 Content-Transfer-Encoding is designed to represent arbitrary sequences of octets in a form that is not humanly readable. A 65-character subset of the ISO/IEC 646: 1991 standard is used, enabling 6 bits to be represented per printable character. (The extra 65th character, '=', is used to signify a special processing function.)

This subset has the important property that it is represented identically in all versions of the ISO/IEC 646:1991 standard, including US ASCII, and all characters in the subset are also represented identically in all versions of EBCDIC. The historical *uuencode* algorithm does not share this property, which is the reason that a second algorithm was added to the ISO POSIX-2 standard.

The string "====" was used for the termination instead of the end used in the original format because the latter is a string that could be valid encoded input.

In an early draft, the **-m** option was named **-b** (for Base64), but it was renamed to reflect its relationship to the RFC 2045. A **-u** was also present to invoke the default algorithm, but since this was not historical practice, it was omitted as being unnecessary.

37604 See the RATIONALE section in *uudecode* for the derivation of the /**dev/stdout** symbol.

uuencode Utilities

37605 **FUTURE DIRECTIONS**

37606 None.

37607 SEE ALSO

37608 chmod, mailx, uudecode

37609 CHANGE HISTORY

First released in Issue 4.

37611 **Issue 6**

37612 This utility is marked as part of the User Portability Utilities option.

37613 The Base64 algorithm and the ability to output to /dev/stdout are added as specified in the

37614 IEEE P1003.2b draft standard.

Utilities uustat

37615 **NAME** 37616 uustat — uucp status inquiry and job control 37617 SYNOPSIS 37618 XSI uustat [-q| -k jobid| -r jobid] 37619 uustat [-s system] [-u user] 37620 37621 **DESCRIPTION** 37622 The *uustat* utility shall display the status of, or cancel, previously specified *uucp* requests, or 37623 provide general status on *uucp* connections to other systems. When no options are given, *uustat* shall write to standard output the status of all *uucp* requests 37624 37625 issued by the current user. Typical implementations of this utility require a communications line configured to use the Base 37626 Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface, but other 37627 communications means may be used. On systems where there are no available communications 37628 means (either temporarily or permanently), this utility shall write an error message describing 37629 the problem and exit with a non-zero exit status. 37630 37631 OPTIONS 37632 The *uustat* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines. 37633 37634 The following options shall be supported: Write the jobs queued for each machine. 37635 -q -k jobid Kill the *uucp* request whose job identification is *jobid*. The application shall ensure 37636 that the killed *uucp* request belongs to the person invoking *uustat* unless that user 37637 37638 has appropriate privileges. -r jobid Rejuvenate jobid. The files associated with jobid are touched so that their 37639 37640 modification time is set to the current time. This prevents the cleanup program from deleting the job until the jobs modification time reaches the limit imposed by 37641 37642 the program. Write the status of all *uucp* requests for remote system *system*. 37643 -s system 37644 -u user Write the status of all *uucp* requests issued by *user*. 37645 OPERANDS 37646 None. 37647 **STDIN** Not used. 37648 37649 INPUT FILES 37650 None. 37651 ENVIRONMENT VARIABLES 37652 The following environment variables shall affect the execution of *uustat*: LANG Provide a default value for the internationalization variables that are unset or null. 37653 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 37654

used to determine the values of locale categories.)

37655

37656

Internationalization Variables for the precedence of internationalization variables

uustat Utilities

LC ALL 37657 If set to a non-empty string value, override the values of all the other 37658 internationalization variables. Determine the locale for the interpretation of sequences of bytes of text data as 37659 LC_CTYPE characters (for example, single-byte as opposed to multi-byte characters in 37660 arguments). 37661 37662 LC_MESSAGES Determine the locale that should be used to affect the format and contents of 37663 diagnostic messages written to standard error, and informative messages written 37664 to standard output. 37665 37666 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 37667 ASYNCHRONOUS EVENTS Default. 37668 37669 **STDOUT** The standard output shall consist of information about each job selected, in an unspecified 37670 format. The information shall include at least the job ID, the user ID or name, and the remote 37671 system name. 37672 37673 STDERR 37674 The standard error shall be used only for diagnostic messages. 37675 OUTPUT FILES 37676 None. 37677 EXTENDED DESCRIPTION 37678 None. 37679 EXIT STATUS The following exit values shall be returned: 37680 0 Successful completion. 37681 37682 >0 An error occurred. 37683 CONSEQUENCES OF ERRORS 37684 Default. 37685 APPLICATION USAGE 37686 None. 37687 EXAMPLES 37688 None. 37689 RATIONALE None. 37690 37691 FUTURE DIRECTIONS None. 37692 37693 SEE ALSO 37694 uucp 37695 CHANGE HISTORY First released in Issue 2. 37696

Utilities uustat

37697 Issue 6	
37698	The normative text is reworded to avoid use of the term "must" for application requirements.
37699	The LC_TIME and TZ entries are removed from the ENVIRONMENT VARIABLES section.
37700	The UN margin code and associated shading are removed from the $-\mathbf{q}$ option in response to The
37701	Open Group Base Resolution bwg2001-003.

uux Utilities

```
37702 NAME 37703
```

uux — remote command execution

37704 SYNOPSIS

```
37705 XSI uux [-np] command-string

37706 uux [-jnp] command-string
```

37708 DESCRIPTION

The *uux* utility shall gather zero or more files from various systems, execute a shell pipeline (see Section 2.9 (on page 47)) on a specified system, and then send the standard output of the command to a file on a specified system. Only the first command of a pipeline can have a *system-name*! prefix. All other commands in the pipeline shall be executed on the system of the first command.

The following restrictions are applicable to the shell pipeline processed by *uux*:

• In gathering files from different systems, pathname expansion shall not be performed by *uux*. Thus, a request such as:

```
uux "c99 remsys!~/*.c"
```

would attempt to copy the file named literally *.c to the local system.

- The redirection operators ">>", "<<", ">| ", and ">&" shall not be accepted. Any use of these redirection operators shall cause this utility to write an error message describing the problem and exit with a non-zero exit status.
- The reserved word! cannot be used at the head of the pipeline to modify the exit status. (See the *command-string* operand description below.)
- Alias substitution shall not be performed.

A filename can be specified as for *uucp*; it can be an absolute pathname, a pathname preceded by *name* (which is replaced by the corresponding login directory), a pathname specified as ~/ *dest* (*dest* is prefixed by the public directory called *PUBDIR*; the actual location of *PUBDIR* is implementation-defined), or a simple filename (which is prefixed by *uux* with the current directory). See *uucp* for the details.

The execution of commands on remote systems shall take place in an execution directory known to the *uucp* system. All files required for the execution shall be put into this directory unless they already reside on that machine. Therefore, the application shall ensure that non-local filenames (without path or machine reference) are unique within the *uux* request.

The *uux* utility shall attempt to get all files to the execution system. For files that are output files, the application shall ensure that the filename is escaped using parentheses.

The remote system shall notify the user by mail if the requested command on the remote system was disallowed or the files were not accessible. This notification can be turned off by the $-\mathbf{n}$ option.

Typical implementations of this utility require a communications line configured to use the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface, but other communications means may be used. On systems where there are no available communications means (either temporarily or permanently), this utility shall write an error message describing the problem and exit with a non-zero exit status.

The *uux* utility cannot guarantee support for all character encodings in all circumstances. For example, transmission data may be restricted to 7 bits by the underlying network, 8-bit data and

Utilities **uux**

37746 filenames need not be portable to non-internationalized systems, and so on. Under these 37747 circumstances, it is recommended that only characters defined in the ISO/IEC 646:1991 standard International Reference Version (equivalent to ASCII) 7-bit range of characters be used 37748 37749 and that only characters defined in the portable filename character set be used for naming files. 37750 OPTIONS The uux utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 37751 37752 12.2, Utility Syntax Guidelines. The following options shall be supported: 37753 37754 Make the standard input to *uux* the standard input to the *command-string*. -p -j Write the job identification string to standard output. This job identification can be 37755 37756 used by *uustat* to obtain the status or terminate a job. Do not notify the user if the command fails. 37757 37758 OPERANDS 37759 The following operand shall be supported: 37760 command-string A string made up of one or more arguments that are similar to normal command 37761 37762 arguments, except that the command and any filenames can be prefixed by 37763 system-name!. A null system-name shall be interpreted as the local system. 37764 **STDIN** The standard input shall not be used unless the '-' or $-\mathbf{p}$ option is specified; in those cases, the 37765 standard input shall be made the standard input of the *command-string*. 37766 37767 INPUT FILES Input files shall be selected according to the contents of *command-string*. 37768 37769 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *uux*: 37770 LANG 37771 Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 37772 37773 Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 37774 LC ALL 37775 If set to a non-empty string value, override the values of all the other internationalization variables. 37776 37777 LC_CTYPE 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 37778 37779 arguments). LC_MESSAGES 37780 Determine the locale that should be used to affect the format and contents of 37781 37782 diagnostic messages written to standard error. 37783 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 37784 ASYNCHRONOUS EVENTS

Default.

uux Utilities

37786 STDOUT

The standard output shall not be used unless the **-j** option is specified; in that case, the job identification string shall be written to standard output in the following format:

37789 "%s\n", <jobid>

37790 STDERR

37791 The standard error shall be used only for diagnostic messages.

37792 OUTPUT FILES

Output files shall be created or written, or both, according to the contents of *command-string*.

If **-n** is not used, mail files shall be modified following any command or file-access failures on the remote system.

37796 EXTENDED DESCRIPTION

37797 None.

37798 EXIT STATUS

37799 The following exit values shall be returned:

37800 0 Successful completion.

37801 >0 An error occurred.

37802 CONSEQUENCES OF ERRORS

37803 Default.

37804 APPLICATION USAGE

Note that, for security reasons, many installations limit the list of commands executable on behalf of an incoming request from *uux*. Many sites permit little more than the receipt of mail via *uux*.

Any characters special to the command interpreter should be quoted either by quoting the entire *command-string* or quoting the special characters as individual arguments.

As noted in *uucp*, shell pattern matching notation characters appearing in pathnames are expanded on the appropriate local system. This is done under the control of local settings of *LC_COLLATE* and *LC_CTYPE*. Thus, care should be taken when using bracketed filename patterns, as collation and typing rules may vary from one system to another. Also be aware that certain types of expression (that is, equivalence classes, character classes, and collating symbols) need not be supported on non-internationalized systems.

37816 EXAMPLES

37810

37811

37812

37813 37814

37815

37817

37818

37819

37820

37821

37822

37823

37824 37825

37826

37827

1. The following command gets **file1** from system **a** and **file2** from system **b**, executes *diff* on the local system, and puts the results in **file.diff** in the local *PUBDIR* directory. (*PUBDIR* is the *uucp* public directory on the local system.)

```
uux "!diff a!/usr/file1 b!/a4/file2 >!~/file.diff"
```

2. The following command fails because *uux* places all files copied to a system in the same working directory. Although the files **xyz** are from two different systems, their filenames are the same and conflict.

```
uux "!diff a!/usr1/xyz b!/usr2/xyz >!~/xyz.diff"
```

3. The following command succeeds (assuming *diff* is permitted on system **a**) because the file local to system **a** is not copied to the working directory, and hence does not conflict with the file from system **c**.

Utilities uux

37828 uux "a!diff a!/usr/xyz c!/usr/xyz >!~/xyz.diff" 37829 RATIONALE None. 37830 37831 FUTURE DIRECTIONS None. 37832 37833 **SEE ALSO** Chapter 2 (on page 29), uucp, uuencode, uustat 37834 37835 CHANGE HISTORY First released in Issue 2. 37836 37837 **Issue 6** The obsolescent SYNOPSIS is removed. 37838 The normative text is reworded to avoid use of the term "must" for application requirements. 37839 The UN margin code and associated shading are removed from the $-\mathbf{j}$ option in response to The 37840

Open Group Base Resolution bwg2001-003.

val Utilities

37842 NAME	:							
37843	val — valio	val — validate SCCS files (DEVELOPMENT)						
37844 SYNO	37844 SYNOPSIS							
37845 XSI	val -							
37846	val [-s]	[-m name][-r SID][-y type] file						
37847								
	37848 DESCRIPTION							
37849 37850		tility shall determine whether the specified <i>file</i> is an SCCS file meeting the tics specified by the options.						
37851 OPTIC		the specified by the options.						
37851 OPTIC 37852		ity shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2,						
37853		tax Guidelines, except that the usage of the '-' operand is not strictly as intended by						
37854	the guideli	nes (that is, reading options and operands from standard input).						
37855	The follow	ing options shall be supported:						
37856	-m name	Specify a <i>name</i> , which is compared with the SCCS %M% keyword in <i>file</i> ; see <i>get</i> .						
37857	−r SID	Specify a SID (SCCS Identification String), an SCCS delta number. A check shall be						
37858		made to determine whether the <i>SID</i> is ambiguous (for example, -r 1 is ambiguous						
37859		because it physically does not exist but implies 1.1, 1.2, and so on, which may						
37860 37861		exist) or invalid (for example, - r 1.0 or - r 1.1.0 are invalid because neither case can exist as a valid delta number). If the <i>SID</i> is valid and not ambiguous, a check shall						
37862		be made to determine whether it actually exists.						
37863 37864	-s	Silence the diagnostic message normally written to standard output for any error that is detected while processing each named file on a given command line.						
37865 37866	−y type	Specify a <i>type</i> , which shall be compared with the SCCS % Y % keyword in <i>file</i> ; see <i>get</i> .						
37867 OPER	ANDS							
37868	The follow	ing operands shall be supported:						
37869	file	A pathname of an existing SCCS file. If exactly one file operand appears, and it is						
37870		'-', the standard input shall be read: each line shall be independently processed						
37871 37872		as if it were a command line argument list. (However, the line is not subjected to any of the shell word expansions, such as parameter expansion or quote removal.)						
	т	any of the shell word expansions, such as parameter expansion of quote removal.						
37873 STDIN 37874		rd input shall be a text file used only when the <i>file</i> operand is specified as $'-'$.						
		it input shall be a text life used only when the life operand is specified as						
37875 INPUT 37876		files processed shall be files of an unspecified format.						
37877 ENVIR	ONMENT V	ARIABLES						
37878	The follow	ing environment variables shall affect the execution of <i>val</i> :						
37879	LANG	Provide a default value for the internationalization variables that are unset or null.						
37880		(See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,						
37881		Internationalization Variables for the precedence of internationalization variables						

used to determine the values of locale categories.)

internationalization variables.

 LC_ALL

37882

37883

37884

If set to a non-empty string value, override the values of all the other

Utilities val

37885 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 37886 characters (for example, single-byte as opposed to multi-byte characters in arguments and input files). 37887 LC MESSAGES 37888 Determine the locale that should be used to affect the format and contents of 37889 diagnostic messages written to standard error, and informative messages written 37890 to standard output. 37891 **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 37892 37893 ASYNCHRONOUS EVENTS 37894 Default. 37895 **STDOUT** The standard output shall consist of informative messages about either: 37896 1. Each file processed 37897 Each command line read from standard input 37898 If the standard input is not used, for each file operand yielding a discrepancy, the output line 37899 shall have the following format: 37900 "%s: %s\n", <pathname>, <unspecified string> 37901 If standard input is used, a line of input shall be written before each of the preceding lines for 37902 37903 files containing discrepancies: "%s:\n", <input line> 37904 **37905 STDERR** Not used. 37906 37907 OUTPUT FILES 37908 None 37909 EXTENDED DESCRIPTION None. 37910 37911 EXIT STATUS The 8-bit code returned by val shall be a disjunction of the possible errors; that is, it can be 37912 37913 interpreted as a bit string where set bits are interpreted as follows: 0x80Missing file argument. 37914 = 37915 0x40Unknown or duplicate option. = 0x2037916 Corrupted SCCS file. 0x10Cannot open file or file not SCCS. 37917 = 0x08SID is invalid or ambiguous. 37918 = 0x04SID does not exist. 37919 = 0x02%**Y**%, –**y** mismatch. 37920 = 0x01%**M**%, –**m** mismatch. 37921 Note that val can process two or more files on a given command line and can process multiple 37922 command lines (when reading the standard input). In these cases an aggregate code shall be 37923

returned: a logical OR of the codes generated for each command line and file processed.

val Utilities

37925 CONSEQUENCES OF ERRORS Default. 37927 APPLICATION USAGE 37928 Since the val exit status sets the 0x80 bit, shell applications checking "\$?" cannot tell if it 37929 terminated due to a missing file argument or receipt of a signal. 37930 EXAMPLES In a directory with three SCCS files—s.x (of t type "text"), s.y, and s.z (a corrupted file)—the 37931 37932 following command could produce the output shown: 37933 val - <<EOF 37934 -y source s.x 37935 -m y s.y 37936 s.z 37937 EOF 37938 -y source s.x s.x: %Y%, -y mismatch 37939 37940 s.z 37941 s.z: corrupted SCCS file 37942 RATIONALE None. 37943 37944 FUTURE DIRECTIONS None. 37945 37946 SEE ALSO admin, delta, get, prs 37947 37948 CHANGE HISTORY First released in Issue 2. 37949 37950 **Issue 6** The Open Group Corrigendum U025/4 is applied, correcting a typographical error in the EXIT 37951

STATUS.

37953 **NAME**

37954 vi — screen-oriented (visual) display editor

37955 SYNOPSIS

37956 UP vi [-rR] [-c command] [-t tagstring] [-w size] [file ...]
37957

37958 **DESCRIPTION**

This utility shall be provided on systems that both support the User Portability Utilities option and define the POSIX2_CHAR_TERM symbol. On other systems it is optional.

The *vi* (visual) utility is a screen-oriented text editor. Only the open and visual modes of the editor are described in IEEE Std 1003.1-2001; see the line editor *ex* for additional editing capabilities used in *vi*. The user can switch back and forth between *vi* and *ex* and execute *ex* commands from within *vi*.

This reference page uses the term *edit buffer* to describe the current working text. No specific implementation is implied by this term. All editing changes are performed on the edit buffer, and no changes to it shall affect any file until an editor command writes the file.

When using *vi*, the terminal screen acts as a window into the editing buffer. Changes made to the editing buffer shall be reflected in the screen display; the position of the cursor on the screen shall indicate the position within the editing buffer.

Certain terminals do not have all the capabilities necessary to support the complete *vi* definition. When these commands cannot be supported on such terminals, this condition shall not produce an error message such as "not an editor command" or report a syntax error. The implementation may either accept the commands and produce results on the screen that are the result of an unsuccessful attempt to meet the requirements of this volume of IEEE Std 1003.1-2001 or report an error describing the terminal-related deficiency.

37977 OPTIONS

37968

37969

37970

37971

37972

37973

37974 37975

37976

The *vi* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

37980 The following options shall be supported:

 $-\mathbf{c}$ command See the ex command description of the $-\mathbf{c}$ option.

37982 $-\mathbf{r}$ See the ex command description of the $-\mathbf{r}$ option.

37983 $-\mathbf{R}$ See the ex command description of the $-\mathbf{R}$ option.

-t *tagstring* See the *ex* command description of the −**t** option.

37985 -w size See the ex command description of the -w option.

37986 OPERANDS

See the OPERANDS section of the *ex* command for a description of the operands supported by the *vi* command.

37989 **STDIN**

If standard input is not a terminal device, the results are undefined. The standard input consists of a series of commands and input text, as described in the EXTENDED DESCRIPTION section.

If a read from the standard input returns an error, or if the editor detects an end-of-file condition from the standard input, it shall be equivalent to a SIGHUP asynchronous event.

37994 INPUT FILES

See the INPUT FILES section of the *ex* command for a description of the input files supported by the *vi* command.

37997 ENVIRONMENT VARIABLES

See the ENVIRONMENT VARIABLES section of the *ex* command for the environment variables that affect the execution of the *vi* command.

38000 ASYNCHRONOUS EVENTS

See the ASYNCHRONOUS EVENTS section of the *ex* for the asynchronous events that affect the execution of the *vi* command.

38003 STDOUT

38004 If standard output is not a terminal device, undefined results occur.

Standard output may be used for writing prompts to the user, for informational messages, and for writing lines from the file.

38007 STDERR

38008 If standard output is not a terminal device, undefined results occur.

38009 The standard error shall be used only for diagnostic messages.

38010 OUTPUT FILES

See the OUTPUT FILES section of the *ex* command for a description of the output files supported by the *vi* command.

38013 EXTENDED DESCRIPTION

If the terminal does not have the capabilities necessary to support an unspecified portion of the vi definition, implementations shall start initially in ex mode or open mode. Otherwise, after initialization, vi shall be in command mode; text input mode can be entered by one of several commands used to insert or change text. In text input mode, <ESC> can be used to return to command mode; other uses of <ESC> are described later in this section; see **Terminate Command or Input Mode** (on page 996).

38020 Initialization in ex and vi

See **Initialization in ex and vi** (on page 358) for a description of *ex* and *vi* initialization for the *vi* utility.

38023 Command Descriptions in vi

The following symbols are used in this reference page to represent arguments to commands.

buffer See the description of buffer in the EXTENDED DESCRIPTION section of the ex utility; see Command Descriptions in ex (on page 368).

In open and visual mode, when a command synopsis shows both [buffer] and [count] preceding the command name, they can be specified in either order.

38029 count A positive integer used as an optional argument to most commands, either to give a repeat count or as a size. This argument is optional and shall default to 1 unless otherwise specified.

The Synopsis lines for the *vi* commands <control>-G, <control>-L, <control>-R, <control>-], %, &, ^, D, m, M, Q, u, U, and **ZZ** do not have *count* as an optional argument. Regardless, it shall not be an error to specify a *count* to these commands, and any specified *count* shall be ignored.

38032

38033

38036 38037 38038 38039 38040 38041	motion	An optional trailing argues to indicate the region of either one of the commodisted in the following text matched by repeating command specifies the results.	text than delay the classical text (a) and classical text (b) and the classical text (b) and the classical text (b) and the classical text (b) and the classical text (b) and the classical text (b) and the classical text (b) and the classical text (c) and the class	hat sha naracte Each c comm	all be a ers repe of the a nand; ea	ffected be eated or applicable ach com	y the comma one of seven e commands	and. The motion ral other <i>vi</i> con a specifies the re	n can be nmands egion of
38042 38043 38044 38045 38046		Commands that take <i>mo</i> on the circumstances. We within the text region specharacters, only the exact motion command specific	Vhen ecifie ct cha	operat d for tl racters	ing on he com in the	lines, al mand sh specified	ll lines that that the factor of the lines that the	fall partially or ed. When opera	wholly ating on
38047 38048		When commands that not they shall set the current							nmands,
38049		The following command	ls shal	l be va	lid cur	sor motic	on command	ls:	
38050		<apostrophe></apostrophe>	(_	j	Н			
38051		<pre><carriage-return></carriage-return></pre>)	\$	k	L			
38052		<comma></comma>]]	%	1	M			
38053		<control>-H</control>]]		n	N			
38054		<control>-N</control>	{	- ;	t	Т			
38055		<control>-P</control>	}	?	W	W			
38056		<pre><grave accent=""></grave></pre>	,	b	В				
38057		<newline></newline>	+	e	E				
38058		<pre><space></space></pre>	i	f	F				
38059		<zero></zero>	/	h	G				
			•						
38060		Any <i>count</i> that is specifie							
38061		be applied to the motion						n the command	and its
38062		associated motion comm	nand, t	the effe	ect shal	l be mul	tiplicative.		
38063	The follo	owing symbols are used i	n this :	section	to spe	cify loca	tions in the e	edit buffer:	
38064 38065	current o	character The character that is cur	rently	indica	ted by	the curso	or.		
38066 38067 38068 38069	end of a	line The point located bety <newline> of a line. For line.</newline>							_
38070 38071	end of th	e edit buffer The location correspond	ing to	the en	d of the	e last line	e in the edit b	ouffer.	
38072	The foll	owing symbols are used in	_						
38073		In the POSIX locale, vi sh			•	v			
38074 38075		A maximal sequent beginning or end or e					ed and follov	ved by <blank></blank>	s or the

2. One or more sequential blank lines

The first character in the edit buffer

4. The last non-<newline> in the edit buffer

38076

38079	word In th	ne POSIX locale, <i>vi</i> shall recognize five kinds of words:
38080	1.	A maximal sequence of letters, digits, and underscores, delimited at both ends by:
38081		— Characters other than letters, digits, or underscores
38082		— The beginning or end of a line
38083		— The beginning or end of the edit buffer
38084 38085	2.	A maximal sequence of characters other than letters, digits, underscores, or digits, underscores, or
38086		— A letter, digit, underscore
38087		— <blank>s</blank>
38088		— The beginning or end of a line
38089		— The beginning or end of the edit buffer
38090	3.	One or more sequential blank lines
38091	4.	The first character in the edit buffer
38092	5.	The last non- <newline> in the edit buffer</newline>
38093 38094	section bounda A se	ary action boundary is one of the following:
38095	1.	A line whose first character is a <form-feed></form-feed>
38096	2.	A line whose first character is an open curly brace (' $\{\ '\)$
38097 38098	3.	A line whose first character is a period and whose second and third characters match a two-character pair in the sections edit option (see <i>ed</i>)
38099 38100 38101	4.	A line whose first character is a period and whose only other character matches the first character of a two-character pair in the sections edit option, where the second character of the two-character pair is a <space></space>
38102	5.	The first line of the edit buffer
38103 38104 38105	6.	The last line of the edit buffer if the last line of the edit buffer is empty or if it is a <code>]]</code> or <code>}</code> command; otherwise, the last non- <newline> of the last line of the edit buffer</newline>
38106 38107	paragraph bou A pa	ndary aragraph boundary is one of the following:
38108	1.	A section boundary
38109 38110	2.	A line whose first character is a period and whose second and third characters match a two-character pair in the paragraphs edit option (see <i>ed</i>)
38111 38112 38113	3.	A line whose first character is a period and whose only other character matches the first character of a two-character pair in the <i>paragraphs</i> edit option, where the second character of the two-character pair is a <space></space>
38114	4.	One or more sequential blank lines
38115 38116	remembered see See t	arch direction the description of remembered search direction in ed.

sentence boundary

A *sentence boundary* is one of the following:

- 1. A paragraph boundary
- 2. The first non-<blank> that occurs after a paragraph boundary
- 3. The first non-

 slank> that occurs after a period (' . '), exclamation mark (' ! '),

 or question mark ('?'), followed by two <space>s or the end of a line; any

 number of closing parenthesis (')'), closing brackets (']'), double quote ('"'),

 or single quote (''') characters can appear between the punctuation mark and

 the two <space>s or end-of-line

In the remainder of the description of the *vi* utility, the term "buffer line" refers to a line in the edit buffer and the term "display line" refers to the line or lines on the display screen used to display one buffer line. The term "current line" refers to a specific "buffer line".

If there are display lines on the screen for which there are no corresponding buffer lines because they correspond to lines that would be after the end of the file, they shall be displayed as a single tilde (' $^{\circ}$) character, plus the terminating <newline>.

The last line of the screen shall be used to report errors or display informational messages. It shall also be used to display the input for "line-oriented commands" (/, ?, :, and !). When a line-oriented command is executed, the editor shall enter text input mode on the last line on the screen, using the respective command characters as prompt characters. (In the case of the ! command, the associated motion shall be entered by the user before the editor enters text input mode.) The line entered by the user shall be terminated by a <newline>, a non-<control>-V-escaped <carriage-return>, or unescaped <ESC>. It is unspecified if more characters than require a display width minus one column number of screen columns can be entered.

If any command is executed that overwrites a portion of the screen other than the last line of the screen (for example, the *ex* **suspend** or ! commands), other than the *ex* **shell** command, the user shall be prompted for a character before the screen is refreshed and the edit session continued.

<tab>s shall take up the number of columns on the screen set by the **tabstop** edit option (see *ed*), unless there are less than that number of columns before the display margin that will cause the displayed line to be folded; in this case, they shall only take up the number of columns up to that boundary.

The cursor shall be placed on the current line and relative to the current column as specified by each command described in the following sections.

In open mode, if the current line is not already displayed, then it shall be displayed.

In visual mode, if the current line is not displayed, then the lines that are displayed shall be expanded, scrolled, or redrawn to cause an unspecified portion of the current line to be displayed. If the screen is redrawn, no more than the number of display lines specified by the value of the **window** edit option shall be displayed (unless the current line cannot be completely displayed in the number of display lines specified by the **window** edit option) and the current line shall be positioned as close to the center of the displayed lines as possible (within the constraints imposed by the distance of the line from the beginning or end of the edit buffer). If the current line is before the first line in the display and the screen is scrolled, an unspecified portion of the current line shall be placed on the first line of the display. If the current line is after the last line in the display and the screen is scrolled, an unspecified portion of the current line shall be placed on the last line of the display.

In visual mode, if a line from the edit buffer (other than the current line) does not entirely fit into the lines at the bottom of the display that are available for its presentation, the editor may

choose not to display any portion of the line. The lines of the display that do not contain text from the edit buffer for this reason shall each consist of a single '@' character.

In visual mode, the editor may choose for unspecified reasons to not update lines in the display to correspond to the underlying edit buffer text. The lines of the display that do not correctly correspond to text from the edit buffer for this reason shall consist of a single '@' character (plus the terminating <newline>), and the <control>-R command shall cause the editor to update the screen to correctly represent the edit buffer.

Open and visual mode commands that set the current column set it to a column position in the display, and not a character position in the line. In this case, however, the column position in the display shall be calculated for an infinite width display; for example, the column related to a character that is part of a line that has been folded onto additional screen lines will be offset from the display line column where the buffer line begins, not from the beginning of a particular display line.

The display cursor column in the display is based on the value of the current column, as follows, with each rule applied in turn:

- 1. If the current column is after the last display line column used by the displayed line, the display cursor column shall be set to the last display line column occupied by the last non-<newline> in the current line; otherwise, the display cursor column shall be set to the current column.
- 2. If the character of which some portion is displayed in the display line column specified by the display cursor column requires more than a single display line column:
 - a. If in text input mode, the display cursor column shall be adjusted to the first display line column in which any portion of that character is displayed.
 - b. Otherwise, the display cursor column shall be adjusted to the last display line column in which any portion of that character is displayed.

The current column shall not be changed by these adjustments to the display cursor column.

If an error occurs during the parsing or execution of a *vi* command:

- The terminal shall be alerted. Execution of the *vi* command shall stop, and the cursor (for example, the current line and column) shall not be further modified.
- Unless otherwise specified by the following command sections, it is unspecified whether an
 informational message shall be displayed.
- Any partially entered *vi* command shall be discarded.
- If the *vi* command resulted from a **map** expansion, all characters from that **map** expansion shall be discarded, except as otherwise specified by the **map** command (see *ed*).
- If the *vi* command resulted from the execution of a buffer, no further commands caused by the execution of the buffer shall be executed.

38199	Page Backwards
38200	Synopsis: [count] <control>-B</control>
38201 38202	If in open mode, the $<$ control $>$ -B command shall behave identically to the z command. Otherwise, if the current line is the first line of the edit buffer, it shall be an error.
38203 38204	If the window edit option is less than 3, display a screen where the last line of the display shall be some portion of:
38205	(current first line) -1
38206	otherwise, display a screen where the first line of the display shall be some portion of:
38207	(current first line) - count x ((window edit option) -2)
38208 38209	If this calculation would result in a line that is before the first line of the edit buffer, the first line of the display shall display some portion of the first line of the edit buffer.
38210 38211	<i>Current line</i> : If no lines from the previous display remain on the screen, set to the last line of the display; otherwise, set to (<i>line</i> – the number of new lines displayed on this screen).
38212	Current column: Set to non- <blank>.</blank>
38213	Scroll Forward
38214	Synopsis: [count] <control>-D</control>
38215	If the current line is the last line of the edit buffer, it shall be an error.
38216 38217 38218	If no <i>count</i> is specified, <i>count</i> shall default to the <i>count</i> associated with the previous <control>-D or <control>-U command. If there was no previous <control>-D or <control>-U command, <i>count</i> shall default to the value of the scroll edit option.</control></control></control></control>
38219 38220	If in open mode, write lines starting with the line after the current line, until <i>count</i> lines or the last line of the file have been written.
38221 38222	<i>Current line</i> : If the current line + $count$ is past the last line of the edit buffer, set to the last line of the edit buffer; otherwise, set to the current line + $count$.
38223	Current column: Set to non- <blank>.</blank>
38224	Scroll Forward by Line
38225	Synopsis: [count] <control>-E</control>
38226	Display the line count lines after the last line currently displayed.
38227 38228 38229	If the last line of the edit buffer is displayed, it shall be an error. If there is no line <i>count</i> lines after the last line currently displayed, the last line of the display shall display some portion of the last line of the edit buffer.
38230 38231	<i>Current line</i> : Unchanged if the previous current character is displayed; otherwise, set to the first line displayed.

Current column: Unchanged.

38233 **Page Forward** [count] <control>-F 38234 Synopsis: If in open mode, the <control>-F command shall behave identically to the z command. 38235 Otherwise, if the current line is the last line of the edit buffer, it shall be an error. 38236 If the **window** edit option is less than 3, display a screen where the first line of the display shall 38237 38238 be some portion of: (current last line) +1 38239 38240 otherwise, display a screen where the first line of the display shall be some portion of: (current first line) + count x ((window edit option) -2) 38241 38242 If this calculation would result in a line that is after the last line of the edit buffer, the last line of the display shall display some portion of the last line of the edit buffer. 38243 38244 Current line: If no lines from the previous display remain on the screen, set to the first line of the 38245 display; otherwise, set to (*line* + the number of new lines displayed on this screen). 38246 Current column: Set to non-<blank>. **Display Information** 38247 38248 Synopsis: <control>-G 38249 This command shall be equivalent to the *ex* **file** command. **Move Cursor Backwards** 38250 38251 Synopsis: [count] <control>-H 38252 [count] h 38253 the current erase character (see stty) If there are no characters before the current character on the current line, it shall be an error. If 38254 there are less than *count* previous characters on the current line, *count* shall be adjusted to the 38255 38256 number of previous characters on the line. If used as a motion command: 38257 1. The text region shall be from the character before the starting cursor up to and including 38258 the *count*th character before the starting cursor. 38259 38260 2. Any text copied to a buffer shall be in character mode. If not used as a motion command: 38261 38262 Current line: Unchanged.

38263 38264

with the previous current column).

Current column: Set to (column - the number of columns occupied by count characters ending

Move Down 38265 Synopsis: [count] <newline> 38266 38267 [count] <control>-J [count] <control>-M 38268 [count] <control>-N 38269 38270 [count] j 38271 [count] <carriage-return> 38272 [count] + If there are less than *count* lines after the current line in the edit buffer, it shall be an error. 38273 38274 If used as a motion command: 1. The text region shall include the starting line and the next *count* – 1 lines. 38275 38276 2. Any text copied to a buffer shall be in line mode. 38277 If not used as a motion command: Current line: Set to current line+ count. 38278 Current column: Set to non-
 -Solank> for the <carriage-return>, <control>-M, and + commands; 38279 38280 otherwise, unchanged. **Clear and Redisplay** 38281 38282 Synopsis: <control>-L If in open mode, clear the screen and redisplay the current line. Otherwise, clear and redisplay 38283 38284 the screen. Current line: Unchanged. 38285 38286 Current column: Unchanged. 38287 Move Up 38288 Synopsis: [count] <control>-P 38289 [count] k 38290 [count] -If there are less than *count* lines before the current line in the edit buffer, it shall be an error. 38291 If used as a motion command: 38292 38293 1. The text region shall include the starting line and the previous *count* lines. 2. Any text copied to a buffer shall be in line mode. 38294 38295 If not used as a motion command: 38296 Current line: Set to current line – count.

Current column: Set to non-<blank> for the – command; otherwise, unchanged.

38298	Redraw Screen
38299	Synopsis: <control>-R</control>
38300 38301 38302	If any lines have been deleted from the display screen and flagged as deleted on the terminal using the @ convention (see the beginning of the EXTENDED DESCRIPTION section), they shall be redisplayed to match the contents of the edit buffer.
38303 38304	It is unspecified whether lines flagged with @ because they do not fit on the terminal display shall be affected.
38305	Current line: Unchanged.
38306	Current column: Unchanged.
38307	Scroll Backward
38308	Synopsis: [count] <control>-U</control>
38309	If the current line is the first line of the edit buffer, it shall be an error.
38310 38311 38312	If no <i>count</i> is specified, <i>count</i> shall default to the <i>count</i> associated with the previous <control>-D or <control>-U command. If there was no previous <control>-D or <control>-U command, <i>count</i> shall default to the value of the scroll edit option.</control></control></control></control>
38313 38314	<i>Current line</i> : If <i>count</i> is greater than the current line, set to 1; otherwise, set to the current line – <i>count</i> .
38315	Current column: Set to non- <blank>.</blank>
38316	Scroll Backward by Line
38317	Synopsis: [count] <control>-Y</control>
38318	Display the line <i>count</i> lines before the first line currently displayed.
38319 38320 38321	If the current line is the first line of the edit buffer, it shall be an error. If this calculation would result in a line that is before the first line of the edit buffer, the first line of the display shall display some portion of the first line of the edit buffer.
38322 38323	<i>Current line</i> : Unchanged if the previous current character is displayed; otherwise, set to the first line displayed.
38324	Current column: Unchanged.
38325	Edit the Alternate File
38326	Synopsis: <control>-^</control>
38327 38328	This command shall be equivalent to the ex edit command, with the alternate pathname as its argument.
38329	Terminate Command or Input Mode
38330	Synopsis: <esc></esc>
38331 38332	If a partial <i>vi</i> command (as defined by at least one, non- <i>count</i> character) has been entered, discard the <i>count</i> and the command character(s).
38333 38334 38335	Otherwise, if no command characters have been entered, and the <esc> was the result of a map expansion, the terminal shall be alerted and the <esc> character shall be discarded, but it shall not be an error.</esc></esc>

38336 Otherwise, it shall be an error. Current line: Unchanged. 38337 Current column: Unchanged. 38338 38339 Search for tagstring 38340 Synopsis: <control>-] If the current character is not a word or <blank>, it shall be an error. 38341 38342 This command shall be equivalent to the ex tag command, with the argument to that command defined as follows. 38343 If the current character is a <blank>: 38344 1. Skip all

Skip all

blank>s after the cursor up to the end of the line. 38345 38346 2. If the end of the line is reached, it shall be an error. Then, the argument to the ex tag command shall be the current character and all subsequent 38347 characters, up to the first non-word character or the end of the line. 38348 **Move Cursor Forward** 38349 38350 Synopsis: [count] <space> [count] 1 (ell) 38351 If there are less than *count* non-<newline>s after the cursor on the current line, *count* shall be 38352 adjusted to the number of non-<newline>s after the cursor on the line. 38353 38354 If used as a motion command: 38355 1. If the current or *count*th character after the cursor is the last non-<newline> in the line, the 38356 text region shall be comprised of the current character up to and including the last non-<newline> in the line. Otherwise, the text region shall be from the current character up to, 38357 but not including, the *count*th character after the cursor. 38358 2. Any text copied to a buffer shall be in character mode. 38359 If not used as a motion command: 38360 38361 If there are no non-<newline>s after the current character on the current line, it shall be an error. 38362 Current line: Unchanged. Current column: Set to the last column that displays any portion of the countth character after the 38363 current character. 38364 **Replace Text with Results from Shell Command** 38365 38366 [count] ! motion shell-commands <newline> If the motion command is the! command repeated: 38367 1. If the edit buffer is empty and no *count* was supplied, the command shall be the equivalent 38368 of the *ex*:**read!** command, with the text input, and no text shall be copied to any buffer. 38369 38370 2. Otherwise:

an error.

38371 38372 a. If there are less than *count* –1 lines after the current line in the edit buffer, it shall be

38373 b. The text region shall be from the current line up to and including the next *count* –1 38374 38375 Otherwise, the text region shall be the lines in which any character of the text region specified by the motion command appear. 38376 38377 Any text copied to a buffer shall be in line mode. 38378 This command shall be equivalent to the *ex*! command for the specified lines. Move Cursor to End-of-Line 38379 38380 Synopsis: [count] \$ It shall be an error if there are less than (*count* –1) lines after the current line in the edit buffer. 38381 If used as a motion command: 38382 1. If *count* is 1: 38383 a. It shall be an error if the line is empty. 38384 38385 b. Otherwise, the text region shall consist of all characters from the starting cursor to 38386 the last non-<newline> in the line, inclusive, and any text copied to a buffer shall be in character mode. 38387 2. Otherwise, if the starting cursor position is at or before the first non-<blank> in the line, 38388 the text region shall consist of the current and the next count -1 lines, and any text saved to 38389 a buffer shall be in line mode. 38390 3. Otherwise, the text region shall consist of all characters from the starting cursor to the last 38391 non-<newline> in the line that is *count* –1 lines forward from the current line, and any text 38392 copied to a buffer shall be in character mode. 38393 If not used as a motion command: 38394 *Current line*: Set to the *current line* + *count*-1. 38395 Current column: The current column is set to the last display line column of the last non-38396 38397 <newline> in the line, or column position 1 if the line is empty. The current column shall be adjusted to be on the last display line column of the last non-38398 <newline> of the current line as subsequent commands change the current line, until a 38399 command changes the current column. 38400 38401 **Move to Matching Character** 38402 Synopsis: If the character at the current position is not a parenthesis, bracket, or curly brace, search 38403 forward in the line to the first one of those characters. If no such character is found, it shall be an 38404 38405 The matching character shall be the parenthesis, bracket, or curly brace matching the 38406 parenthesis, bracket, or curly brace, respectively, that was at the current position or that was 38407 found on the current line. 38408 Matching shall be determined as follows, for an open parenthesis: 38409

38410 38411 1. Set a counter to 1.

2. Search forwards until a parenthesis is found or the end of the edit buffer is reached.

- 38412 3. If the end of the edit buffer is reached, it shall be an error.
- 38413 4. If an open parenthesis is found, increment the counter by 1.
 - 5. If a close parenthesis is found, decrement the counter by 1.
 - 6. If the counter is zero, the current character is the matching character.

Matching for a close parenthesis shall be equivalent, except that the search shall be backwards, from the starting character to the beginning of the buffer, a close parenthesis shall increment the counter by 1, and an open parenthesis shall decrement the counter by 1.

Matching for brackets and curly braces shall be equivalent, except that searching shall be done for open and close brackets or open and close curly braces. It is implementation-defined whether other characters are searched for and matched as well.

If used as a motion command:

- 1. If the matching cursor was after the starting cursor in the edit buffer, and the starting cursor position was at or before the first non-
blank> non-<newline> in the starting line, and the matching cursor position was at or after the last non-
blank> non-<newline> in the matching line, the text region shall consist of the current line to the matching line, inclusive, and any text copied to a buffer shall be in line mode.
- 2. If the matching cursor was before the starting cursor in the edit buffer, and the starting cursor position was at or after the last non-<blank> non-<newline> in the starting line, and the matching cursor position was at or before the first non-<blank> non-<newline> in the matching line, the text region shall consist of the current line to the matching line, inclusive, and any text copied to a buffer shall be in line mode.
- 3. Otherwise, the text region shall consist of the starting character to the matching character, inclusive, and any text copied to a buffer shall be in character mode.

If not used as a motion command:

Current line: Set to the line where the matching character is located.

Current column: Set to the last column where any portion of the matching character is displayed.

Repeat Substitution

38439 *Synopsis*: &

38414 38415

38416

38417

38418 38419

38420

38421 38422

38423 38424

38425 38426

38427

38428 38429

38430 38431

38432

38433 38434

38435 38436

38437

38438

38440 38441

38442

38444

38445 38446

38447

38448 38449 Repeat the previous substitution command. This command shall be equivalent to the *ex* & command with the current line as its addresses, and without *options*, *count*, or *flags*.

Return to Previous Context at Beginning of Line

38443 Synopsis: 'character

It shall be an error if there is no line in the edit buffer marked by *character*.

If used as a motion command:

- 1. If the starting cursor is after the marked cursor, then the locations of the starting cursor and the marked cursor in the edit buffer shall be logically swapped.
- 2. The text region shall consist of the starting line up to and including the marked line, and any text copied to a buffer shall be in line mode.
- 38450 If not used as a motion command:

Vİ Utilities

Current line: Set to the line referenced by the mark.

38452 Current column: Set to non-

- slank>.

Return to Previous Context

38454 Synopsis: `character

It shall be an error if the marked line is no longer in the edit buffer. If the marked line no longer contains a character in the saved numbered character position, it shall be as if the marked position is the first non-

| the marked line is no longer in the edit buffer. If the marked line no longer contains a character in the saved numbered character position, it shall be as if the marked position is the first non-

| the marked line is no longer in the edit buffer. If the marked line no longer in the edit buffer. If the marked line no longer in the edit buffer. If the marked line no longer contains a character in the saved numbered character position, it shall be as if the marked line no longer contains a character in the saved numbered character position, it shall be as if the marked position is the first non-

| the marked line is no longer in the edit buffer. If the marked line no longer contains a character in the saved numbered character position, it shall be as if the marked line is no longer line is not longer line in the longer line is not longer line in the longer line is not longer line in the longer line is not longer line in the longer line is not longer line in the longer line is not longer line in the longer line is not longer line in the longer line in the longer line is not longer line in the longer line is not longer line in the longer line is not longer line in the longer line in the longer line in the longer line is not longer line in the longer line in th

If used as a motion command:

- 1. It shall be an error if the marked cursor references the same character in the edit buffer as the starting cursor.
- 2. If the starting cursor is after the marked cursor, then the locations of the starting cursor and the marked cursor in the edit buffer shall be logically swapped.
- 3. If the starting line is empty or the starting cursor is at or before the first non-

 <newline> of the starting line, and the marked cursor line is empty or the marked cursor references the first character of the marked cursor line, the text region shall consist of all lines containing characters from the starting cursor to the line before the marked cursor line, inclusive, and any text copied to a buffer shall be in line mode.
- 4. Otherwise, if the marked cursor line is empty or the marked cursor references a character at or before the first non-
blank> non-<newline> of the marked cursor line, the region of text shall be from the starting cursor to the last non-<newline> of the line before the marked cursor line, inclusive, and any text copied to a buffer shall be in character mode.
- 5. Otherwise, the region of text shall be from the starting cursor (inclusive), to the marked cursor (exclusive), and any text copied to a buffer shall be in character mode.

If not used as a motion command:

Current line: Set to the line referenced by the mark.

Current column: Set to the last column in which any portion of the character referenced by the mark is displayed.

Return to Previous Section

Synopsis: [count] [[

Move the cursor backward through the edit buffer to the first character of the previous section boundary, *count* times.

If used as a motion command:

- 1. If the starting cursor was at the first character of the starting line or the starting line was empty, and the first character of the boundary was the first character of the boundary line, the text region shall consist of the current line up to and including the line where the *count*th next boundary starts, and any text copied to a buffer shall be in line mode.
- 2. If the boundary was the last line of the edit buffer or the last non-<newline> of the last line of the edit buffer, the text region shall consist of the last character in the edit buffer up to and including the starting character, and any text saved to a buffer shall be in character mode.

38491 3. Otherwise, the text region shall consist of the starting character up to but not including the 38492 first character in the countth next boundary, and any text copied to a buffer shall be in character mode. 38493 If not used as a motion command: 38494 38495 *Current line*: Set to the line where the *count*th next boundary in the edit buffer starts. Current column: Set to the last column in which any portion of the first character of the countth 38496 38497 next boundary is displayed, or column position 1 if the line is empty. Move to Next Section 38498 38499 Synopsis: [count]]] Move the cursor forward through the edit buffer to the first character of the next section 38500 boundary, count times. 38501 38502 If used as a motion command: 1. If the starting cursor was at the first character of the starting line or the starting line was 38503 38504 empty, and the first character of the boundary was the first character of the boundary line, 38505 the text region shall consist of the current line up to and including the line where the *count*th previous boundary starts, and any text copied to a buffer shall be in line mode. 38506 2. If the boundary was the first line of the edit buffer, the text region shall consist of the first 38507 character in the edit buffer up to but not including the starting character, and any text 38508 38509 copied to a buffer shall be in character mode. 3. Otherwise, the text region shall consist of the first character in the *count*th previous section 38510 boundary up to but not including the starting character, and any text copied to a buffer 38511 shall be in character mode. 38512 If not used as a motion command: 38513 *Current line*: Set to the line where the *count*th previous boundary in the edit buffer starts. 38514 *Current column*: Set to the last column in which any portion of the first character of the *count*th 38515 38516 previous boundary is displayed, or column position 1 if the line is empty. 38517 Synopsis: 38518 38519 If used as a motion command: 1. If the line has no non-
-\text{olank} non-\text{-newline}\text{s, or if the cursor is at the first non-\text{-\text{blank}}\text{>} 38520 non-<newline> of the line, it shall be an error. 38521 If the cursor is before the first non-
blank> non-<newline> of the line, the text region shall 38522 be comprised of the current character, up to, but not including, the first non-
blank> non-38523 <newline> of the line. 38524 3. If the cursor is after the first non-
-
blank> non-<newline> of the line, the text region shall 38525 be from the character before the starting cursor up to and including the first non-<blank> 38526

4. Any text copied to a buffer shall be in character mode.

non-<newline> of the line.

If not used as a motion command:

38527

38528

38530	Current line: Unchanged.
38531	Current column: Set to non- <blank>.</blank>
38532	Current and Line Above
38533	Synopsis: [count] _
38534	If there are less than $count-1$ lines after the current line in the edit buffer, it shall be an error.
38535	If used as a motion command:
38536	1. If <i>count</i> is less than 2, the text region shall be the current line.
38537	2. Otherwise, the text region shall include the starting line and the next $count - 1$ lines.
38538	3. Any text copied to a buffer shall be in line mode.
38539	If not used as a motion command:
38540	<i>Current line</i> : Set to current line + <i>count</i> −1.
38541	Current column: Set to non- <blank>.</blank>
38542	Move Back to Beginning of Sentence
38543	Synopsis: [count] (
38544 38545 38546	Move backward to the beginning of a sentence. This command shall be equivalent to the [command, with the exception that sentence boundaries shall be used instead of section boundaries.
38547	Move Forward to Beginning of Sentence
38548	Synopsis: [count])
38549 38550 38551	Move forward to the beginning of a sentence. This command shall be equivalent to the J command, with the exception that sentence boundaries shall be used instead of section boundaries.
38552	Move Back to Preceding Paragraph
38553	Synopsis: [count] {
38554 38555 38556	Move back to the beginning of the preceding paragraph. This command shall be equivalent to the [[command, with the exception that paragraph boundaries shall be used instead of section boundaries.
38557	Move Forward to Next Paragraph
38558	Synopsis: [count] }
38559 38560 38561	Move forward to the beginning of the next paragraph. This command shall be equivalent to the [] command, with the exception that paragraph boundaries shall be used instead of section boundaries.

Move to Specific Column Position

38563 Synopsis: [count]

38562

38566

38567 38568

38569

38570

38571

38572

38573 38574

38575 38576

38577

38578

For the purposes of this command, lines that are too long for the current display and that have been folded shall be treated as having a single, 1-based, number of columns.

If there are less than *count* columns in which characters from the current line are displayed on the screen, *count* shall be adjusted to be the last column in which any portion of the line is displayed on the screen.

If used as a motion command:

- 1. If the line is empty, or the cursor character is the same as the character on the *count*th column of the line, it shall be an error.
- 2. If the cursor is before the *count*th column of the line, the text region shall be comprised of the current character, up to but not including the character on the *count*th column of the line.
- 3. If the cursor is after the *count*th column of the line, the text region shall be from the character before the starting cursor up to and including the character on the *count*th column of the line.
- 4. Any text copied to a buffer shall be in character mode.
- 38579 If not used as a motion command:
- 38580 *Current line*: Unchanged.
- 38581 *Current column*: Set to the last column in which any portion of the character that is displayed in the *count* column of the line is displayed.

38583 Reverse Find Character

- Synopsis: [count],
- If the last **F**, **f**, **T**, or **t** command was **F**, **f**, **T**, or **t**, this command shall be equivalent to an **f**, **F**, **t**, or **T** command, respectively, with the specified *count* and the same search character.
- If there was no previous **F**, **f**, **T**, or **t** command, it shall be an error.

38588 Repeat

- 38589 Synopsis: [count]
- Repeat the last !, <, >, A, C, D, I, J, O, P, R, S, X, Y, a, c, d, i, o, p, r, s, x, y, or ~ command. It shall be an error if none of these commands have been executed. Commands (other than commands that enter text input mode) executed as a result of map expansions, shall not change the value of the last repeatable command.
- Repeated commands with associated motion commands shall repeat the motion command as well; however, any specified *count* shall replace the *count*(s) that were originally specified to the repeated command or its associated motion command.
- If the motion component of the repeated command is **f**, **F**, **t**, or **T**, the repeated command shall not set the remembered search character for the ; and , commands.
- If the repeated command is **p** or **P**, and the buffer associated with that command was a numeric buffer named with a number less than 9, the buffer associated with the repeated command shall be set to be the buffer named by the name of the previous buffer logically incremented by 1.

If the repeated character is a text input command, the input text associated with that command is repeated literally:

Input characters are neither macro or abbreviation-expanded.

• Input characters are not interpreted in any special way with the exception that <newline>, <carriage-return>, and <control>-T behave as described in **Input Mode Commands in vi** (on page 1022).

Current line: Set as described for the repeated command.

Current column: Set as described for the repeated command.

Find Regular Expression

Synopsis:

If the input line contains no non-<newline>s, it shall be equivalent to a line containing only the last regular expression encountered. The enhanced regular expressions supported by *vi* are described in **Regular Expressions in ex** (on page 391).

Otherwise, the line shall be interpreted as one or more regular expressions, optionally followed by an address offset or a *vi* **z** command.

If the regular expression is not the last regular expression on the line, or if a line offset or **z** command is specified, the regular expression shall be terminated by an unescaped '/' character, which shall not be used as part of the regular expression. If the regular expression is not the first regular expression on the line, it shall be preceded by zero or more
blank>s, a semicolon, zero or more
blank>s, and a leading '/' character, which shall not be interpreted as part of the regular expression. It shall be an error to precede any regular expression with any characters other than these.

Each search shall begin from the character after the first character of the last match (or, if it is the first search, after the cursor). If the **wrapscan** edit option is set, the search shall continue to the character before the starting cursor character; otherwise, to the end of the edit buffer. It shall be an error if any search fails to find a match, and an informational message to this effect shall be displayed.

An optional address offset (see **Addressing in ex** (on page 361)) can be specified after the last regular expression by including a trailing '/' character after the regular expression and specifying the address offset. This offset will be from the line containing the match for the last regular expression specified. It shall be an error if the line offset would indicate a line address less than 1 or greater than the last line in the edit buffer. An address offset of zero shall be supported. It shall be an error to follow the address offset with any other characters than

| Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank | Shank |

If not used as a motion command, an optional z command (see **Redraw Window** (on page 1021)) can be specified after the last regular expression by including a trailing $^\prime$ / $^\prime$ character after the regular expression, zero or more

<code>slank>s</code>, a $^\prime$ z $^\prime$, zero or more

<code>slank>s</code>, an optional new **window** edit option value, zero or more

<code>slank>s</code>, and a location character. The effect shall be as if the z command was executed after the / command. It shall be an error to follow the z command with any other characters than

<code>slank>s</code>.

The remembered search direction shall be set to forward.

If used as a motion command:

1. It shall be an error if the last match references the same character in the edit buffer as the starting cursor.

2. If any address offset is specified, the last match shall be adjusted by the specified offset as described previously.

- 3. If the starting cursor is after the last match, then the locations of the starting cursor and the last match in the edit buffer shall be logically swapped.
- 4. If any address offset is specified, the text region shall consist of all lines containing characters from the starting cursor to the last match line, inclusive, and any text copied to a buffer shall be in line mode.
- 5. Otherwise, if the starting line is empty or the starting cursor is at or before the first non-

 klank> non-<newline> of the starting line, and the last match line is empty or the last

 match starts at the first character of the last match line, the text region shall consist of all

 lines containing characters from the starting cursor to the line before the last match line,

 inclusive, and any text copied to a buffer shall be in line mode.
- 6. Otherwise, if the last match line is empty or the last match begins at a character at or before the first non-

 shank> non-< newline> of the last match line, the region of text shall be from the current cursor to the last non-< newline> of the line before the last match line, inclusive, and any text copied to a buffer shall be in character mode.
- 7. Otherwise, the region of text shall be from the current cursor (inclusive), to the first character of the last match (exclusive), and any text copied to a buffer shall be in character mode.

38665 If not used as a motion command:

Current line: If a match is found, set to the last matched line plus the address offset, if any; otherwise, unchanged.

Current column: Set to the last column on which any portion of the first character in the last matched string is displayed, if a match is found; otherwise, unchanged.

Move to First Character in Line

Synopsis: 0 (zero)

Move to the first character on the current line. The character '0' shall not be interpreted as a command if it is immediately preceded by a digit.

If used as a motion command:

- 1. If the cursor character is the first character in the line, it shall be an error.
- 2. The text region shall be from the character before the cursor character up to and including the first character in the line.
- 3. Any text copied to a buffer shall be in character mode.
- 38679 If not used as a motion command:
- *Current line*: Unchanged.
- *Current column*: The last column in which any portion of the first character in the line is displayed, or if the line is empty, unchanged.

38683 **Execute an ex Command** Synopsis: 38684 Execute one or more ex commands. 38685 If any portion of the screen other than the last line of the screen was overwritten by any ex 38686 command (except **shell**), vi shall display a message indicating that it is waiting for an input from 38687 the user, and shall then read a character. This action may also be taken for other, unspecified 38688 38689 reasons. If the next character entered is a ':', another ex command shall be accepted and executed. Any 38690 38691 other character shall cause the screen to be refreshed and *vi* shall return to command mode. Current line: As specified for the ex command. 38692 Current column: As specified for the ex command. 38693 38694 Repeat Find Synopsis: 38695 [count]; This command shall be equivalent to the last F, f, T, or t command, with the specified *count*, and 38696 with the same search character used for the last **F**, **f**, **T**, or **t** command. If there was no previous **F**, 38697 f. T. or t command, it shall be an error. 38698 Shift Left 38699 38700 Synopsis: [count] < motion 38701 If the motion command is the < command repeated: 1. If there are less than *count* –1 lines after the current line in the edit buffer, it shall be an 38702 38703 error. 2. The text region shall be from the current line, up to and including the next *count* –1 lines. 38704 Shift any line in the text region specified by the *count* and motion command one shiftwidth (see 38705 the ex shiftwidth option) toward the start of the line, as described by the ex < command. The 38706 38707 unshifted lines shall be copied to the unnamed buffer in line mode. 38708 Current line: If the motion was from the current cursor position toward the end of the edit 38709 buffer, unchanged. Otherwise, set to the first line in the edit buffer that is part of the text region specified by the motion command. 38710 38711 Current column: Set to non-
 - slank>. 38712 **Shift Right** Synopsis: [count] > motion 38713 38714 If the motion command is the > command repeated: 1. If there are less than *count* –1 lines after the current line in the edit buffer, it shall be an 38715 38716 error.

38717

38718 38719

38720

2. The text region shall be from the current line, up to and including the next *count* –1 lines.

Shift any line with characters in the text region specified by the *count* and motion command one

shiftwidth (see the ex shiftwidth option) away from the start of the line, as described by the ex >

command. The unshifted lines shall be copied into the unnamed buffer in line mode.

38721 Current line: If the motion was from the current cursor position toward the end of the edit 38722 buffer, unchanged. Otherwise, set to the first line in the edit buffer that is part of the text region 38723 specified by the motion command. 38724 Current column: Set to non-

- slank>. Scan Backwards for Regular Expression 38725 38726 Synopsis: 38727 Scan backwards; the ? command shall be equivalent to the / command (see Find Regular 38728 **Expression** (on page 1004)) with the following exceptions: 1. The input prompt shall be a '?'. 38729 38730 2. Each search shall begin from the character before the first character of the last match (or, if it is the first search, the character before the cursor character). 38731 38732 The search direction shall be from the cursor toward the beginning of the edit buffer, and 38733 the wrapscan edit option shall affect whether the search wraps to the end of the edit buffer and continues. 38734 4. The remembered search direction shall be set to backward. 38735 38736 Execute @buffer 38737 Synopsis: If the buffer is specified as @, the last buffer executed shall be used. If no previous buffer has been 38738 executed, it shall be an error. 38739 38740 Behave as if the contents of the named buffer were entered as standard input. After each line of a 38741 line-mode buffer, and all but the last line of a character mode buffer, behave as if a <newline> were entered as standard input. 38742 38743 If an error occurs during this process, an error message shall be written, and no more characters 38744 resulting from the execution of this command shall be processed. If a count is specified, behave as if that count were entered as user input before the characters 38745 from the @ buffer were entered. 38746 *Current line*: As specified for the individual commands. 38747 *Current column*: As specified for the individual commands. 38748 Reverse Case 38749 [count] ~ 38750 Synopsis: Reverse the case of the current character and the next *count* –1 characters, such that lowercase 38751 characters that have uppercase counterparts shall be changed to uppercase characters, and 38752 38753 uppercase characters that have lowercase counterparts shall be changed to lowercase characters, as prescribed by the current locale. No other characters shall be affected by this command. 38754 If there are less than *count* –1 characters after the cursor in the edit buffer, *count* shall be adjusted 38755 to the number of characters after the cursor in the edit buffer minus 1. 38756 38757 For the purposes of this command, the next character after the last non-<newline> on the line shall be the next character in the edit buffer. 38758

Current line: Set to the line including the (*count*–1)th character after the cursor.

38760 Current column: Set to the last column in which any portion of the (count-1)th character after the 38761 cursor is displayed. **Append** 38762 38763 Synopsis: [count] a Enter text input mode after the current cursor position. No characters already in the edit buffer 38764 shall be affected by this command. A *count* shall cause the input text to be appended *count* –1 38765 more times to the end of the input. 38766 38767 Current line/column: As specified for the text input commands (see Input Mode Commands in vi (on page 1022)). 38768 38769 **Append at End-of-Line** Synopsis: 38770 [count] A This command shall be equivalent to the *vi* command: 38771 38772 \$ [count] a 38773 (see **Append**). Move Backward to Preceding Word 38774 Synopsis: 38775 [count] b With the exception that words are used as the delimiter instead of bigwords, this command shall 38776 be equivalent to the **B** command. 38777 **Move Backward to Preceding Bigword** 38778 Synopsis: [count] B 38779 If the edit buffer is empty or the cursor is on the first character of the edit buffer, it shall be an 38780 error. If less than count bigwords begin between the cursor and the start of the edit buffer, count 38781 shall be adjusted to the number of bigword beginnings between the cursor and the start of the 38782 38783 edit buffer. If used as a motion command: 38784 1. The text region shall be from the first character of the *count*th previous bigword beginning 38785 up to but not including the cursor character. 38786 2. Any text copied to a buffer shall be in character mode. 38787 If not used as a motion command: 38788 *Current line*: Set to the line containing the *current column*. 38789 Current column: Set to the last column upon which any part of the first character of the countth 38790

38791

previous bigword is displayed.

38792	Change					
38793	Synopsis: [buffer][count] c motion					
38794	If the motion command is the c command repeated:					
38795	1. The buffer text shall be in line mode.					
38796 38797	2. If there are less than $count-1$ lines after the current line in the edit buffer, it shall be at error.					
38798	3. The text region shall be from the current line up to and including the next $count-1$ lines.					
38799	Otherwise, the buffer text mode and text region shall be as specified by the motion command.					
38800 38801 38802	The replaced text shall be copied into <i>buffer</i> , if specified, and into the unnamed buffer. If the text to be replaced contains characters from more than a single line, or the buffer text is in line mode, the replaced text shall be copied into the numeric buffers as well.					
38803	If the buffer text is in line mode:					
38804 38805	 Any lines that contain characters in the region shall be deleted, and the editor shall enter text input mode at the beginning of a new line which shall replace the first line deleted. 					
38806 38807	If the autoindent edit option is set, autoindent characters equal to the autoindent characters on the first line deleted shall be inserted as if entered by the user.					
38808	Otherwise, if characters from more than one line are in the region of text:					
38809	1. The text shall be deleted.					
38810 38811	Any text remaining in the last line in the text region shall be appended to the first line in the region, and the last line in the region shall be deleted.					
38812 38813	The editor shall enter text input mode after the last character not deleted from the first line in the text region, if any; otherwise, on the first column of the first line in the region.					
38814	Otherwise:					
38815 38816	1. If the glyph for ' $$'$ ' is smaller than the region, the end of the region shall be marked with a ' $$'$ '.					
38817	2. The editor shall enter text input mode, overwriting the region of text.					
38818 38819	Current line/column: As specified for the text input commands (see Input Mode Commands in vi (on page 1022)).					
38820	Change to End-of-Line					
38821	Synopsis: [buffer] [count] C					
38822	This command shall be equivalent to the <i>vi</i> command:					

[buffer][count] c\$

See the c command.

38825	Delete					
38826	Synopsis: [buffer] [count] d motion					
38827	If the motion command is the d command repeated:					
38828	1. The buffer text shall be in line mode.					
38829 38830	If there are less than $count-1$ lines after the current line in the edit buffer, it shall be an error.					
38831	3. The text region shall be from the current line up to and including the next $count-1$ lines.					
38832	Otherwise, the buffer text mode and text region shall be as specified by the motion command.					
38833 38834	If in open mode, and the current line is deleted, and the line remains on the display, an '@' character shall be displayed as the first glyph of that line.					
38835 38836 38837	Delete the region of text into <i>buffer</i> , if specified, and into the unnamed buffer. If the text to be deleted contains characters from more than a single line, or the buffer text is in line mode, the deleted text shall be copied into the numeric buffers, as well.					
38838 38839 38840	<i>Current line</i> : Set to the first text region line that appears in the edit buffer, unless that line has been deleted, in which case it shall be set to the last line in the edit buffer, or line 1 if the edit buffer is empty.					
38841	Current column:					
38842	1. If the line is empty, set to column position 1.					
38843 38844	2. Otherwise, if the buffer text is in line mode or the motion was from the cursor toward the end of the edit buffer:					
38845 38846	a. If a character from the current line is displayed in the current column, set to the last column that displays any portion of that character.					
38847 38848	b. Otherwise, set to the last column in which any portion of any character in the line is displayed.					
38849 38850	3. Otherwise, if a character is displayed in the column that began the text region, set to the last column that displays any portion of that character.					
38851 38852	4. Otherwise, set to the last column in which any portion of any character in the line is displayed.					
38853	Delete to End-of-Line					
38854	Synopsis: [buffer] D					
38855	Delete the text from the current position to the end of the current line; equivalent to the <i>vi</i>					

38856

38857

command:
[buffer] d\$

Move to End-of-Word

38858

38893

38859	Synopsis: [count] e					
38860 38861	With the exception that words are used instead of bigwords as the delimiter, this command shall be equivalent to the ${\bf E}$ command.					
38862	Move to End-of-Bigword					
38863	Synopsis: [count] E					
38864 38865 38866	If the edit buffer is empty it shall be an error. If less than <i>count</i> bigwords end between the cursor and the end of the edit buffer, <i>count</i> shall be adjusted to the number of bigword endings between the cursor and the end of the edit buffer.					
38867	If used as a motion command:					
38868 38869	1. The text region shall be from the last character of the <i>count</i> th next bigword up to and including the cursor character.					
38870	2. Any text copied to a buffer shall be in character mode.					
38871	If not used as a motion command:					
38872	Current line: Set to the line containing the current column.					
38873 38874	<i>Current column</i> : Set to the last column upon which any part of the last character of the <i>count</i> th next bigword is displayed.					
38875	Find Character in Current Line (Forward)					
38876	Synopsis: [count] f character					
38877	It shall be an error if <i>count</i> occurrences of the character do not occur after the cursor in the line.					
38878	If used as a motion command:					
38879 38880	 The text range shall be from the cursor character up to and including the countth occurrence of the specified character after the cursor. 					
38881	2. Any text copied to a buffer shall be in character mode.					
38882	If not used as a motion command:					
38883	Current line: Unchanged.					
38884 38885	<i>Current column</i> : Set to the last column in which any portion of the <i>count</i> th occurrence of the specified character after the cursor appears in the line.					
38886	Find Character in Current Line (Reverse)					
38887	Synopsis: [count] F character					
38888	It shall be an error if <i>count</i> occurrences of the character do not occur before the cursor in the line.					
38889	If used as a motion command:					
38890 38891	1. The text region shall be from the <i>count</i> th occurrence of the specified character before the cursor, up to, but not including the cursor character.					
38892	2. Any text copied to a buffer shall be in character mode.					

If not used as a motion command:

38894 Current line: Unchanged. Current column: Set to the last column in which any portion of the countth occurrence of the 38895 38896 specified character before the cursor appears in the line. Move to Line 38897 Synopsis: [count] G 38898 38899 If *count* is not specified, it shall default to the last line of the edit buffer. If *count* is greater than the last line of the edit buffer, it shall be an error. 38900 38901 If used as a motion command: 1. The text region shall be from the cursor line up to and including the specified line. 38902 2. Any text copied to a buffer shall be in line mode. 38903 If not used as a motion command: 38904 *Current line*: Set to *count* if *count* is specified; otherwise, the last line. 38905 Current column: Set to non-<blank>. 38906 Move to Top of Screen 38907 Synopsis: 38908 [count] H 38909 If the beginning of the line *count* greater than the first line of which any portion appears on the display does not exist, it shall be an error. 38910 38911 If used as a motion command: 1. If in open mode, the text region shall be the current line. 38912 38913 2. Otherwise, the text region shall be from the starting line up to and including (the first line of the display + count - 1). 38914 38915 3. Any text copied to a buffer shall be in line mode. 38916 If not used as a motion command: 38917 If in open mode, this command shall set the current column to non-

-| shank > and do nothing else. 38918 Otherwise, it shall set the current line and current column as follows. *Current line*: Set to (the first line of the display + *count* -1). 38919 38920 Current column: Set to non-

- slank>. **Insert Before Cursor** 38921 Synopsis: [count] i 38922 Enter text input mode before the current cursor position. No characters already in the edit buffer 38923 shall be affected by this command. A *count* shall cause the input text to be appended *count* –1 38924 38925 more times to the end of the input. Current line/column: As specified for the text input commands (see Input Mode Commands in vi 38926 (on page 1022)). 38927

38928	Insert at Beginning of Line					
38929	Synopsis: [count] I					
38930	This command shall be equivalent to the vi command $\hat{\ }[count]i$.					
38931	Join					
38932	Synopsis: [count] J					
38933	If the current line is the last line in the edit buffer, it shall be an error.					
38934 38935 38936 38937	This command shall be equivalent to the <i>ex</i> join command with no addresses, and an <i>ex</i> command <i>count</i> value of 1 if <i>count</i> was not specified or if a <i>count</i> of 1 was specified, and an <i>ex</i> command <i>count</i> value of <i>count</i> –1 for any other value of <i>count</i> , except that the current line and column shall be set as follows.					
38938	Current line: Unchanged.					
38939 38940 38941	<i>Current column</i> : The last column in which any portion of the character following the last character in the initial line is displayed, or the last non- <newline> in the line if no characters were appended.</newline>					
38942	Move to Bottom of Screen					
38943	Synopsis: [count] L					
38944 38945	If the beginning of the line <i>count</i> less than the last line of which any portion appears on the display does not exist, it shall be an error.					
38946	If used as a motion command:					
38947	1. If in open mode, the text region shall be the current line.					
38948 38949	2. Otherwise, the text region shall include all lines from the starting cursor line to (the last line of the display $-(count - 1)$).					
38950	3. Any text copied to a buffer shall be in line mode.					
38951	If not used as a motion command:					
38952 38953	 If in open mode, this command shall set the current column to non-<blank> and do nothing else.</blank> 					
38954	2. Otherwise, it shall set the current line and current column as follows.					
38955	Current line: Set to (the last line of the display $-(count - 1)$).					
38956	Current column: Set to non- <blank>.</blank>					
38957	Mark Position					
38958	Synopsis: m letter					
38959	This command shall be equivalent to the ex mark command with the specified character as an					

38960

argument.

38961 Move to Middle of Screen Synopsis: 38962 M The middle line of the display shall be calculated as follows: 38963 (the top line of the display) + (((number of lines displayed) +1) /2) -1 38964 If used as a motion command: 38965 38966 1. If in open mode, the text region shall be the current line. 2. Otherwise, the text region shall include all lines from the starting cursor line up to and 38967 38968 including the middle line of the display. 3. Any text copied to a buffer shall be in line mode. 38969 38970 If not used as a motion command: If in open mode, this command shall set the current column to non-
 -| shank > and do nothing else. 38971 38972 Otherwise, it shall set the current line and current column as follows. 38973 *Current line*: Set to the middle line of the display. *Current column*: Set to non-<blank>. 38974 Repeat Regular Expression Find (Forward) 38975 38976 Synopsis: If the remembered search direction was forward, the $\bf n$ command shall be equivalent to the vi38977 38978 command with no characters entered by the user. Otherwise, it shall be equivalent to the vi? 38979 command with no characters entered by the user. If the n command is used as a motion command for the! command, the editor shall not enter 38980 38981 text input mode on the last line on the screen, and shall behave as if the user entered a single '!' character as the text input. 38982 38983 Repeat Regular Expression Find (Reverse) Synopsis: 38984 Scan for the next match of the last pattern given to / or ?, but in the reverse direction; this is the 38985 reverse of **n**. 38986 If the remembered search direction was forward, the N command shall be equivalent to the vi? 38987 command with no characters entered by the user. Otherwise, it shall be equivalent to the vi / 38988 command with no characters entered by the user. If the N command is used as a motion 38989 command for the! command, the editor shall not enter text input mode on the last line on the 38990 screen, and shall behave as if the user entered a single! character as the text input. 38991 **Insert Empty Line Below** 38992 Synopsis: 38993 Enter text input mode in a new line appended after the current line. A *count* shall cause the input 38994 text to be appended *count* –1 more times to the end of the already added text, each time starting 38995 on a new, appended line. 38996

(on page 1022)).

38997 38998 Current line/column: As specified for the text input commands (see Input Mode Commands in vi

38999 Insert Empty Line Above

Synopsis: ○

39004 39005

39006

39008

39009 39010

39011

39012

39013 39014

39015

39016 39017

39018 39019

39020

39021

39022

39023

39025 39026

39027 39028

39029

39030

39031

39032

39033

39034

39035

39036

39037

39038

39039

Enter text input mode in a new line inserted before the current line. A *count* shall cause the input text to be appended *count* –1 more times to the end of the already added text, each time starting on a new, appended line.

Current line/column: As specified for the text input commands (see **Input Mode Commands in vi** (on page 1022)).

Put from Buffer Following

39007 Synopsis: [buffer] p

If no buffer is specified, the unnamed buffer shall be used.

If the buffer text is in line mode, the text shall be appended below the current line, and each line of the buffer shall become a new line in the edit buffer. A *count* shall cause the buffer text to be appended *count* –1 more times to the end of the already added text, each time starting on a new, appended line.

If the buffer text is in character mode, the text shall be appended into the current line after the cursor, and each line of the buffer other than the first and last shall become a new line in the edit buffer. A *count* shall cause the buffer text to be appended *count* –1 more times to the end of the already added text, each time starting after the last added character.

Current line: If the buffer text is in line mode, set the line to line +1; otherwise, unchanged.

Current column: If the buffer text is in line mode:

- 1. If there is a non-<blank> in the first line of the buffer, set to the last column on which any portion of the first non-
blank> in the line is displayed.
- 2. If there is no non-<blank> in the first line of the buffer, set to the last column on which any portion of the last non-<newline> in the first line of the buffer is displayed.

If the buffer text is in character mode:

- 1. If the text in the buffer is from more than a single line, then set to the last column on which any portion of the first character from the buffer is displayed.
- 2. Otherwise, if the buffer is the unnamed buffer, set to the last column on which any portion of the last character from the buffer is displayed.
- 3. Otherwise, set to the first column on which any portion of the first character from the buffer is displayed.

Put from Buffer Before

Synopsis: [buffer] P

If no *buffer* is specified, the unnamed buffer shall be used.

If the buffer text is in line mode, the text shall be inserted above the current line, and each line of the buffer shall become a new line in the edit buffer. A *count* shall cause the buffer text to be appended *count* –1 more times to the end of the already added text, each time starting on a new, appended line.

If the buffer text is in character mode, the text shall be inserted into the current line before the cursor, and each line of the buffer other than the first and last shall become a new line in the edit buffer. A *count* shall cause the buffer text to be appended *count* –1 more times to the end of the

39040 already added text, each time starting after the last added character.

Current line: Unchanged.

Current column: If the buffer text is in line mode:

1. If there is a non-<blank> in the first line of the buffer, set to the last column on which any portion of that character is displayed.

2. If there is no non-<blank> in the first line of the buffer, set to the last column on which any portion of the last non-<newline> in the first line of the buffer is displayed.

If the buffer text is in character mode:

- 1. If the buffer is the unnamed buffer, set to the last column on which any portion of the last character from the buffer is displayed.
- 2. Otherwise, set to the first column on which any portion of the first character from the buffer is displayed.

Enter ex Mode

Synopsis: Q

Leave visual or open mode and enter *ex* command mode.

Current line: Unchanged.

Current column: Unchanged.

Replace Character

Synopsis: [count] r character

Replace the *count* characters at and after the cursor with the specified character. If there are less than *count* non-<newline>s at and after the cursor on the line, it shall be an error.

If character is <control>-V, any next character other than the <newline> shall be stripped of any special meaning and used as a literal character.

If character is <ESC>, no replacement shall be made and the current line and current column shall be unchanged.

If character is <carriage-return> or <newline>, *count* new lines shall be appended to the current line. All but the last of these lines shall be empty. *count* characters at and after the cursor shall be discarded, and any remaining characters after the cursor in the current line shall be moved to the last of the new lines. If the **autoindent** edit option is set, they shall be preceded by the same number of **autoindent** characters found on the line from which the command was executed.

Current line: Unchanged unless the replacement character is a <carriage-return> or <newline>, in which case it shall be set to line + *count*.

Current column: Set to the last column position on which a portion of the last replaced character is displayed, or if the replacement character caused new lines to be created, set to non-
blank>.

39074 **Replace Characters** 39075 Synopsis: Enter text input mode at the current cursor position possibly replacing text on the current line. A 39076 39077 *count* shall cause the input text to be appended *count* –1 more times to the end of the input. Current line/column: As specified for the text input commands (see Input Mode Commands in vi 39078 (on page 1022)). 39079 **Substitute Character** 39080 39081 Synopsis: [buffer] [count] s This command shall be equivalent to the *vi* command: 39082 39083 [buffer] [count] c<space> **Substitute Lines** 39084 Synopsis: [buffer] [count] S 39085 This command shall be equivalent to the *vi* command: 39086 [buffer][count] c 39087 39088 **Move Cursor to Before Character (Forward)** 39089 Synopsis: [count] t character It shall be an error if *count* occurrences of the character do not occur after the cursor in the line. 39090 If used as a motion command: 39091 39092 The text region shall be from the cursor up to but not including the *count*th occurrence of the specified character after the cursor. 39093 2. Any text copied to a buffer shall be in character mode. 39094 39095 If not used as a motion command: Current line: Unchanged. 39096 Current column: Set to the last column in which any portion of the character before the countth 39097 occurrence of the specified character after the cursor appears in the line. 39098 **Move Cursor to After Character (Reverse)** 39099 Synopsis: [count] T character 39100 It shall be an error if *count* occurrences of the character do not occur before the cursor in the line. 39101 If used as a motion command: 39102 39103 1. If the character before the cursor is the specified character, it shall be an error. The text region shall be from the character before the cursor up to but not including the 39104 *count*th occurrence of the specified character before the cursor. 39105 3. Any text copied to a buffer shall be in character mode. 39106

If not used as a motion command:

39108	Current line: Unchanged.					
39109 39110	<i>Current column</i> : Set to the last column in which any portion of the character after the <i>count</i> th occurrence of the specified character before the cursor appears in the line.					
39111	Undo					
39112	Synopsis: u					
39113 39114	This command shall be equivalent to the <i>ex</i> undo command except that the current line and current column shall be set as follows:					
39115 39116	<i>Current line</i> : Set to the first line added or changed if any; otherwise, move to the line preceding any deleted text if one exists; otherwise, move to line 1.					
39117	Current column: If undoing an ex command, set to the first non- <blank>.</blank>					
39118	Otherwise, if undoing a text input command:					
39119 39120	1. If the command was a C, c, O, o, R, S, or s command, the current column shall be set to the value it held when the text input command was entered.					
39121 39122 39123 39124	2. Otherwise, set to the last column in which any portion of the first character after the deleted text is displayed, or, if no non- <newline>s follow the text deleted from this line, set to the last column in which any portion of the last non-<newline> in the line is displayed, or 1 if the line is empty.</newline></newline>					
39125	Otherwise, if a single line was modified (that is, not added or deleted) by the ${\bf u}$ command:					
39126 39127	 If text was added or changed, set to the last column in which any portion of the first character added or changed is displayed. 					
39128 39129 39130 39131	2. If text was deleted, set to the last column in which any portion of the first character after the deleted text is displayed, or, if no non- <newline>s follow the deleted text, set to the last column in which any portion of the last non-<newline> in the line is displayed, or 1 if the line is empty.</newline></newline>					
39132	Otherwise, set to non- <blank>.</blank>					
39133	Undo Current Line					
39134	Synopsis: U					
39135 39136	Restore the current line to its state immediately before the most recent time that it became the current line.					
39137	Current line: Unchanged.					
39138 39139	<i>Current column</i> : Set to the first column in the line in which any portion of the first character in the line is displayed.					
39140	Move to Beginning of Word					
39141	Synopsis: [count] w					
39142	With the exception that words are used as the delimiter instead of bigwords, this command shall					

be equivalent to the **W** command.

39144	Move to Beginning of Bigword					
39145	Synopsis: [count] W					
39146 39147 39148	If the edit buffer is empty, it shall be an error. If there are less than <i>count</i> bigwords between the cursor and the end of the edit buffer, <i>count</i> shall be adjusted to move the cursor to the last bigword in the edit buffer.					
39149	If used as a motion command:					
39150 39151	 If the associated command is c, count is 1, and the cursor is on a <blank>, the region of to shall be the current character and no further action shall be taken.</blank> 					
39152 39153 39154	2. If there are less than <i>count</i> bigwords between the cursor and the end of the edit buffer, the command shall succeed, and the region of text shall include the last character of the edit buffer.					
39155 39156 39157	3. If there are <black>s or an end-of-line that precede the <i>count</i>th bigword, and the associated command is c, the region of text shall be up to and including the last character before the preceding <black>s or end-of-line.</black></black>					
39158 39159 39160	4. If there are <black>s or an end-of-line that precede the bigword, and the associated command is d or y, the region of text shall be up to and including the last <black> before the start of the bigword or end-of-line.</black></black>					
39161	5. Any text copied to a buffer shall be in character mode.					
39162	If not used as a motion command:					
39163	1. If the cursor is on the last character of the edit buffer, it shall be an error.					
39164	Current line: Set to the line containing the current column.					
39165 39166	<i>Current column</i> : Set to the last column in which any part of the first character of the <i>count</i> th next bigword is displayed.					
39167	Delete Character at Cursor					
39168	Synopsis: [buffer][count] x					
39169 39170	Delete the <i>count</i> characters at and after the current character into <i>buffer</i> , if specified, and into the unnamed buffer.					
39171 39172 39173	If the line is empty, it shall be an error. If there are less than <i>count</i> non- <newline>s at and after the cursor on the current line, <i>count</i> shall be adjusted to the number of non-<newline>s at and after the cursor.</newline></newline>					
39174	Current line: Unchanged.					

Current column: If the line is empty, set to column position 1. Otherwise, if there were count or

less non-<newline>s at and after the cursor on the current line, set to the last column that

displays any part of the last non-<newline> of the line. Otherwise, unchanged.

39175

39176

39178 **Delete Character Before Cursor** Synopsis: [buffer][count] X 39179 Delete the *count* characters before the current character into *buffer*, if specified, and into the 39180 unnamed buffer. 39181 If there are no characters before the current character on the current line, it shall be an error. If 39182 there are less than count previous characters on the current line, count shall be adjusted to the 39183 number of previous characters on the line. 39184 Current line: Unchanged. 39185 39186 *Current column*: Set to (current column – the width of the deleted characters). Yank 39187 Synopsis: [buffer] [count] y motion 39188 Copy (yank) the region of text into *buffer*, if specified, and into the unnamed buffer. 39189 If the motion command is the **y** command repeated: 39190 1. The buffer shall be in line mode. 39191 2. If there are less than *count* –1 lines after the current line in the edit buffer, it shall be an 39192 39193 error. 39194 3. The text region shall be from the current line up to and including the next *count* –1 lines. Otherwise, the buffer text mode and text region shall be as specified by the motion command. 39195 Current line: If the motion was from the current cursor position toward the end of the edit 39196 buffer, unchanged. Otherwise, set to the first line in the edit buffer that is part of the text region 39197 39198 specified by the motion command. Current column: 39199 1. If the motion was from the current cursor position toward the end of the edit buffer, 39200 39201 unchanged. 2. Otherwise, if the current line is empty, set to column position 1. 39202 Otherwise, set to the last column that displays any part of the first character in the file that 39203 is part of the text region specified by the motion command. 39204 39205 **Yank Current Line** Synopsis: [buffer] [count] Y 39206 This command shall be equivalent to the *vi* command: 39207

39208

[buffer] [count] y_

39209 **Redraw Window** If in open mode, the **z** command shall have the Synopsis: 39210 39211 Synopsis: [count] z 39212 If *count* is not specified, it shall default to the **window** edit option -1. The **z** command shall be equivalent to the ex z command, with a type character of = and a count of count -2, except that 39213 39214 the current line and current column shall be set as follows, and the **window** edit option shall not 39215 be affected. If the calculation for the *count* argument would result in a negative number, the count argument to the exz command shall be zero. A blank line shall be written after the last line 39216 39217 is written. Current line: Unchanged. 39218 39219 Current column: Unchanged. If not in open mode, the **z** command shall have the following Synopsis: 39220 39221 Synopsis: [line] z [count] character 39222 If line is not specified, it shall default to the current line. If line is specified, but is greater than the number of lines in the edit buffer, it shall default to the number of lines in the edit buffer. 39223 If *count* is specified, the value of the **window** edit option shall be set to *count* (as described in the 39224 ex window command), and the screen shall be redrawn. 39225 *line* shall be placed as specified by the following characters: 39226 39227 <newline>, <carriage-return> Place the beginning of the line on the first line of the display. 39228 39229 Place the beginning of the line in the center of the display. The middle line of the display shall be calculated as described for the M command. 39230 Place an unspecified portion of the line on the last line of the display. 39231 39232 If line was specified, equivalent to the <newline> case. If line was not specified, display a screen where the first line of the display shall be (current last line) +1. If there are no lines 39233 39234 after the last line in the display, it shall be an error. If line was specified, display a screen where the last line of the display shall contain an 39235 unspecified portion of the first line of a display that had an unspecified portion of the 39236 specified line on the last line of the display. If this calculation results in a line before the 39237 39238 beginning of the edit buffer, display the first screen of the edit buffer. Otherwise, display a screen where the last line of the display shall contain an unspecified 39239 portion of (current first line –1). If this calculation results in a line before the beginning of 39240 the edit buffer, it shall be an error. 39241 *Current line*: If *line* and the ' ^ ' character were specified: 39242 1. If the first screen was displayed as a result of the command attempting to display lines 39243 before the beginning of the edit buffer: if the first screen was already displayed, 39244 39245 unchanged; otherwise, set to (current first line −1). 2. Otherwise, set to the last line of the display. 39246

Otherwise, if *line* was specified, set to *line*.

If *line* and the '+' character were specified, set to the first line of the display.

39247

39249 Otherwise, unchanged. Current column: Set to non-<blank>. 39250 Exit 39251 39252 Synopsis: ZZ39253 This command shall be equivalent to the ex xit command with no addresses, trailing !, or 39254 filename (see the ex xit command). 39255 **Input Mode Commands in vi** In text input mode, the current line shall consist of zero or more of the following categories, plus 39256 39257 the terminating <newline>: 1. Characters preceding the text input entry point 39258 39259 Characters in this category shall not be modified during text input mode. 2. **autoindent** characters 39260 autoindent characters shall be automatically inserted into each line that is created in text 39261 input mode, either as a result of entering a <newline> or <carriage-return> while in text 39262 input mode, or as an effect of the command itself; for example, O or o (see the ex 39263 39264 **autoindent** command), as if entered by the user. 39265 It shall be possible to erase autoindent characters with the <control>-D command; it is unspecified whether they can be erased by <control>-H, <control>-U, and <control>-W 39266 characters. Erasing any autoindent character turns the glyph into erase-columns and 39267 deletes the character from the edit buffer, but does not change its representation on the 39268 39269 screen. 3. Text input characters 39270 39271 Text input characters are the characters entered by the user. Erasing any text input character turns the glyph into erase-columns and deletes the character from the edit buffer, 39272 39273 but does not change its representation on the screen. Each text input character entered by the user (that does not have a special meaning) shall 39274 39275 be treated as follows: The text input character shall be appended to the last character in the edit buffer 39276 39277 from the first, second, or third categories. b. If there are no erase-columns on the screen, the text input command was the R 39278 command, and characters in the fifth category from the original line follow the 39279 cursor, the next such character shall be deleted from the edit buffer. If the **slowopen** 39280 edit option is not set, the corresponding glyph on the screen shall become erase-39281 columns. 39282 c. If there are erase-columns on the screen, as many columns as they occupy, or as are 39283 necessary, shall be overwritten to display the text input character. (If only part of a 39284 multi-column glyph is overwritten, the remainder shall be left on the screen, and 39285 continue to be treated as erase-columns; it is unspecified whether the remainder of 39286 the glyph is modified in any way.) 39287 d. If additional display line columns are needed to display the text input character: 39288 39289 1. If the **slowopen** edit option is set, the text input characters shall be displayed 39290 on subsequent display line columns, overwriting any characters displayed in

 those columns.

 2. Otherwise, any characters currently displayed on or after the column on the display line where the text input character is to be displayed shall be pushed ahead the number of display line columns necessary to display the rest of the text input character.

39296 4. Erase-columns

Erase-columns are not logically part of the edit buffer, appearing only on the screen, and may be overwritten on the screen by subsequent text input characters. When text input mode ends, all erase-columns shall no longer appear on the screen.

Erase-columns are initially the region of text specified by the **c** command (see **Change** (on page 1009)); however, erasing **autoindent** or text input characters causes the glyphs of the erased characters to be treated as erase-columns.

5. Characters following the text region for the c command, or the text input entry point for all other commands

Characters in this category shall not be modified during text input mode, except as specified in category 3.b. for the **R** text input command, or as
blank>s deleted when a <newline> or <carriage-return> is entered.

It is unspecified whether it is an error to attempt to erase past the beginning of a line that was created by the entry of a <newline> or <carriage-return> during text input mode. If it is not an error, the editor shall behave as if the erasing character was entered immediately after the last text input character entered on the previous line, and all of the non-<newline>s on the current line shall be treated as erase-columns.

When text input mode is entered, or after a text input mode character is entered (except as specified for the special characters below), the cursor shall be positioned as follows:

- 1. On the first column that displays any part of the first erase-column, if one exists
- 2. Otherwise, if the **slowopen** edit option is set, on the first display line column after the last character in the first, second, or third categories, if one exists
- 3. Otherwise, the first column that displays any part of the first character in the fifth category, if one exists
- 4. Otherwise, the display line column after the last character in the first, second, or third categories, if one exists
- 5. Otherwise, on column position 1

The characters that are updated on the screen during text input mode are unspecified, other than that the last text input character shall always be updated, and, if the **slowopen** edit option is not set, the current cursor character shall always be updated.

The following specifications are for command characters entered during text input mode.

NUL 39327 NUL 39328 Synopsis: If the first character of the text input is a NUL, the most recently input text shall be input as if 39329 39330 entered by the user, and then text input mode shall be exited. The text shall be input literally; 39331 that is, characters are neither macro or abbreviation expanded, nor are any characters interpreted in any special manner. It is unspecified whether implementations shall support more than 256 39332 bytes of remembered input text. 39333 <control>-D 39334 Synopsis: <control>-D 39335 39336 The <control>-D character shall have no special meaning when in text input mode for a lineoriented command (see **Command Descriptions in vi** (on page 988)). 39337 This command need not be supported on block-mode terminals. 39338 If the cursor does not follow an autoindent character, or an autoindent character and a '0' or 39339 ' ^ ' character: 39340 1. If the cursor is in column position 1, the <control>-D character shall be discarded and no 39341 39342 further action taken. 2. Otherwise, the <control>-D character shall have no special meaning. 39343 39344 If the last input character was a '0', the cursor shall be moved to column position 1. Otherwise, if the last input character was a ' ^ ', the cursor shall be moved to column position 1. 39345 39346 In addition, the autoindent level for the next input line shall be derived from the same line from which the **autoindent** level for the current input line was derived. 39347 39348 Otherwise, the cursor shall be moved back to the column after the previous shiftwidth (see the 39349 ex shiftwidth command) boundary. 39350 All of the glyphs on columns between the starting cursor position and (inclusively) the ending cursor position shall become erase-columns as described in Input Mode Commands in vi (on 39351 39352 page 1022). Current line: Unchanged. 39353 39354 Current column: Set to 1 if the <control>-D was preceded by a '^' or '0'; otherwise, set to (column -1) - ((column -2) % **shiftwidth**).39355 <control>-H 39356 Synopsis: <control>-H 39357 If in text input mode for a line-oriented command, and there are no characters to erase, text 39358 input mode shall be terminated, no further action shall be done for this command, and the 39359 39360 current line and column shall be unchanged. If there are characters other than autoindent characters that have been input on the current line 39361 before the cursor, the cursor shall move back one character. 39362 Otherwise, if there are autoindent characters on the current line before the cursor, it is 39363 implementation-defined whether the <control>-H command is an error or if the cursor moves 39364 back one **autoindent** character. 39365

39366

39367

Otherwise, if the cursor is in column position 1 and there are previous lines that have been input,

it is implementation-defined whether the <control>-H command is an error or if it is equivalent

39368	to entering <control>-H after the last input character on the previous input line.</control>				
39369	Otherwise, it shall be an error.				
39370 39371 39372	All of the glyphs on columns between the starting cursor position and (inclusively) the ending cursor position shall become erase-columns as described in Input Mode Commands in vi (on page 1022).				
39373 39374 39375	The current erase character (see <i>stty</i>) shall cause an equivalent action to the <control>-H command, unless the previously inserted character was a backslash, in which case it shall be as if the literal current erase character had been inserted instead of the backslash.</control>				
39376 39377	<i>Current line</i> : Unchanged, unless previously input lines are erased, in which case it shall be set to line -1 .				
39378 39379	Current column: Set to the first column that displays any portion of the character backed up over.				
39380	<newline></newline>				
39381 39382 39383 39384	<pre>Synopsis: <newline></newline></pre>				
39385 39386	If input was part of a line-oriented command, text input mode shall be terminated and the command shall continue execution with the input provided.				
39387 39388 39389	Otherwise, terminate the current line. If there are no characters other than autoindent characters on the line, all characters on the line shall be discarded. Otherwise, it is unspecified whether the autoindent characters in the line are modified by entering these characters.				
39390 39391 39392 39393	Continue text input mode on a new line appended after the current line. If the slowopen edit option is set, the lines on the screen below the current line shall not be pushed down, but the first of them shall be cleared and shall appear to be overwritten. Otherwise, the lines of the screen below the current line shall be pushed down.				
39394 39395	If the autoindent edit option is set, an appropriate number of autoindent characters shall be added as a prefix to the line as described by the <i>ex</i> autoindent edit option.				
39396 39397	All columns after the cursor that are erase-columns (as described in Input Mode Commands in vi (on page 1022)) shall be discarded.				
39398 39399	If the autoindent edit option is set, all simmediately following the cursor shall be discarded.				
39400 39401	All remaining characters after the cursor shall be transferred to the new line, positioned after any autoindent characters.				
39402	Current line: Set to current line +1.				
39403 39404 39405	<i>Current column</i> : Set to the first column that displays any portion of the first character after the autoindent characters on the new line, if any, or the first column position after the last autoindent character, if any, or column position 1.				

39406	<control>-T</control>					
39407	Synopsis: <control>-T</control>					
39408 39409	The <control>-T character shall have no special meaning when in text input mode for a line-oriented command (see Command Descriptions in vi (on page 988)).</control>					
39410	This command need not be supported on block-mode terminals.					
39411 39412 39413	Behave as if the user entered the minimum number of <blank>s necessary to move the cursor forward to the column position after the next shiftwidth (see the <i>ex</i> shiftwidth command) boundary.</blank>					
39414	Current line: Unchanged.					
39415	<i>Current column</i> : Set to <i>column</i> + shiftwidth − ((column −1) % shiftwidth).					
39416	<control>-U</control>					
39417	Synopsis: <control>-U</control>					
39418 39419 39420	If there are characters other than autoindent characters that have been input on the current line before the cursor, the cursor shall move to the first character input after the autoindent characters.					
39421 39422 39423	Otherwise, if there are autoindent characters on the current line before the cursor, it is implementation-defined whether the <control>-U command is an error or if the cursor moves to the first column position on the line.</control>					
39424 39425 39426	Otherwise, if the cursor is in column position 1 and there are previous lines that have been input, it is implementation-defined whether the <control>-U command is an error or if it is equivalent to entering <control>-U after the last input character on the previous input line.</control></control>					
39427	Otherwise, it shall be an error.					
39428 39429 39430	All of the glyphs on columns between the starting cursor position and (inclusively) the ending cursor position shall become erase-columns as described in Input Mode Commands in vi (on page 1022).					
39431 39432 39433	The current <i>kill</i> character (see <i>stty</i>) shall cause an equivalent action to the <control>-U command, unless the previously inserted character was a backslash, in which case it shall be as if the literal current <i>kill</i> character had been inserted instead of the backslash.</control>					
39434 39435	<i>Current line</i> : Unchanged, unless previously input lines are erased, in which case it shall be set to line –1.					
39436 39437	<i>Current column</i> : Set to the first column that displays any portion of the last character backed up over.					
39438	<control>-V</control>					
39439 39440	Synopsis: <control>-V <control>-Q</control></control>					
39441 39442 39443 39444 39445	Allow the entry of any subsequent character, other than <code><control>-J</control></code> or the <code><newline></newline></code> , as a literal character, removing any special meaning that it may have to the editor in text input mode. If a <code><control>-V</control></code> or <code><control>-Q</control></code> is entered before a <code><control>-J</control></code> or <code><newline></newline></code> , the <code><control>-V</control></code> or <code><control>-Q</control></code> character shall be discarded, and the <code><control>-J</control></code> or <code><newline></newline></code> shall behave as described in the <code><newline></newline></code> command character during input mode.					

39446 For purposes of the display only, the editor shall behave as if a ' ^' character was entered, and 39447 the cursor shall be positioned as if overwriting the ' ^' character. When a subsequent character is entered, the editor shall behave as if that character was entered instead of the original 39448 <control>-V or <control>-Q character. 39449 39450 Current line: Unchanged. Current column: Unchanged. 39451 <control>-W 39452 39453 Synopsis: <control>-W If there are characters other than autoindent characters that have been input on the current line 39454 39455 before the cursor, the cursor shall move back over the last word preceding the cursor (including 39456 any <blank>s between the end of the last word and the current cursor); the cursor shall not move to before the first character after the end of any **autoindent** characters. 39457 Otherwise, if there are autoindent characters on the current line before the cursor, it is 39458 implementation-defined whether the <control>-W command is an error or if the cursor moves to 39459 39460 the first column position on the line. Otherwise, if the cursor is in column position 1 and there are previous lines that have been input, 39461 it is implementation-defined whether the <control>-W command is an error or if it is equivalent 39462 39463 to entering <control>-W after the last input character on the previous input line. 39464 Otherwise, it shall be an error. 39465 All of the glyphs on columns between the starting cursor position and (inclusively) the ending 39466 cursor position shall become erase-columns as described in Input Mode Commands in vi (on page 1022). 39467 39468 Current line: Unchanged, unless previously input lines are erased, in which case it shall be set to 39469 39470 Current column: Set to the first column that displays any portion of the last character backed up over. 39471 <ESC> 39472 Synopsis: <ESC> 39473 If input was part of a line-oriented command: 39474 39475 1. If *interrupt* was entered, text input mode shall be terminated and the editor shall return to command mode. The terminal shall be alerted. 39476 2. If <ESC> was entered, text input mode shall be terminated and the command shall 39477 continue execution with the input provided. 39478 39479 Otherwise, terminate text input mode and return to command mode. Any autoindent characters entered on newly created lines that have no other non-<newline>s 39480 39481 shall be deleted. Any leading autoindent and

blank>s on newly created lines shall be rewritten to be the 39482

The screen shall be redisplayed as necessary to match the contents of the edit buffer.

minimum number of <blank>s possible.

Current line: Unchanged.

39483

39486 *Current column*:

1. If there are text input characters on the current line, the column shall be set to the last column where any portion of the last text input character is displayed.

- 2. Otherwise, if a character is displayed in the current column, unchanged.
- 39490 3. Otherwise, set to column position 1.

39491 EXIT STATUS

39489

The following exit values shall be returned:

- 39493 0 Successful completion.
- 39494 >0 An error occurred.

39495 CONSEQUENCES OF ERRORS

When any error is encountered and the standard input is not a terminal device file, *vi* shall not write the file or return to command or text input mode, and shall terminate with a non-zero exit status.

Otherwise, when an unrecoverable error is encountered it shall be equivalent to a SIGHUP asynchronous event.

Otherwise, when an error is encountered, the editor shall behave as specified in **Command**Descriptions in vi (on page 988).

39503 APPLICATION USAGE

39504 None.

39505 EXAMPLES

39506 None.

39507 RATIONALE

39508

39509 39510

39511

39513

39514

39515 39516

39517

39518

39519

39520 39521

39522 39523

39524

39525

39526 39527

39528

39529

See the RATIONALE for *ex* for more information on *vi*. Major portions of the *vi* utility specification point to *ex* to avoid inadvertent divergence. While *ex* and *vi* have historically been implemented as a single utility, this is not required by IEEE Std 1003.1-2001.

It is recognized that portions of *vi* would be difficult, if not impossible, to implement satisfactorily on a block-mode terminal, or a terminal without any form of cursor addressing, thus it is not a mandatory requirement that such features should work on all terminals. It is the intention, however, that a *vi* implementation should provide the full set of capabilities on all terminals capable of supporting them.

Historically, *vi* exited immediately if the standard input was not a terminal. IEEE Std 1003.1-2001 permits, but does not require, this behavior. An end-of-file condition is not equivalent to an end-of-file character. A common end-of-file character, <control>-D, is historically a *vi* command.

The text in the STDOUT section reflects the usage of the verb *display* in this section; some implementations of *vi* use standard output to write to the terminal, but IEEE Std 1003.1-2001 does not require that to be the case.

Historically, implementations reverted to open mode if the terminal was incapable of supporting full visual mode. IEEE Std 1003.1-2001 requires this behavior. Historically, the open mode of *vi* behaved roughly equivalently to the visual mode, with the exception that only a single line from the edit buffer (one "buffer line") was kept current at any time. This line was normally displayed on the next-to-last line of a terminal with cursor addressing (and the last line performed its normal visual functions for line-oriented commands and messages). In addition, some few commands behaved differently in open mode than in visual mode. IEEE Std 1003.1-2001 requires conformance to historical practice.

Historically, *ex* and *vi* implementations have expected text to proceed in the usual European/Latin order of left to right, top to bottom. There is no requirement in IEEE Std 1003.1-2001 that this be the case. The specification was deliberately written using words like "before", "after", "first", and "last" in order to permit implementations to support the natural text order of the language.

Historically, lines past the end of the edit buffer were marked with single tilde ($'^{\sim}$) characters; that is, if the one-based display was 20 lines in length, and the last line of the file was on line one, then lines 2-20 would contain only a single $'^{\sim}$ ' character.

Historically, the vi editor attempted to display only complete lines at the bottom of the screen (it did display partial lines at the top of the screen). If a line was too long to fit in its entirety at the bottom of the screen, the screen lines where the line would have been displayed were displayed as single '@' characters, instead of displaying part of the line. IEEE Std 1003.1-2001 permits, but does not require, this behavior. Implementations are encouraged to attempt always to display a complete line at the bottom of the screen when doing scrolling or screen positioning by buffer lines.

Historically, lines marked with '@' were also used to minimize output to dumb terminals over slow lines; that is, changes local to the cursor were updated, but changes to lines on the screen that were not close to the cursor were simply marked with an '@' sign instead of being updated to match the current text. IEEE Std 1003.1-2001 permits, but does not require this feature because it is used ever less frequently as terminals become smarter and connections are faster.

Initialization in ex and vi

 Historically, *vi* always had a line in the edit buffer, even if the edit buffer was "empty". For example:

- 1. The *ex* command = executed from visual mode wrote "1" when the buffer was empty.
- 2. Writes from visual mode of an empty edit buffer wrote files of a single character (a <newline>), while writes from *ex* mode of an empty edit buffer wrote empty files.
- 3. Put and read commands into an empty edit buffer left an empty line at the top of the edit buffer.

For consistency, IEEE Std 1003.1-2001 does not permit any of these behaviors.

Historically, *vi* did not always return the terminal to its original modes; for example, ICRNL was modified if it was not originally set. IEEE Std 1003.1-2001 does not permit this behavior.

Command Descriptions in vi

Motion commands are among the most complicated aspects of vi to describe. With some exceptions, the text region and buffer type effect of a motion command on a vi command are described on a case-by-case basis. The descriptions of text regions in IEEE Std 1003.1-2001 are not intended to imply direction; that is, an inclusive region from line n to line n+5 is identical to a region from line n+5 to line n. This is of more than academic interest—movements to marks can be in either direction, and, if the **wrapscan** option is set, so can movements to search points. Historically, lines are always stored into buffers in text order; that is, from the start of the edit buffer to the end. IEEE Std 1003.1-2001 requires conformance to historical practice.

Historically, command counts were applied to any associated motion, and were multiplicative to any supplied motion count. For example, **2cw** is the same as **c2w**, and **2c3w** is the same as **c6w**. IEEE Std 1003.1-2001 requires this behavior. Historically, *vi* commands that used bigwords, words, paragraphs, and sentences as objects treated groups of empty lines, or lines that contained only
blank>s, inconsistently. Some commands treated them as a single entity, while

others treated each line separately. For example, the w, W, and B commands treated groups of empty lines as individual words; that is, the command would move the cursor to each new empty line. The e and E commands treated groups of empty lines as a single word; that is, the first use would move past the group of lines. The b command would just beep at the user, or if done from the start of the line as a motion command, fail in unexpected ways. If the lines contained only (or ended with)

blank>s, the w and W commands would just beep at the user, the E and e commands would treat the group as a single word, and the B and b commands would treat the lines as individual words. For consistency and simplicity of specification, IEEE Std 1003.1-2001 requires that all vi commands treat groups of empty or blank lines as a single entity, and that movement through lines ending with

blank>s be consistent with other movements.

Historically, *vi* documentation indicated that any number of double quotes were skipped after punctuation marks at sentence boundaries; however, implementations only skipped single quotes. IEEE Std 1003.1-2001 requires both to be skipped.

Historically, the first and last characters in the edit buffer were word boundaries. This historical practice is required by IEEE Std 1003.1-2001.

Historically, *vi* attempted to update the minimum number of columns on the screen possible, which could lead to misleading information being displayed. IEEE Std 1003.1-2001 makes no requirements other than that the current character being entered is displayed correctly, leaving all other decisions in this area up to the implementation.

Historically, lines were arbitrarily folded between columns of any characters that required multiple column positions on the screen, with the exception of tabs, which terminated at the right-hand margin. IEEE Std 1003.1-2001 permits the former and requires the latter. Implementations that do not arbitrarily break lines between columns of characters that occupy multiple column positions should not permit the cursor to rest on a column that does not contain any part of a character.

The historical *vi* had a problem in that all movements were by buffer lines, not by display or screen lines. This is often the right thing to do; for example, single line movements, such as **j** or **k**, should work on buffer lines. Commands like **dj**, or **j**., where . is a change command, only make sense for buffer lines. It is not, however, the right thing to do for screen motion or scrolling commands like <control>-D, <control>-F, and **H**. If the window is fairly small, using buffer lines in these cases can result in completely random motion; for example, **1**<control>-D can result in a completely changed screen, without any overlap. This is clearly not what the user wanted. The problem is even worse in the case of the **H**, **L**, and **M** commands—as they position the cursor at the first non-

-\solon blank> of the line, they may all refer to the same location in large lines, and will result in no movement at all.

In addition, if the line is larger than the screen, using buffer lines can make it impossible to display parts of the line—there are not any commands that do not display the beginning of the line in historical vi, and if both the beginning and end of the line cannot be on the screen at the same time, the user suffers. Finally, the page and half-page scrolling commands historically moved to the first non-
blank> in the new line. If the line is approximately the same size as the screen, this is inadequate because the cursor before and after a <control>-D command will refer to the same location on the screen.

Implementations of *ex* and *vi* exist that do not have these problems because the relevant commands (<control>-B, <control>-D, <control>-F, <control>-U, <control>-Y, <control>-E, H, L, and M) operate on display (screen) lines, not (edit) buffer lines.

IEEE Std 1003.1-2001 does not permit this behavior by default because the standard developers believed that users would find it too confusing. However, historical practice has been relaxed.

For example, ex and vi historically attempted, albeit sometimes unsuccessfully, to never put part of a line on the last lines of a screen; for example, if a line would not fit in its entirety, no part of the line was displayed, and the screen lines corresponding to the line contained single '@' characters. This behavior is permitted, but not required by IEEE Std 1003.1-2001, so that it is possible for implementations to support long lines in small screens more reasonably without changing the commands to be oriented to the display (instead of oriented to the buffer). IEEE Std 1003.1-2001 also permits implementations to refuse to edit any edit buffer containing a line that will not fit on the screen in its entirety.

The display area (for example, the value of the **window** edit option) has historically been "grown", or expanded, to display new text when local movements are done in displays where the number of lines displayed is less than the maximum possible. Expansion has historically been the first choice, when the target line is less than the maximum possible expansion value away. Scrolling has historically been the next choice, done when the target line is less than half a display away, and otherwise, the screen was redrawn. There were exceptions, however, in that *ex* commands generally always caused the screen to be redrawn. IEEE Std 1003.1-2001 does not specify a standard behavior because there may be external issues, such as connection speed, the number of characters necessary to redraw as opposed to scroll, or terminal capabilities that implementations will have to accommodate.

The current line in IEEE Std 1003.1-2001 maps one-to-one to a buffer line in the file. The current column does not. There are two different column values that are described by IEEE Std 1003.1-2001. The first is the current column value as set by many of the *vi* commands. This value is remembered for the lifetime of the editor. The second column value is the actual position on the screen where the cursor rests. The two are not always the same. For example, when the cursor is backed by a multi-column character, the actual cursor position on the screen has historically been the last column of the character in command mode, and the first column of the character in input mode.

Commands that set the current line, but that do not set the current cursor value (for example, \mathbf{j} and \mathbf{k}) attempt to get as close as possible to the remembered column position, so that the cursor tends to restrict itself to a vertical column as the user moves around in the edit buffer. IEEE Std 1003.1-2001 requires conformance to historical practice, requiring that the display location of the cursor on the display line be adjusted from the current column value as necessary to support this historical behavior.

Historically, only a single line (and for some terminals, a single line minus 1 column) of characters could be entered by the user for the line-oriented commands; that is, :, !, /, or ?. IEEE Std 1003.1-2001 permits, but does not require, this limitation.

Historically, "soft" errors in *vi* caused the terminal to be alerted, but no error message was displayed. As a general rule, no error message was displayed for errors in command execution in *vi*, when the error resulted from the user attempting an invalid or impossible action, or when a searched-for object was not found. Examples of soft errors included **h** at the left margin, <control>-B or [[at the beginning of the file, **2G** at the end of the file, and so on. In addition, errors such as %,]], }, N, n, f, F, t, and T failing to find the searched-for object were soft as well. Less consistently, / and ? displayed an error message if the pattern was not found, /, ?, N, and n displayed an error message if no previous regular expression had been specified, and ; did not display an error message if no previous f, F, t, or T command had occurred. Also, behavior in this area might reasonably be based on a runtime evaluation of the speed of a network connection. Finally, some implementations have provided error messages for soft errors in order to assist naive users, based on the value of a verbose edit option. IEEE Std 1003.1-2001 does not list specific errors for which an error message shall be displayed. Implementations should conform to historical practice in the absence of any strong reason to diverge.

Page Backwards

The <code><control>-B</code> and <code><control>-F</code> commands historically considered it an error to attempt to page past the beginning or end of the file, whereas the <code><control>-D</code> and <code><control>-U</code> commands simply moved to the beginning or end of the file. For consistency, IEEE Std 1003.1-2001 requires the latter behavior for all four commands. All four commands still consider it an error if the current line is at the beginning (<code><control>-B</code>, <code><control>-U</code>) or end (<code><control>-F</code>, <code><control>-D</code>) of the file. Historically, the <code><control>-B</code> and <code><control>-F</code> commands skip two lines in order to include overlapping lines when a single command is entered. This makes less sense in the presence of a <code>count</code>, as there will be, by definition, no overlapping lines. The actual calculation used by historical implementations of the <code>vi</code> editor for <code><control>-B</code> was:

This calculation does not work well when intermixing commands with and without counts; for example, **3**<control>-F is not equivalent to entering the <control>-F command three times, and is not reversible by entering the <control>-B command three times. For consistency with other *vi* commands that take counts, IEEE Std 1003.1-2001 requires a different calculation.

Scroll Forward

The 4BSD and System V implementations of *vi* differed on the initial value used by the **scroll** command. 4BSD used:

```
39692 ((window edit option) +1) /2
```

while System V used the value of the **scroll** edit option. The System V version is specified by IEEE Std 1003.1-2001 because the standard developers believed that it was more intuitive and permitted the user a method of setting the scroll value initially without also setting the number of lines that are displayed.

Scroll Forward by Line

Historically, the <control>-E and <control>-Y commands considered it an error if the last and first lines, respectively, were already on the screen. IEEE Std 1003.1-2001 requires conformance to historical practice. Historically, the <control>-E and <control>-Y commands had no effect in open mode. For simplicity and consistency of specification, IEEE Std 1003.1-2001 requires that they behave as usual, albeit with a single line screen.

Clear and Redisplay

The historical <code><control>-L</code> command refreshed the screen exactly as it was supposed to be currently displayed, replacing any <code>'@'</code> characters for lines that had been deleted but not updated on the screen with refreshed <code>'@'</code> characters. The intent of the <code><control>-L</code> command is to refresh when the screen has been accidentally overwritten; for example, by a <code>write</code> command from another user, or modem noise.

39709 Redraw Screen

 The historical <control>-R command redisplayed only when necessary to update lines that had been deleted but not updated on the screen and that were flagged with '@' characters. There is no requirement that the screen be in any way refreshed if no lines of this form are currently displayed. IEEE Std 1003.1-2001 permits implementations to extend this command to refresh lines on the screen flagged with '@' characters because they are too long to be displayed in the current framework; however, the current line and column need not be modified.

Search for tagstring

Historically, the first non-<black> at or after the cursor was the first character, and all subsequent characters that were word characters, up to the end of the line, were included. For example, with the cursor on the leading space or on the '#' character in the text "#bar@", the tag was "#bar". On the character 'b' it was "bar", and on the 'a' it was "ar". IEEE Std 1003.1-2001 requires this behavior.

Replace Text with Results from Shell Command

Historically, the <, >, and ! commands considered most cursor motions other than line-oriented motions an error; for example, the command >/foo<CR> succeeded, while the command >l failed, even though the text region described by the two commands might be identical. For consistency, all three commands only consider entire lines and not partial lines, and the region is defined as any line that contains a character that was specified by the motion.

Move to Matching Character

Other matching characters have been left implementation-defined in order to allow extensions such as matching '<' and '>' for searching HTML, or **#ifdef**, **#else**, and **#endif** for searching C source.

Repeat Substitution

IEEE Std 1003.1-2001 requires that any **c** and **g** flags specified to the previous substitute command be ignored; however, the **r** flag may still apply, if supported by the implementation.

Return to Previous (Context or Section)

The [[,]], (,), {, and } commands are all affected by "section boundaries", but in some historical implementations not all of the commands recognize the same section boundaries. This is a bug, not a feature, and a unique section-boundary algorithm was not described for each command. One special case that is preserved is that the sentence command moves to the end of the last line of the edit buffer while the other commands go to the beginning, in order to preserve the traditional character cut semantics of the sentence command. Historically, *vi* section boundaries at the beginning and end of the edit buffer were the first non-
blank> on the first and last lines of the edit buffer if one exists; otherwise, the last character of the first and last lines of the edit buffer, or the first and the last lines of the edit buffer if they are empty.

Sentence boundaries were problematic in the historical *vi*. They were not only the boundaries as defined for the section and paragraph commands, but they were the first non-
blank> that occurred after those boundaries, as well. Historically, the *vi* section commands were documented as taking an optional window size as a *count* preceding the command. This was not implemented in historical versions, so IEEE Std 1003.1-2001 requires that the *count* repeat the command, for consistency with other *vi* commands.

39753 Repeat

Historically, mapped commands other than text input commands could not be repeated using the **period** command. IEEE Std 1003.1-2001 requires conformance to historical practice.

The restrictions on the interpretation of special characters (for example, <control>-H) in the repetition of text input mode commands is intended to match historical practice. For example, given the input sequence:

iab<control>-H<control>-H<control>-Hdef<escape>

the user should be informed of an error when the sequence is first entered, but not during a command repetition. The character <control>-T is specifically exempted from this restriction. Historical implementations of *vi* ignored <control>-T characters that were input in the original command during command repetition. IEEE Std 1003.1-2001 prohibits this behavior.

Find Regular Expression

Historically, commands did not affect the line searched to or from if the motion command was a search (/, ?, N, n) and the final position was the start/end of the line. There were some special cases and *vi* was not consistent. IEEE Std 1003.1-2001 does not permit this behavior, for consistency. Historical implementations permitted but were unable to handle searches as motion commands that wrapped (that is, due to the edit option **wrapscan**) to the original location. IEEE Std 1003.1-2001 requires that this behavior be treated as an error.

Historically, the syntax "/RE/0" was used to force the command to cut text in line mode. IEEE Std 1003.1-2001 requires conformance to historical practice.

Historically, in open mode, a **z** specified to a search command redisplayed the current line instead of displaying the current screen with the current line highlighted. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior.

Historically, trailing **z** commands were permitted and ignored if entered as part of a search used as a motion command. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior.

Execute an ex Command

Historically, vi implementations restricted the commands that could be entered on the colon command line (for example, **append** and **change**), and some other commands were known to cause them to fail catastrophically. For consistency, IEEE Std 1003.1-2001 does not permit these restrictions. When executing an ex command by entering;, it is not possible to enter a <newline> as part of the command because it is considered the end of the command. A different approach is to enter ex command mode by using the vi Q command (and later resuming visual mode with the ex vi command). In ex command mode, the single-line limitation does not exist. So, for example, the following is valid:

39788 Q 39789 s/break here/break\ 39790 here/ 39791 vi

IEEE Std 1003.1-2001 requires that, if the *ex* command overwrites any part of the screen that would be erased by a refresh, *vi* pauses for a character from the user. Historically, this character could be any character; for example, a character input by the user before the message appeared, or even a mapped character. This is probably a bug, but implementations that have tried to be more rigorous by requiring that the user enter a specific character, or that the user enter a character after the message was displayed, have been forced by user indignation back into

historical behavior. IEEE Std 1003.1-2001 requires conformance to historical practice.

Shift Left (Right)

Execute

Historically, buffers could execute other buffers, and loops, infinite and otherwise, were possible. IEEE Std 1003.1-2001 requires conformance to historical practice. The *buffer syntax of ex is not required in vi, because it is not historical practice and has been used in some vi implementations to support additional scripting languages.

Reverse Case

Historically, the ~ command ignored any associated *count*, and acted only on the characters in the current line. For consistency with other *vi* commands, IEEE Std 1003.1-2001 requires that an associated *count* act on the next *count* characters, and that the command move to subsequent lines if warranted by *count*, to make it possible to modify large pieces of text in a reasonably efficient manner. There exist *vi* implementations that optionally require an associated motion command for the ~ command. Implementations supporting this functionality are encouraged to base it on the **tildedop** edit option and handle the text regions and cursor positioning identically to the **yank** command.

Append

Historically, *counts* specified to the **A**, **a**, **I**, and **i** commands repeated the input of the first line *count* times, and did not repeat the subsequent lines of the input text. IEEE Std 1003.1-2001 requires that the entire text input be repeated *count* times.

Move Backward to Preceding Word

Historically, *vi* became confused if word commands were used as motion commands in empty files. IEEE Std 1003.1-2001 requires that this be an error. Historical implementations of *vi* had a large number of bugs in the word movement commands, and they varied greatly in behavior in the presence of empty lines, "words" made up of a single character, and lines containing only
blank>s. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior.

Change to End-of-Line

Some historical implementations of the **C** command did not behave as described by IEEE Std 1003.1-2001 when the \$ key was remapped because they were implemented by pushing the \$ key onto the input queue and reprocessing it. IEEE Std 1003.1-2001 does not permit this behavior. Historically, the **C**, **S**, and **s** commands did not copy replaced text into the numeric buffers. For consistency and simplicity of specification, IEEE Std 1003.1-2001 requires that they behave like their respective **c** commands in all respects.

39839 Delete

Historically, lines in open mode that were deleted were scrolled up, and an @ glyph written over the beginning of the line. In the case of terminals that are incapable of the necessary cursor motions, the editor erased the deleted line from the screen. IEEE Std 1003.1-2001 requires conformance to historical practice; that is, if the terminal cannot display the '@' character, the line cannot remain on the screen.

Delete to End-of-Line

Some historical implementations of the **D** command did not behave as described by IEEE Std 1003.1-2001 when the \$ key was remapped because they were implemented by pushing the \$ key onto the input queue and reprocessing it. IEEE Std 1003.1-2001 does not permit this behavior.

Join

An historical oddity of *vi* is that the commands **J**, **1J**, and **2J** are all equivalent. IEEE Std 1003.1-2001 requires conformance to historical practice. The *vi* **J** command is specified in terms of the *ex* **join** command with an *ex* command *count* value. The address correction for a *count* that is past the end of the edit buffer is necessary for historical compatibility for both *ex* and *vi*.

Mark Position

Historical practice is that only lowercase letters, plus ' ' and ' ' ', could be used to mark a cursor position. IEEE Std 1003.1-2001 requires conformance to historical practice, but encourages implementations to support other characters as marks as well.

Repeat Regular Expression Find (Forward and Reverse)

Historically, the **N** and **n** commands could not be used as motion components for the **c** command. With the exception of the **cN** command, which worked if the search crossed a line boundary, the text region would be discarded, and the user would not be in text input mode. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior.

Insert Empty Line (Below and Above)

Historically, counts to the **O** and **o** commands were used as the number of physical lines to open, if the terminal was dumb and the **slowopen** option was not set. This was intended to minimize traffic over slow connections and repainting for dumb terminals. IEEE Std 1003.1-2001 does not permit this behavior, requiring that a *count* to the open command behave as for other text input commands. This change to historical practice was made for consistency, and because a superset of the functionality is provided by the **slowopen** edit option.

Put from Buffer (Following and Before)

Historically, *counts* to the **p** and **P** commands were ignored if the buffer was a line mode buffer, but were (mostly) implemented as described in IEEE Std 1003.1-2001 if the buffer was a character mode buffer. Because implementations exist that do not have this limitation, and because pasting lines multiple times is generally useful, IEEE Std 1003.1-2001 requires that *count* be supported for all **p** and **P** commands.

Historical implementations of vi were widely known to have major problems in the \mathbf{p} and \mathbf{P} commands, particularly when unusual regions of text were copied into the edit buffer. The standard developers viewed these as bugs, and they are not permitted for consistency and

39881 simplicity of specification.

Historically, a **P** or **p** command (or an *ex* **put** command executed from open or visual mode)
executed in an empty file, left an empty line as the first line of the file. For consistency and
simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior.

Replace Character

Historically, the **r** command did not correctly handle the *erase* and *word erase* characters as arguments, nor did it handle an associated *count* greater than 1 with a <carriage-return> argument, for which it replaced *count* characters with a single <newline>. IEEE Std 1003.1-2001 does not permit these inconsistencies.

Historically, the **r** command permitted the <control>-V escaping of entered characters, such as <ESC> and the <carriage-return>; however, it required two leading <control>-V characters instead of one. IEEE Std 1003.1-2001 requires that this be changed for consistency with the other text input commands of *vi*.

Historically, it is an error to enter the ${\bf r}$ command if there are less than *count* characters at or after the cursor in the line. While a reasonable and unambiguous extension would be to permit the ${\bf r}$ command on empty lines, it would require that too large a *count* be adjusted to match the number of characters at or after the cursor for consistency, which is sufficiently different from historical practice to be avoided. IEEE Std 1003.1-2001 requires conformance to historical practice.

Replace Characters

Historically, if there were **autoindent** characters in the line on which the **R** command was run, and **autoindent** was set, the first <newline> would be properly indented and no characters would be replaced by the <newline>. Each additional <newline> would replace *n* characters, where *n* was the number of characters that were needed to indent the rest of the line to the proper indentation level. This behavior is a bug and is not permitted by IEEE Std 1003.1-2001.

Undo

Historical practice for cursor positioning after undoing commands was mixed. In most cases, when undoing commands that affected a single line, the cursor was moved to the start of added or changed text, or immediately after deleted text. However, if the user had moved from the line being changed, the column was either set to the first non-

being changed, the column was either set to the first non-

being changed, the column was either set to the first non-

being changed, the column was either set to the first non-

blank>, returned to the origin of the command, or remained unchanged. When undoing commands that affected multiple lines or entire lines, the cursor was moved to the first character in the first line restored. As an example of how inconsistent this was, a search, followed by an o text input command, followed by an undo would return the cursor to the location where the o command was entered, but a cw command followed by an o command followed by an undo would return the cursor to the first non-

blank> of the line. IEEE Std 1003.1-2001 requires the most useful of these behaviors, and discards the least useful, in the interest of consistency and simplicity of specification.

Yank

Historically, the <code>yank</code> command did not move to the end of the motion if the motion was in the forward direction. It moved to the end of the motion if the motion was in the backward direction, except for the <code>_</code> command, or for the <code>G</code> and ' commands when the end of the motion was on the current line. This was further complicated by the fact that for a number of motion commands, the <code>yank</code> command moved the cursor but did not update the screen; for example, a subsequent command would move the cursor from the end of the motion, even though the cursor on the screen had not reflected the cursor movement for the <code>yank</code> command. IEEE Std 1003.1-2001 requires that all <code>yank</code> commands associated with backward motions move the cursor to the end of the motion for consistency, and specifically, to make ' commands as motions consistent with search patterns as motions.

Yank Current Line

Some historical implementations of the \mathbf{Y} command did not behave as described by IEEE Std 1003.1-2001 when the '_' key was remapped because they were implemented by pushing the '_' key onto the input queue and reprocessing it. IEEE Std 1003.1-2001 does not permit this behavior.

Redraw Window

Historically, the **z** command always redrew the screen. This is permitted but not required by IEEE Std 1003.1-2001, because of the frequent use of the **z** command in macros such as **map n nz**. for screen positioning, instead of its use to change the screen size. The standard developers believed that expanding or scrolling the screen offered a better interface for users. The ability to redraw the screen is preserved if the optional new window size is specified, and in the <control>-L and <control>-R commands.

The semantics of \mathbf{z} are confusing at best. Historical practice is that the screen before the screen that ended with the specified line is displayed. IEEE Std 1003.1-2001 requires conformance to historical practice.

Historically, the **z** command would not display a partial line at the top or bottom of the screen. If the partial line would normally have been displayed at the bottom of the screen, the command worked, but the partial line was replaced with '@' characters. If the partial line would normally have been displayed at the top of the screen, the command would fail. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior.

Historically, the **z** command with a line specification of 1 ignored the command. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior.

Historically, the **z** command did not set the cursor column to the first non-

slank> for the character if the first screen was to be displayed, and was already displayed. For consistency and simplicity of specification, IEEE Std 1003.1-2001 does not permit this behavior.

Input Mode Commands in vi

Historical implementations of *vi* did not permit the user to erase more than a single line of input, or to use normal erase characters such as *line erase*, *worderase*, and *erase* to erase **autoindent** characters. As there exist implementations of *vi* that do not have these limitations, both behaviors are permitted, but only historical practice is required. In the case of these extensions, *vi* is required to pause at the **autoindent** and previous line boundaries.

Historical implementations of *vi* updated only the portion of the screen where the current cursor character was displayed. For example, consider the *vi* input keystrokes:

39962 iabcd<escape>0C<tab>

Historically, the <tab> would overwrite the characters "abcd" when it was displayed. Other implementations replace only the 'a' character with the <tab>, and then push the rest of the characters ahead of the cursor. Both implementations have problems. The historical implementation is probably visually nicer for the above example; however, for the keystrokes:

39967 iabcd<ESC>0R<tab><ESC>

the historical implementation results in the string "bcd" disappearing and then magically reappearing when the <ESC> character is entered. IEEE Std 1003.1-2001 requires the former behavior when overwriting erase-columns—that is, overwriting characters that are no longer logically part of the edit buffer—and the latter behavior otherwise.

Historical implementations of *vi* discarded the <control>-D and <control>-T characters when they were entered at places where their command functionality was not appropriate. IEEE Std 1003.1-2001 requires that the <control>-T functionality always be available, and that <control>-D be treated as any other key when not operating on **autoindent** characters.

NUL

 Some historical implementations of *vi* limited the number of characters entered using the NUL input character to 256 bytes. IEEE Std 1003.1-2001 permits this limitation; however, implementations are encouraged to remove this limit.

<control>-D

See also Rationale for the input mode command <newline>. The hidden assumptions in the <control>-D command (and in the *vi* autoindent specification in general) is that <space>s take up a single column on the screen and that <tab>s are comprised of an integral number of <space>s.

<newline>

Implementations are permitted to rewrite **autoindent** characters in the line when <newline>, <carriage-return>, <control>-D, and <control>-T are entered, or when the **shift** commands are used, because historical implementations have both done so and found it necessary to do so. For example, a <control>-D when the cursor is preceded by a single <tab>, with **tabstop** set to 8, and **shiftwidth** set to 3, will result in the <tab> being replaced by several <space>s.

<control>-T

See also the Rationale for the input mode command <newline>. Historically, <control>-T only worked if no non-
blank>s had yet been input in the current input line. In addition, the characters inserted by <control>-T were treated as **autoindent** characters, and could not be erased using normal user erase characters. Because implementations exist that do not have these limitations, and as moving to a column boundary is generally useful, IEEE Std 1003.1-2001 requires that both limitations be removed.

39998	<control>-V</control>					
39999 40000	Historically, <i>vi</i> used ^V , regardless of the value of the literal-next character of the terminal. IEEE Std 1003.1-2001 requires conformance to historical practice.					
40001 40002 40003 40004 40005 40006 40007 40008	The uses described for <control>-V can also be accomplished with <control>-Q, which is useful on terminals that use <control>-V for the down-arrow function. However, most historical implementations use <control>-Q for the <i>termios</i> START character, so the editor will generally not receive the <control>-Q unless stty ixon mode is set to off. (In addition, some historical implementations of <i>vi</i> explicitly set ixon mode to on, so it was difficult for the user to set it to off.) Any of the command characters described in IEEE Std 1003.1-2001 can be made ineffective by their selection as <i>termios</i> control characters, using the <i>stty</i> utility or other methods described</control></control></control></control></control>					
40009	<esc></esc>					
40010 40011	Historically, SIGINT alerted the terminal when used to end input mode. This behavior is permitted, but not required, by IEEE Std 1003.1-2001.					
40012 FUTUR 40013	None.					
40014 SEE AL						
40015	ed, ex, stty					
40016 CHAN 0 40017	GE HISTORY First released in Issue 2.					
40018 Issue 5 40019	The FUTURE DIRECTIONS section is added.					
40020 Issue 6	This sailte is something a some of the Heavy Department of the History and the					
40021	This utility is marked as part of the User Portability Utilities option. The APPLICATION USA CE section is added.					
40022	The APPLICATION USAGE section is added.					
40023	The obsolescent SYNOPSIS is removed.					
40024 40025	The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:					
40026	• The reindent command description is added.					
40027 40028	The <i>vi</i> utility has been extensively rewritten for alignment with the IEEE P1003.2b draft standard.					
40029	IEEE PASC Interpretations 1003.2 #57, #62, #63, #64, #78, and #188 are applied.					
40030 40031	IEEE PASC Interpretation $1003.2~\#207$ is applied, clarifying the description of the R command in a manner similar to the descriptions of other text input mode commands such as i , o , and O .					
40032	The -l option is removed.					
40033 40034	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/41 is applied, adding [count] to the Synopsis for [[.					
40035 40036	IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/42 is applied, adding [count] to the Synopsis for]].					

Utilities wait

40037 NAME				1		
40038	wait — awai	it proc	ess comp	letion		
40039 SYNOP		_				
40040	wait [pid]				
40041 DESCR		1	12.4	(C (1 0.001 (70))		
40042 40043		When an asynchronous list (see Section 2.9.3.1 (on page 50)) is started by the shell, the process ID of the last command in each element of the asynchronous list shall become known in the current				
40044		shell execution environment; see Section 2.12 (on page 61).				
40045 40046		If the <i>wait</i> utility is invoked with no operands, it shall wait until all process IDs known to the invoking shell have terminated and exit with a zero exit status.				
40047	Ü			are specified that represent known process IDs, the wait utility shall		
40048		_	-	terminated. If one or more <i>pid</i> operands are specified that represent		
40049				shall treat them as if they were known process IDs that exited with		
40050 40051		exit status 127. The exit status returned by the <i>wait</i> utility shall be the exit status of the process requested by the last <i>pid</i> operand.				
40052		The known process IDs are applicable only for invocations of <i>wait</i> in the current shell execution				
40053	environment.					
40054 OPTIO	NS					
40055	None.					
40056 OPERA						
40057	The following	ıg ope	rand shal	l be supported:		
40058	pid	One	of the fol	lowing:		
40059 40060		1.		igned decimal integer process ID of a command, for which the utility t for the termination.		
40061		2.		ntrol job ID (see the Base Definitions volume of IEEE Std 1003.1-2001,		
40062				3.203, Job Control Job ID) that identifies a background process group		
40063 40064				raited for. The job control job ID notation is applicable only for ons of <i>wait</i> in the current shell execution environment; see Section		
40065				page 61). The exit status of <i>wait</i> shall be determined by the last		
40066				nd in the pipeline.		
40067 40068			Note:	The job control job ID type of \emph{pid} is only available on systems supporting the User Portability Utilities option.		
40069 STDIN						
40070	Not used.					
40071 INPUT 40072	FILES None.					
40073 ENVIR 040074				variables shall affect the execution of <i>wait</i> :		
40075	LANG	Prov	ide a defa	ault value for the internationalization variables that are unset or null.		
40076				ase Definitions volume of IEEE Std 1003.1-2001, Section 8.2,		
40077				zation Variables for the precedence of internationalization variables		
40078		used	to detern	nine the values of locale categories.)		

internationalization variables.

 LC_ALL

40079 40080 If set to a non-empty string value, override the values of all the other

wait Utilities

40081 *LC_CTYPE* 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).

40084 *LC_MESSAGES*

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

40087 XSI NLSPATH Determine the location of message catalogs for the processing of LC_MESSAGES.

40088 ASYNCHRONOUS EVENTS

40089 Default.

40090 STDOUT

40091 Not used.

40092 STDERR

40093 The standard error shall be used only for diagnostic messages.

40094 OUTPUT FILES

40095 None.

40096 EXTENDED DESCRIPTION

40097 None.

40098 EXIT STATUS

40099 40100

40101 40102

40103

40104

40105

40106

40107

40108

40109

40120 40121

40122

40123

40124

If one or more operands were specified, all of them have terminated or were not known by the invoking shell, and the status of the last operand specified is known, then the exit status of *wait* shall be the exit status information of the command indicated by the last operand specified. If the process terminated abnormally due to the receipt of a signal, the exit status shall be greater than 128 and shall be distinct from the exit status generated by other signals, but the exact value is unspecified. (See the *kill* –l option.) Otherwise, the *wait* utility shall exit with one of the following values:

0 The *wait* utility was invoked with no operands and all process IDs known by the invoking shell have terminated.

1-126 The *wait* utility detected an error.

127 The command identified by the last *pid* operand specified is unknown.

40110 CONSEQUENCES OF ERRORS

40111 Default.

40112 APPLICATION USAGE

40113 On most implementations, *wait* is a shell built-in. If it is called in a subshell or separate utility execution environment, such as one of the following:

```
40115 (wait)
40116 nohup wait ...
40117 find . -exec wait ... \;
```

40118 it returns immediately because there are no known process IDs to wait for in those environments.

Historical implementations of interactive shells have discarded the exit status of terminated background processes before each shell prompt. Therefore, the status of background processes was usually lost unless it terminated while *wait* was waiting for it. This could be a serious problem when a job that was expected to run for a long time actually terminated quickly with a syntax or initialization error because the exit status returned was usually zero if the requested

Utilities wait

process ID was not found. This volume of IEEE Std 1003.1-2001 requires the implementation to keep the status of terminated jobs available until the status is requested, so that scripts like:

```
40127 j1&
40128 p1=$!
40129 j2&
40130 wait $p1
40131 echo Job 1 exited with status $?
40132 wait $!
40133 echo Job 2 exited with status $?
```

work without losing status on any of the jobs. The shell is allowed to discard the status of any process if it determines that the application cannot get the process ID for that process from the shell. It is also required to remember only {CHILD_MAX} number of processes in this way. Since the only way to get the process ID from the shell is by using the '!' shell parameter, the shell is allowed to discard the status of an asynchronous list if "\$!" was not referenced before another asynchronous list was started. (This means that the shell only has to keep the status of the last asynchronous list started if the application did not reference "\$!". If the implementation of the shell is smart enough to determine that a reference to "\$!" was not saved anywhere that the application can retrieve it later, it can use this information to trim the list of saved information. Note also that a successful call to wait with no operands discards the exit status of all asynchronous lists.)

If the exit status of *wait* is greater than 128, there is no way for the application to know if the waited-for process exited with that value or was killed by a signal. Since most utilities exit with small values, there is seldom any ambiguity. Even in the ambiguous cases, most applications just need to know that the asynchronous job failed; it does not matter whether it detected an error and failed or was killed and did not complete its job normally.

40150 EXAMPLES

Although the exact value used when a process is terminated by a signal is unspecified, if it is known that a signal terminated a process, a script can still reliably determine which signal by using *kill* as shown by the following script:

```
40154
             sleep 1000&
             pid=$!
40155
             kill -kill $pid
40156
             wait $pid
40157
             echo $pid was terminated by a SIG$(kill -1 $?) signal.
40158
             If the following sequence of commands is run in less than 31 seconds:
40159
40160
             sleep 257 | sleep 31 &
             jobs -1 %%
40161
             either of the following commands returns the exit status of the second sleep in the pipeline:
40162
             wait <pid of sleep 31>
40163
```

40165 RATIONALE

wait %%

The description of *wait* does not refer to the *waitpid()* function from the System Interfaces volume of IEEE Std 1003.1-2001 because that would needlessly overspecify this interface. However, the wording means that *wait* is required to wait for an explicit process when it is given an argument so that the status information of other processes is not consumed. Historical implementations use the *wait()* function defined in the System Interfaces volume of IEEE Std 1003.1-2001 until *wait()* returns the requested process ID or finds that the requested

wait Utilities

```
40172
              process does not exist. Because this means that a shell script could not reliably get the status of
40173
              all background children if a second background job was ever started before the first job finished,
              it is recommended that the wait utility use a method such as the functionality provided by the
40174
              waitpid() function.
40175
              The ability to wait for multiple pid operands was adopted from the KornShell.
40176
40177
              This new functionality was added because it is needed to determine the exit status of any
              asynchronous list accurately. The only compatibility problem that this change creates is for a
40178
              script like
40179
40180
              while sleep 60 do
                   job& echo Job started $(date) as $! done
40181
40182
              which causes the shell to monitor all of the jobs started until the script terminates or runs out of
              memory. This would not be a problem if the loop did not reference "$!" or if the script would
40183
40184
              occasionally wait for jobs it started.
40185 FUTURE DIRECTIONS
              None.
40186
40187 SEE ALSO
40188
              Chapter 2 (on page 29), kill, sh, the System Interfaces volume of IEEE Std 1003.1-2001, wait(),
40189
              waitpid()
40190 CHANGE HISTORY
              First released in Issue 2.
40191
```

Utilities WC

40192 NAME 40193	wc — word, line, and byte or character count				
40194 SYNOP	•				
40194 311101 40195		[-lw][file]			
40196 DESCR	·				
40197 40198	The <i>wc</i> utilit	y shall read one or more input files and, by default, write the number of <newline>s, bytes contained in each input file to the standard output.</newline>			
40199 40200	The utility a specified.	also shall write a total count for all named files, if more than one input file is			
40201 40202	The <i>wc</i> utility white space.	ty shall consider a word to be a non-zero-length string of characters delimited by			
40203 OPTIO	NS				
40204 40205		y shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, x Guidelines.			
40206	The following	ng options shall be supported:			
40207	-с	Write to the standard output the number of bytes in each input file.			
40208	- l	Write to the standard output the number of <newline>s in each input file.</newline>			
40209	-m	Write to the standard output the number of characters in each input file.			
40210	- w	Write to the standard output the number of words in each input file.			
40211 40212	When any option is specified, wc shall report only the information requested by the specified options.				
40213 OPERA 40214	ERANDS The following operand shall be supported:				
40215 40216	file	A pathname of an input file. If no $\it file$ operands are specified, the standard input shall be used.			
40217 STDIN					
40218 40219	The standard input shall be used only if no <i>file</i> operands are specified. See the INPUT FILES section.				
40220 INPUT	40220 INPUT FILES				
40221	The input files may be of any type.				
40222 ENVIR 40223	22 ENVIRONMENT VARIABLES 23 The following environment variables shall affect the execution of <i>wc</i> :				
40224 40225 40226 40227	LANG	Provide a default value for the internationalization variables that are unset or null (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2 Internationalization Variables for the precedence of internationalization variable used to determine the values of locale categories.)			
40228 40229	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.			
40230 40231 40232	LC_CTYPE	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) and which characters are defined as white space characters			

characters.

Utilities wc

40234 40235 40236 40237	LC_MESSAGES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.			
40238 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .		
40239 ASYNO 40240	CHRONOUS Default.	EVENTS		
40241 STDOU				
40242	By default, t	he standard output shall contain an entry for each input file of the form:		
40243	"%d %d %d	%s\n", <newlines>, <words>, <bytes>, <file></file></bytes></words></newlines>		
40244 40245	If the -m of format.	ption is specified, the number of characters shall replace the <i><bytes></bytes></i> field in this		
40246 40247	If any option not be written	ns are specified and the $-\mathbf{l}$ option is not specified, the number of <newline>s shall en.</newline>		
40248 40249	If any options are specified and the –w option is not specified, the number of words shall not be written.			
40250 40251	If any options are specified and neither $-c$ nor $-m$ is specified, the number of bytes or characters shall not be written.			
40252 40253	If no input <i>file</i> operands are specified, no name shall be written and no <blank>s preceding the pathname shall be written.</blank>			
40254 40255 40256 40257	If more than one input <i>file</i> operand is specified, an additional line shall be written, of the same format as the other lines, except that the word total (in the POSIX locale) shall be written instead of a pathname and the total of each column shall be written as appropriate. Such an additional line, if any, is written at the end of the output.			
40258 STDER				
40259	The standard error shall be used only for diagnostic messages.			
40260 OUTPUT FILES 40261 None.				
40262 EXTEN 40263	40262 EXTENDED DESCRIPTION 40263 None.			
40264 EXIT STATUS				
40265	The following exit values shall be returned:			
40266	0 Success	ful completion.		
40267	>0 An erro	or occurred.		
40268 CONSI 40269	0268 CONSEQUENCES OF ERRORS 0269 Default.			

Utilities WC

40270 APPLICATION USAGE

The $-\mathbf{m}$ option is not a switch, but an option at the same level as $-\mathbf{c}$. Thus, to produce the full

default output with character counts instead of bytes, the command required is:

40273 wc -mlw

40274 EXAMPLES

40275 None.

40276 RATIONALE

40283

40284

40285

40277 The output file format pseudo-*printf()* string differs from the System V version of *wc*:

40278 "%7d%7d%7d %s\n"

which produces possibly ambiguous and unparsable results for very large files, as it assumes no

40280 number shall exceed six digits.

Some historical implementations use only <space>, <tab>, and <newline> as word separators.

40282 The equivalent of the ISO C standard *isspace()* function is more appropriate.

The –c option stands for "character" count, even though it counts bytes. This stems from the sometimes erroneous historical view that bytes and characters are the same size. Due to

international requirements, the -m option (reminiscent of "multi-byte") was added to obtain

40286 actual character counts.

Early proposals only specified the results when input files were text files. The current specification more closely matches historical practice. (Bytes, words, and <newline>s are

counted separately and the results are written when an end-of-file is detected.)

Historical implementations of the wc utility only accepted one argument to specify the options $-\mathbf{c}$, $-\mathbf{l}$, and $-\mathbf{w}$. Some of them also had multiple occurrences of an option cause the

corresponding count to be written multiple times and had the order of specification of the options affect the order of the fields on output, but did not document either of these. Because common usage either specifies no options or only one option, and because none of this was documented, the changes required by this volume of IEEE Std 1003.1-2001 should not break

documented, the changes required by this volume of IEEE sta 1000.1 2001 should not

40296 many historical applications (and do not break any historical conforming applications).

40297 FUTURE DIRECTIONS

40298 None.

40299 SEE ALSO

40300 *cksum*

40301 CHANGE HISTORY

40302 First released in Issue 2.

what

```
40303 NAME
40304
              what — identify SCCS files (DEVELOPMENT)
40305 SYNOPSIS
              what [-s] file...
40306 XSI
40307
40308 DESCRIPTION
              The what utility shall search the given files for all occurrences of the pattern that get (see get)
40309
              substitutes for the %Z% keyword ("@(#)") and shall write to standard output what follows
40310
              until the first occurrence of one of the following:
40311
40312
                        newline
                                           MITT.
40313 OPTIONS
              The what utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section
40314
40315
              12.2, Utility Syntax Guidelines.
40316
              The following option shall be supported:
                           Quit after finding the first occurrence of the pattern in each file.
40317
40318 OPERANDS
40319
              The following operands shall be supported:
              file
40320
                           A pathname of a file to search.
40321 STDIN
40322
              Not used.
40323 INPUT FILES
              The input files shall be of any file type.
40324
40325 ENVIRONMENT VARIABLES
              The following environment variables shall affect the execution of what:
40326
              LANG
40327
                           Provide a default value for the internationalization variables that are unset or null.
                           (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2,
40328
                           Internationalization Variables for the precedence of internationalization variables
40329
                           used to determine the values of locale categories.)
40330
              LC\_ALL
                           If set to a non-empty string value, override the values of all the other
40331
                           internationalization variables.
40332
                           Determine the locale for the interpretation of sequences of bytes of text data as
40333
                           characters (for example, single-byte as opposed to multi-byte characters in
40334
                           arguments and input files).
40335
              LC_MESSAGES
40336
                           Determine the locale that should be used to affect the format and contents of
40337
40338
                           diagnostic messages written to standard error.
              NLSPATH
                           Determine the location of message catalogs for the processing of LC_MESSAGES.
40339
40340 ASYNCHRONOUS EVENTS
              Default.
40341
40342 STDOUT
40343
              The standard output shall consist of the following for each file operand:
```

"%s:\n\t%s\n", <pathname>, <identification string>

Utilities what

```
40345 STDERR
40346
             The standard error shall be used only for diagnostic messages.
40347 OUTPUT FILES
             None.
40348
40349 EXTENDED DESCRIPTION
40350
             None.
40351 EXIT STATUS
40352
             The following exit values shall be returned:
40353
                 Any matches were found.
40354
                 Otherwise.
40355 CONSEQUENCES OF ERRORS
             Default.
40356
40357 APPLICATION USAGE
40358
             The what utility is intended to be used in conjunction with the SCCS command get, which
40359
             automatically inserts identifying information, but it can also be used where the information is
             inserted by any other means.
40360
             When the string "@(#)" is included in a library routine in a shared library, it might not be found
40361
             in an a.out file using that library routine.
40362
40363 EXAMPLES
             If the C-language program in file f.c contains:
40364
40365
             char ident[] = "@(#)identification information";
             and f.c is compiled to yield f.o and a.out, then the command:
40366
             what f.c f.o a.out
40367
             writes:
40368
             f.c:
40369
                  identification information
40370
40371
40372
             f.o:
                  identification information
40373
40374
40375
             a.out:
                  identification information
40376
40377
40378 RATIONALE
40379
             None.
40380 FUTURE DIRECTIONS
40381
             None.
40382 SEE ALSO
40383
             get
40384 CHANGE HISTORY
```

First released in Issue 2.

who Utilities

40386 NAME 40387	who — disp	olay who is on the system				
40388 SYNOI 40389 UP		SIS who [-mTu]				
40390 XSI	who [-mu]	who [-mu]-s[-bHlprt][file]				
40391	who [-mTu][-abdHlprt][file]				
40392	who -q [f	who -q [file]				
40393	who am i					
40394 40395	who am I					
40396 DESCR 40397 40398	The who uti	lity shall list various pieces of information about accessible users. The domain of is implementation-defined.				
40399 XSI 40400 40401		e options given, <i>who</i> can also list the user's name, terminal line, login time, elapsed ctivity occurred on the line, and the process ID of the command interpreter for each em user.				
40402 OPTIO 40403 40404	The <i>who</i> utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.					
40405 40406		The following options shall be supported. The metavariables, such as <i>line</i> >, refer to fields described in the STDOUT section.				
40407 XSI 40408	–a	Process the implementation-defined database or named file with the $-\mathbf{b}$, $-\mathbf{d}$, $-\mathbf{l}$, $-\mathbf{p}$, $-\mathbf{r}$, $-\mathbf{t}$, $-\mathbf{T}$ and $-\mathbf{u}$ options turned on.				
40409 XSI	-b	Write the time and date of the last reboot.				
40410 XSI 40411 40412 40413	−d	Write a list of all processes that have expired and not been respawned by the <i>init</i> system process. The <i><exit></exit></i> field shall appear for dead processes and contain the termination and exit values of the dead process. This can be useful in determining why a process terminated.				
40414 XSI	– H	Write column headings above the regular output.				
40415 XSI 40416 40417	- 1	-l (The letter ell.) List only those lines on which the system is waiting for someone to login. The <i><name></name></i> field shall be LOGIN in such cases. Other fields shall be the same as for user entries except that the <i><state></state></i> field does not exist.				
40418	-m	-m Output only information about the current terminal.				
40419 XSI 40420	- p	-p List any other process that is currently active and has been previously spawned by <i>init</i> .				
40421 XSI 40422	-q	(Quick.) List only the names and the number of users currently logged on. When this option is used, all other options shall be ignored.				
40423 XSI	–r	Write the current <i>run-level</i> of the <i>init</i> process.				
40424 XSI	-s	List only the < <i>name</i> >, < <i>line</i> >, and < <i>time</i> > fields. This is the default case.				
40425 XSI	–t	-t Indicate the last change to the system clock.				

who **Utilities**

40426	-T	Show the state of each terminal, as described in the STDOUT section.
40427 40428 40429 XSI 40430 40431 40432 40433	-u	Write "idle time" for each displayed user in addition to any other information. The idle time is the time since any activity occurred on the user's terminal. The method of determining this is unspecified. This option shall list only those users who are currently logged in. The <name> is the user's login name. The line> is the name of the line as found in the directory /dev. The <time> is the time that the user logged in. The <activity> is the number of hours and minutes since activity last occurred on that particular line. A dot indicates that the terminal has seen activity in the last</activity></time></name>
40434 40435 40436 40437		minute and is therefore "current". If more than twenty-four hours have elapsed or the line has not been used since boot time, the entry shall be marked $< old >$. This field is useful when trying to determine whether a person is working at the terminal or not. The $< pid >$ is the process ID of the user's login process.
40438 OPERA 40439 XSI		g operands shall be supported:
40440 40441	am i, am I	In the POSIX locale, limit the output to describing the invoking user, equivalent to the $-\mathbf{m}$ option. The \mathbf{am} and \mathbf{i} or \mathbf{I} must be separate arguments.

40439 XSI	The following operands shall be supported:		
40440 40441	am i, am I	In the POSIX locale, limit the output to describing the invoking user, equivalent to the $-m$ option. The am and i or I must be separate arguments.	
40442 40443	file	Specify a pathname of a file to substitute for the implementation-defined database of logged-on users that <i>who</i> uses by default.	

40444 **STDIN**

Not used. 40445

40446 INPUT FILES

None. 40447

40448 ENVIRONMENT VARIABLES					
40449	The following environment variables shall affect the execution of who:				
40450 40451 40452 40453	LANG	Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.)			
40454 40455	LC_ALL	If set to a non-empty string value, override the values of all the other internationalization variables.			
40456 40457 40458	LC_CTYPE	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).			
40459 40460 40461	LC_MESSA	GES Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.			
40462	LC_TIME	Determine the locale used for the format and contents of the date and time strings.			
40463 XSI	NLSPATH	Determine the location of message catalogs for the processing of <i>LC_MESSAGES</i> .			
40464 40465	TZ	Determine the timezone used when writing date and time information. If <i>TZ</i> is unset or null, an unspecified default timezone shall be used.			

40466 ASYNCHRONOUS EVENTS

40467 Default. **who** Utilities

40468 STDOUT 40469 The who utility shall write its default format to the standard output in an implementation-40470 defined format, subject only to the requirement of containing the information described above. 40471 XSI OF XSI-conformant systems shall write the default information to the standard output in the 40472 following general format: <name>[<state>]<line><time>[<activity>][<pid>][<comment>][<exit>] 40473 40474 The following format shall be used for the **–T** option: 40475 "%s %c %s %sn" <name>, <terminal state>, <terminal name>, 40476 <time of login> where < terminal state> is one of the following characters: 40477 40478 The terminal allows write access to other users. The terminal denies write access to other users. 40479 2 The terminal write-access state cannot be determined. 40480 40481 In the POSIX locale, the *<time of login>* shall be equivalent in format to the output of: date +"%b %e %H:%M" 40482 If the -u option is used with -T, the idle time shall be added to the end of the previous format in 40483 an unspecified format. 40484 40485 STDERR The standard error shall be used only for diagnostic messages. 40487 OUTPUT FILES None. 40488 40489 EXTENDED DESCRIPTION None. 40490 40491 EXIT STATUS 40492 The following exit values shall be returned: Successful completion. 40493 >0 An error occurred. 40494 40495 CONSEQUENCES OF ERRORS Default. 40496 40497 APPLICATION USAGE 40498 The name init used for the system process is the most commonly used on historical systems, but 40499 The "domain of accessibility" referred to is a broad concept that permits interpretation either on 40500 40501 a very secure basis or even to allow a network-wide implementation like the historical rwho. 40502 EXAMPLES None 40503 40504 RATIONALE 40505 Due to differences between historical implementations, the base options provided were a 40506 compromise to allow users to work with those functions. The standard developers also 40507 considered removing all the options, but felt that these options offered users valuable

40508

functionality. Additional options to match historical systems are available on XSI-conformant

Utilities who

40509	systems.
40510 40511 40512	It is recognized that the <i>who</i> command may be of limited usefulness, especially in a multi-level secure environment. The standard developers considered, however, that having some standard method of determining the "accessibility" of other users would aid user portability.
40513 40514 40515	No format was specified for the default <i>who</i> output for systems not supporting the XSI Extension. In such a user-oriented command, designed only for human use, this was not considered to be a deficiency.
40516 40517	The format of the terminal name is unspecified, but the descriptions of <i>ps, talk</i> , and <i>write</i> require that they use the same format.
40518	It is acceptable for an implementation to produce no output for an invocation of who mil.
40519 FUTUR 40520	RE DIRECTIONS None.
40521 SEE AI	SO
40522	mesg
40523 CHAN 40524	GE HISTORY First released in Issue 2.
40525 Issue 6	
40526	This utility is marked as part of the User Portability Utilities option.
40527	The TZ entry is added to the ENVIRONMENT VARIABLES section.

write Utilities

```
40528 NAME
40529 write — write to another user
40530 SYNOPSIS
```

40531 UP write user name [terminal]

40532

40540

40542

40543

40544

40545

40546

40547

40548

40549

40550

40551

40552

40553 40554

40555

40556

40557 40558

40559

40560

40561

40562

40563

40564

40565

40566

40533 **DESCRIPTION**

The *write* utility shall read lines from the user's standard input and write them to the terminal of another user. When first invoked, it shall write the message:

40536 Message from sender-login-id (sending-terminal) [date]...

to *user_name*. When it has successfully completed the connection, the sender's terminal shall be alerted twice to indicate that what the sender is typing is being written to the recipient's terminal.

If the recipient wants to reply, this can be accomplished by typing:

40541 write sender-login-id [sending-terminal]

upon receipt of the initial message. Whenever a line of input as delimited by an NL, EOF, or EOL special character (see the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface) is accumulated while in canonical input mode, the accumulated data shall be written on the other user's terminal. Characters shall be processed as follows:

- Typing <alert> shall write the alert character to the recipient's terminal.
- Typing the erase and kill characters shall affect the sender's terminal in the manner described by the **termios** interface in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface.
- Typing the interrupt or end-of-file characters shall cause *write* to write an appropriate message ("EOT\n" in the POSIX locale) to the recipient's terminal and exit.
- Typing characters from *LC_CTYPE* classifications **print** or **space** shall cause those characters to be sent to the recipient's terminal.
- When and only when the stty iexten local mode is enabled, the existence and processing of additional special control characters and multi-byte or single-byte functions is implementation-defined.
- Typing other non-printable characters shall cause implementation-defined sequences of printable characters to be written to the recipient's terminal.

To write to a user who is logged in more than once, the *terminal* argument can be used to indicate which terminal to write to; otherwise, the recipient's terminal is selected in an implementation-defined manner and an informational message is written to the sender's standard output, indicating which terminal was chosen.

Permission to be a recipient of a *write* message can be denied or granted by use of the *mesg* utility. However, a user's privilege may further constrain the domain of accessibility of other users' terminals. The *write* utility shall fail when the user lacks the appropriate privileges to perform the requested action.

40567 **OPTIONS**

40568 None.

Utilities write

40569 OPERANDS 40570 The following operands shall be supported: 40571 user name Login name of the person to whom the message shall be written. The application shall ensure that this operand is of the form returned by the *who* utility. 40572 40573 terminal Terminal identification in the same format provided by the *who* utility. 40574 **STDIN** 40575 Lines to be copied to the recipient's terminal are read from standard input. 40576 INPUT FILES 40577 None. 40578 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of write: 40579 LANG Provide a default value for the internationalization variables that are unset or null. 40580 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 40581 40582 Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 40583 LC_ALL If set to a non-empty string value, override the values of all the other 40584 internationalization variables. 40585 Determine the locale for the interpretation of sequences of bytes of text data as 40586 LC_CTYPE characters (for example, single-byte as opposed to multi-byte characters in 40587 arguments and input files). If the recipient's locale does not use an LC_CTYPE 40588 equivalent to the sender's, the results are undefined. 40589 LC MESSAGES 40590 Determine the locale that should be used to affect the format and contents of 40591 40592 diagnostic messages written to standard error and informative messages written to standard output. 40593 40594 XSI **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*. 40595 ASYNCHRONOUS EVENTS 40596 If an interrupt signal is received, write shall write an appropriate message on the recipient's 40597 terminal and exit with a status of zero. It shall take the standard action for all other signals. 40598 STDOUT An informational message shall be written to standard output if a recipient is logged in more 40599 than once. 40600 40601 STDERR 40602 The standard error shall be used only for diagnostic messages. 40603 OUTPUT FILES The recipient's terminal is used for output. 40604 40605 EXTENDED DESCRIPTION None. 40606 40607 EXIT STATUS The following exit values shall be returned: 40608 Successful completion. 40609

>0 The addressed user is not logged on or the addressed user denies permission.

write **Utilities**

40611 CONSEQUENCES OF ERRORS

40612 Default.

40613 APPLICATION USAGE

40614 The *talk* utility is considered by some users to be a more usable utility on full-screen terminals.

40615 EXAMPLES

40616 None.

40617 RATIONALE

40618 The write utility was included in this volume of IEEE Std 1003.1-2001 since it can be implemented on all terminal types. The standard developers considered the talk utility, which 40619 cannot be implemented on certain terminals, to be a "better" communications interface. Both of 40620 40621 these programs are in widespread use on historical implementations. Therefore, the standard developers decided that both utilities should be specified. 40622

40623 The format of the terminal name is unspecified, but the descriptions of ps, talk, who, and write 40624 require that they all use or accept the same format.

40625 FUTURE DIRECTIONS

None. 40626

40627 SEE ALSO

40628 mesg, talk, who, the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 11, General Terminal Interface 40629

40630 CHANGE HISTORY

First released in Issue 2. 40631

40632 Issue 5

The FUTURE DIRECTIONS section is added. 40633

40634 Issue 6

40635 This utility is marked as part of the User Portability Utilities option.

40636 The normative text is reworded to avoid use of the term "must" for application requirements. Utilities xargs

```
40637 NAME
40638 xargs — construct argument lists and invoke utility
40639 SYNOPSIS
40640 XSI xargs [-t] [-p]] [-E eofstr] [-I replstr] [-L number] [-n number [-x]]
40641 [-s size] [utility [argument...]]
```

DESCRIPTION

The *xargs* utility shall construct a command line consisting of the *utility* and *argument* operands specified followed by as many arguments read in sequence from standard input as fit in length and number constraints specified by the options. The *xargs* utility shall then invoke the constructed command line and wait for its completion. This sequence shall be repeated until one of the following occurs:

- An end-of-file condition is detected on standard input.
- The logical end-of-file string (see the **–E** *eofstr* option) is found on standard input after double-quote processing, apostrophe processing, and backslash escape processing (see next paragraph).
- An invocation of a constructed command line returns an exit status of 255.

The application shall ensure that arguments in the standard input are separated by unquoted

 characters and non-<newline>s can be quoted by enclosing them in double-quotes. A string of zero or more non-apostrophe (''') characters and non-<newline>s can be quoted by enclosing them in apostrophes. Any unquoted character can be escaped by preceding it with a backslash. The utility named by *utility* shall be executed one or more times until the end-of-file is reached or the logical end-of file string is found. The results are unspecified if the utility named by *utility* attempts to read from its standard input.

The generated command line length shall be the sum of the size in bytes of the utility name and each argument treated as strings, including a null byte terminator for each of these strings. The *xargs* utility shall limit the command line length such that when the command line is invoked, the combined argument and environment lists (see the *exec* family of functions in the System Interfaces volume of IEEE Std 1003.1-2001) shall not exceed {ARG_MAX}-2 048 bytes. Within this constraint, if neither the -**n** nor the -**s** option is specified, the default command line length shall be at least {LINE_MAX}.

40668 OPTIONS

The *xargs* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported:

		~ .
40672 40673 40674 40675 40676	–E eofstr	Use <i>eofstr</i> as the logical end-of-file string. If $-E$ is not specified, it is unspecified whether the logical end-of-file string is the underscore character ($'_'$) or the end-of-file string capability is disabled. When <i>eofstr</i> is the null string, the logical end-of-file string capability shall be disabled and underscore characters shall be taken literally.
40677 XSI	–I replstr	Insert mode: <i>utility</i> is executed for each line from standard input, taking the entire
40678	•	line as a single argument, inserting it in arguments for each occurrence of replstr. A
40679		maximum of five arguments in arguments can each contain one or more instances
40680		of replstr. Any <blank>s at the beginning of each line shall be ignored.</blank>
40681		Constructed arguments cannot grow larger than 255 bytes. Option -x shall be
40682		forced on.

xargs Utilities

40683 XSI 40684 40685 40686 40687 40688	–L number	The <i>utility</i> shall be executed for each non-empty <i>number</i> lines of arguments from standard input. The last invocation of <i>utility</i> shall be with fewer lines of arguments if fewer than <i>number</i> remain. A line is considered to end with the first <newline> unless the last character of the line is a </newline>			
40689 40690	– n number	Invoke <i>utility</i> using as many standard input arguments as possible, up to <i>number</i> (a positive decimal integer) arguments maximum. Fewer arguments shall be used if:			
40691 40692		• The command line length accumulated exceeds the size specified by the $-s$ option (or {LINE_MAX} if there is no $-s$ option).			
40693		• The last iteration has fewer than <i>number</i> , but not zero, operands remaining.			
40694 40695 40696 40697 40698	- p	Prompt mode: the user is asked whether to execute <i>utility</i> at each invocation. Trace mode (-t) is turned on to write the command instance to be executed, followed by a prompt to standard error. An affirmative response read from /dev/tty shall execute the command; otherwise, that particular invocation of <i>utility</i> shall be skipped.			
40699 40700 40701	−s size	Invoke <i>utility</i> using as many standard input arguments as possible yielding a command line length less than <i>size</i> (a positive decimal integer) bytes. Fewer arguments shall be used if:			
40702		• The total number of arguments exceeds that specified by the $-\mathbf{n}$ option.			
40703 XSI		 The total number of lines exceeds that specified by the –L option. 			
40704		End-of-file is encountered on standard input before size bytes are accumulated.			
40705 40706 40707 40708 40709		Values of <i>size</i> up to at least {LINE_MAX} bytes shall be supported, provided that the constraints specified in the DESCRIPTION are met. It shall not be considered an error if a value larger than that supported by the implementation or exceeding the constraints specified in the DESCRIPTION is given; <i>xargs</i> shall use the largest value it supports within the constraints.			
40710 40711	-t	Enable trace mode. Each generated command line shall be written to standard error just prior to invocation.			
40712 40713 XSI 40714	- x	Terminate if a command line containing <i>number</i> arguments (see the -n option above) or <i>number</i> lines (see the -L option above) will not fit in the implied of specified size (see the -s option above).			
40715 OPERA					
40716	The following operands shall be supported:				
40717 40718 40719 40720 40721	utility	The name of the utility to be invoked, found by search path using the <i>PATH</i> environment variable, described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables. If <i>utility</i> is omitted, the default shall be the <i>echo</i> utility. If the <i>utility</i> operand names any of the special built-in utilities in Section 2.14 (on page 64), the results are undefined.			
40722	argument	An initial option or operand for the invocation of <i>utility</i> .			
40723 STDIN 40724 40725	The standard input shall be a text file. The results are unspecified if an end-of-file condition is detected immediately following an escaped <newline>.</newline>				

Utilities xargs

40726 INPUT FILES The file $\frac{\mathbf{dev}}{\mathbf{tty}}$ shall be used to read responses required by the $-\mathbf{p}$ option. 40727 **40728 ENVIRONMENT VARIABLES** The following environment variables shall affect the execution of *xargs*: 40729 40730 LANG Provide a default value for the internationalization variables that are unset or null. (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 40731 Internationalization Variables for the precedence of internationalization variables 40732 used to determine the values of locale categories.) 40733 LC_ALL 40734 If set to a non-empty string value, override the values of all the other internationalization variables. 40735 40736 LC_COLLATE Determine the locale for the behavior of ranges, equivalence classes, and multi-40737 character collating elements used in the extended regular expression defined for 40738 40739 the **yesexpr** locale keyword in the *LC_MESSAGES* category. LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 40740 40741 characters (for example, single-byte as opposed to multi-byte characters in arguments and input files) and the behavior of character classes used in the 40742 extended regular expression defined for the yesexpr locale keyword in the 40743 40744 *LC_MESSAGES* category. LC_MESSAGES 40745 40746 Determine the locale for the processing of affirmative responses and that should be used to affect the format and contents of diagnostic messages written to standard 40747 40748 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 40749 XSI **PATH** 40750 Determine the location of *utility*, as described in the Base Definitions volume of IEEE Std 1003.1-2001, Chapter 8, Environment Variables. 40751 40752 ASYNCHRONOUS EVENTS Default. 40753 40754 STDOUT Not used. 40755 **40756 STDERR** The standard error shall be used for diagnostic messages and the -t and -p options. If the -t 40757 40758 option is specified, the *utility* and its constructed argument list shall be written to standard error, as it will be invoked, prior to invocation. If $-\mathbf{p}$ is specified, a prompt of the following format 40759 40760 shall be written (in the POSIX locale): 40761 40762 at the end of the line of the output from **–t**. 40763 OUTPUT FILES 40764 None.

40765 EXTENDED DESCRIPTION

None.

xargs Utilities

40767 EXIT STATUS

40768 The following exit values shall be returned:

- 40769 0 All invocations of *utility* returned exit status zero.
- 40770 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.
- 40773 126 The utility specified by *utility* was found but could not be invoked.
- 40774 127 The utility specified by *utility* could not be found.

40775 CONSEQUENCES OF ERRORS

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.

40780 APPLICATION USAGE

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.

Note that input is parsed as lines; <blank>s separate arguments. If xargs is used to bundle output of commands like find dir -print or ls into commands to be executed, unexpected results are likely if any filenames contain any <blank>s or <newline>s. This can be fixed by using find to call a script that converts each file found into a quoted string that is then piped to xargs. 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 and the shell syntax is not fully compatible with it. An easy rule that can be used to transform any string into a quoted form that xargs interprets correctly is to precede each character in the string with a backslash.

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.

The *command, env, nice, nohup, time,* 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.

40806 EXAMPLES

1. The following command combines the output of the parenthesised commands onto one line, which is then written to the end-of-file **log**:

```
(logname; date; printf "%s\n" "$0 $*") | xargs >>log
```

2. The following command invokes *diff* with successive pairs of arguments originally typed as command line arguments (assuming there are no embedded <blank>s in the elements of the original argument list):

Utilities xargs

```
40813 printf "%s\n" "$*" | xargs -n 2 -x diff
```

3. In the following commands, the user is asked which files in the current directory are to be archived. The files are archived into **arch**; *a*, one at a time, or *b*, many at a time.

```
a. ls | xargs -p -L 1 ar -r arch
b. ls | xargs -p -L 1 | xargs ar -r arch
```

4. The following executes with successive pairs of arguments originally typed as command line arguments:

```
echo $* | xargs -n 2 diff
```

5. On XSI-conformant systems, the following moves all files from directory \$1 to directory \$2, and echoes each move command just before doing it:

```
ls $1 | xargs -I \{\} -t mv \{1/\{\}\} \{2/\{\}\}
```

40824 RATIONALE

40837

 The *xargs* utility was usually found only in System V-based systems; BSD systems included an *apply* utility that provided functionality similar to *xargs* –**n** *number*. The SVID lists *xargs* as a software development extension. This volume of IEEE Std 1003.1-2001 does not share the view that it is used only for development, and therefore it is not optional.

The classic application of the *xargs* utility is in conjunction with the *find* utility to reduce the number of processes launched by a simplistic use of the *find* –**exec** combination. The *xargs* utility is also used to enforce an upper limit on memory required to launch a process. With this basis in mind, this volume of IEEE Std 1003.1-2001 selected only the minimal features required.

Although the 255 exit status is mostly an accident of historical implementations, it allows a utility being used by *xargs* to tell *xargs* to terminate if it knows no further invocations using the current data stream shall succeed. Any non-zero exit status from a utility falls into the 1-125 range when *xargs* exits. There is no statement of how the various non-zero utility exit status codes are accumulated by *xargs*. The value could be the addition of all codes, their highest value, the last one received, or a single value such as 1. Since no algorithm is arguably better than the others, and since many of the standard utilities say little more (portably) than "pass/fail", no new algorithm was invented.

Several other *xargs* options were withdrawn because simple alternatives already exist within this volume of IEEE Std 1003.1-2001. For example, the -**i** *replstr* option can be just as efficiently performed using a shell **for** loop. Since *xargs* calls an *exec* function with each input line, the -**i** option does not usually exploit the grouping capabilities of *xargs*.

The requirement that *xargs* never produces command lines such that invocation of *utility* is within 2 048 bytes of hitting the POSIX *exec* {ARG_MAX} limitations is intended to guarantee that the invoked utility has room to modify its environment variables and command line arguments and still be able to invoke another utility. Note that the minimum {ARG_MAX} allowed by the System Interfaces volume of IEEE Std 1003.1-2001 is 4 096 bytes and the minimum value allowed by this volume of IEEE Std 1003.1-2001 is 2 048 bytes; therefore, the 2 048 bytes difference seems reasonable. Note, however, that *xargs* may never be able to invoke a utility if the environment passed in to *xargs* comes close to using {ARG_MAX} bytes.

The version of *xargs* required by this volume of IEEE Std 1003.1-2001 is required to wait for the completion of the invoked command before invoking another command. This was done because historical scripts using *xargs* assumed sequential execution. Implementations wanting to provide parallel operation of the invoked utilities are encouraged to add an option enabling parallel invocation, but should still wait for termination of all of the children before *xargs* terminates normally.

xargs Utilities

The —e option was omitted from the ISO POSIX-2:1993 standard in the belief that the *eofstr* option-argument was recognized only when it was on a line by itself and before quote and escape processing were performed, and that the logical end-of-file processing was only enabled if a —e option was specified. In that case, a simple *sed* script could be used to duplicate the —e functionality. Further investigation revealed that:

- The logical end-of-file string was checked for after quote and escape processing, making a *sed* script that provided equivalent functionality much more difficult to write.
- The default was to perform logical end-of-file processing with an underscore as the logical end-of-file string.

To correct this misunderstanding, the —**E** *eofstr* option was adopted from the X/Open Portability Guide. Users should note that the description of the —**E** option matches historical documentation of the —**e** option (which was not adopted because it did not support the Utility Syntax Guidelines), by saying that if *eofstr* is the null string, logical end-of-file processing is disabled. Historical implementations of *xargs* actually did not disable logical end-of-file processing; they treated a null argument found in the input as a logical end-of-file string. (A null *string* argument could be generated using single or double quotes (' ' or " "). Since this behavior was not documented historically, it is considered to be a bug.

40876 FUTURE DIRECTIONS

40877 None.

40878 SEE ALSO

40859

40860

40861

40862 40863

40864

40865

40866 40867

40868

40869

40870

40871

40872

40873

40874 40875

40879 Chapter 2 (on page 29), echo, find, the System Interfaces volume of IEEE Std 1003.1-2001, exec

40880 CHANGE HISTORY

40881 First released in Issue 2.

40882 Issue 5

40883 A second FUTURE DIRECTION is added.

40884 Issue 6

40888

40891

The obsolescent $-\mathbf{e}$, $-\mathbf{i}$, and $-\mathbf{l}$ options are removed.

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

- The -p option is added.
- In the INPUT FILES section, the file /dev/tty is used to read responses required by the $-\mathbf{p}$ option.
 - The STDERR section is updated to describe the -p option.
- 40892 The description of the –E option is aligned with the ISO POSIX-2: 1993 standard.
- 40893 The normative text is reworded to avoid use of the term "must" for application requirements.

Utilities yacc

40894 **NAME**

40895 yacc — yet another compiler (**DEVELOPMENT**)

40896 SYNOPSIS

40897 CD yacc [-dltv] [-b file_prefix] [-p sym_prefix] grammar
40898

40899 **DESCRIPTION**

The *yacc* utility shall read a description of a context-free grammar in *grammar* and write C source code, conforming to the ISO C standard, to a code file, and optionally header information into a header file, in the current directory. The C code shall define a function and related routines and macros for an automaton that executes a parsing algorithm meeting the requirements in Algorithms (on page 1074).

40905 The form and meaning of the grammar are described in the EXTENDED DESCRIPTION section.

The C source code and header file shall be produced in a form suitable as input for the C compiler (see *c99*).

40908 OPTIONS

40911

40912 40913

40914

40915

40924

40925

40926

40927

40928

40929

40930

40931

40932

40933

The *yacc* utility shall conform to the Base Definitions volume of IEEE Std 1003.1-2001, Section 12.2, Utility Syntax Guidelines.

The following options shall be supported:

- -b file_prefix Use file_prefix instead of y as the prefix for all output filenames. The code file y.tab.c, the header file y.tab.h (created when -d is specified), and the description file y.output (created when -v is specified), shall be changed to file_prefix.tab.c, file_prefix.tab.h, and file_prefix.output, respectively.
- 40916 —**d** Write the header file; by default only the code file is written. The **#define**40917 statements associate the token codes assigned by *yacc* with the user-declared token
 40918 names. This allows source files other than **y.tab.c** to access the token codes.
- 40919 —I Produce a code file that does not contain any **#line** constructs. If this option is not present, it is unspecified whether the code file or header file contains **#line** directives. This should only be used after the grammar and the associated actions are fully debugged.

40923 —**p** *sym_prefix*

Use sym_prefix instead of yy as the prefix for all external names produced by yacc. The names affected shall include the functions yyparse(), yylex(), and yyerror(), and the variables yylval, yychar, and yydebug. (In the remainder of this section, the six symbols cited are referenced using their default names only as a notational convenience.) Local names may also be affected by the $-\mathbf{p}$ option; however, the $-\mathbf{p}$ option shall not affect #define symbols generated by yacc.

- Modify conditional compilation directives to permit compilation of debugging code in the code file. Runtime debugging statements shall always be contained in the code file, but by default conditional compilation directives prevent their compilation.
- 40934 **v** Write a file containing a description of the parser and a report of conflicts generated by ambiguities in the grammar.

yacc Utilities

40936 OPERANDS 40937 The following operand is required: 40938 grammar A pathname of a file containing instructions, hereafter called *grammar*, for which a parser is to be created. The format for the grammar is described in the EXTENDED 40939 DESCRIPTION section. 40940 40941 **STDIN** Not used. 40942 40943 INPUT FILES 40944 The file grammar shall be a text file formatted as specified in the EXTENDED DESCRIPTION 40945 section. 40946 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of *yacc*: 40947 LANG Provide a default value for the internationalization variables that are unset or null. 40948 (See the Base Definitions volume of IEEE Std 1003.1-2001. Section 8.2. 40949 Internationalization Variables for the precedence of internationalization variables 40950 40951 used to determine the values of locale categories.) LC ALL If set to a non-empty string value, override the values of all the other 40952 internationalization variables. 40953 LC_CTYPE Determine the locale for the interpretation of sequences of bytes of text data as 40954 40955 characters (for example, single-byte as opposed to multi-byte characters in arguments and input files). 40956 LC MESSAGES 40957 Determine the locale that should be used to affect the format and contents of 40958 diagnostic messages written to standard error. 40959 NLSPATH Determine the location of message catalogs for the processing of *LC_MESSAGES*. 40960 XSI 40961 The LANG and LC * variables affect the execution of the yacc utility as stated. The main() function defined in **Yacc Library** (on page 1074) shall call: 40962 setlocale(LC ALL, "") 40963 and thus the program generated by yacc shall also be affected by the contents of these variables 40964 at runtime. 40965 **40966 ASYNCHRONOUS EVENTS** 40967 Default. **40968 STDOUT** Not used. 40969 **40970 STDERR** If shift/reduce or reduce/reduce conflicts are detected in grammar, yacc shall write a report of 40971 those conflicts to the standard error in an unspecified format. 40972 40973 Standard error shall also be used for diagnostic messages.

40974 OUTPUT FILES

The code file, the header file, and the description file shall be text files. All are described in the following sections.

Utilities yacc

40977 Code File

This file shall contain the C source code for the *yyparse()* function. It shall contain code for the various semantic actions with macro substitution performed on them as described in the EXTENDED DESCRIPTION section. It also shall contain a copy of the **#define** statements in the header file. If a **%union** declaration is used, the declaration for YYSTYPE shall also be included in this file.

40983 Header File

The header file shall contain **#define** statements that associate the token numbers with the token names. This allows source files other than the code file to access the token codes. If a **%union** declaration is used, the declaration for YYSTYPE and an *extern YYSTYPE yylval* declaration shall also be included in this file.

Description File

The description file shall be a text file containing a description of the state machine corresponding to the parser, using an unspecified format. Limits for internal tables (see **Limits** (on page 1075)) shall also be reported, in an implementation-defined manner. (Some implementations may use dynamic allocation techniques and have no specific limit values to report.)

40994 EXTENDED DESCRIPTION

The *yacc* command accepts a language that is used to define a grammar for a target language to be parsed by the tables and code generated by *yacc*. The language accepted by *yacc* as a grammar for the target language is described below using the *yacc* input language itself.

The input *grammar* includes rules describing the input structure of the target language and code to be invoked when these rules are recognized to provide the associated semantic action. The code to be executed shall appear as bodies of text that are intended to be C-language code. The C-language inclusions are presumed to form a correct function when processed by *yacc* into its output files. The code included in this way shall be executed during the recognition of the target language.

Given a grammar, the *yacc* utility generates the files described in the OUTPUT FILES section. The code file can be compiled and linked using c99. If the declaration and programs sections of the grammar file did not include definitions of main(), yylex(), and yyerror(), the compiled output requires linking with externally supplied versions of those functions. Default versions of main() and yyerror() are supplied in the yacc library and can be linked in by using the -ly operand to c99. The yacc library interfaces need not support interfaces with other than the default yy symbol prefix. The application provides the lexical analyzer function, yylex(); the lex utility is specifically designed to generate such a routine.

Input Language

The application shall ensure that every specification file consists of three sections in order:

declarations, grammar rules, and programs, separated by double percent signs ("%%"). The
declarations and programs sections can be empty. If the latter is empty, the preceding "%%"
mark separating it from the rules section can be omitted.

41017 The input is free form text following the structure of the grammar defined below.

yacc Utilities

Lexical Structure of the Grammar

The <black>s, <newline>s, and <form-feed>s shall be ignored, except that the application shall ensure that they do not appear in names or multi-character reserved symbols. Comments shall be enclosed in "/* . . . */", and can appear wherever a name is valid.

Names are of arbitrary length, made up of letters, periods ('.'), underscores ('_'), and non-initial digits. Uppercase and lowercase letters are distinct. Conforming applications shall not use names beginning in **yy** or **YY** since the *yacc* parser uses such names. Many of the names appear in the final output of *yacc*, and thus they should be chosen to conform with any additional rules created by the C compiler to be used. In particular they appear in **#define** statements.

A literal shall consist of a single character enclosed in single-quotes ('''). All of the escape sequences supported for character constants by the ISO C standard shall be supported by *yacc*.

The relationship with the lexical analyzer is discussed in detail below.

The application shall ensure that the NUL character is not used in grammar rules or literals.

Declarations Section

The declarations section is used to define the symbols used to define the target language and their relationship with each other. In particular, much of the additional information required to resolve ambiguities in the context-free grammar for the target language is provided here.

Usually *yacc* assigns the relationship between the symbolic names it generates and their underlying numeric value. The declarations section makes it possible to control the assignment of these values.

It is also possible to keep semantic information associated with the tokens currently on the parse stack in a user-defined C-language **union**, if the members of the union are associated with the various names in the grammar. The declarations section provides for this as well.

The first group of declarators below all take a list of names as arguments. That list can optionally be preceded by the name of a C union member (called a *tag* below) appearing within '<' and '>'. (As an exception to the typographical conventions of the rest of this volume of IEEE Std 1003.1-2001, in this case <*tag*> does not represent a metavariable, but the literal angle bracket characters surrounding a symbol.) The use of *tag* specifies that the tokens named on this line shall be of the same C type as the union member referenced by *tag*. This is discussed in more detail below.

For lists used to define tokens, the first appearance of a given token can be followed by a positive integer (as a string of decimal digits). If this is done, the underlying value assigned to it for lexical purposes shall be taken to be that number.

The following declares *name* to be a token:

```
%token [<tag>] name [number][name [number]]...
```

If *tag* is present, the C type for all tokens on this line shall be declared to be the type referenced by *tag*. If a positive integer, *number*, follows a *name*, that value shall be assigned to the token.

The following declares *name* to be a token, and assigns precedence to it:

```
%left [<tag>] name [number] [name [number]]...
%right [<tag>] name [number] [name [number]]...
```

One or more lines, each beginning with one of these symbols, can appear in this section. All tokens on the same line have the same precedence level and associativity; the lines are in order

Utilities yacc

41061 of increasing precedence or binding strength. %left denotes that the operators on that line are 41062 left associative, and %right similarly denotes right associative operators. If tag is present, it shall 41063 declare a C type for *names* as described for **%token**. The following declares *name* to be a token, and indicates that this cannot be used associatively: 41064 %nonassoc [<tag>] name [number][name [number]]... 41065 If the parser encounters associative use of this token it reports an error. If tag is present, it shall 41066 41067 declare a C type for *names* as described for **%token**. 41068 The following declares that union member *names* are non-terminals, and thus it is required to 41069 have a *tag* field at its beginning: 41070 %type <tag> name... Because it deals with non-terminals only, assigning a token number or using a literal is also 41071 prohibited. If this construct is present, yacc shall perform type checking; if this construct is not 41072 present, the parse stack shall hold only the **int** type. 41073 Every name used in grammar not defined by a %token, %left, %right, or %nonassoc declaration 41074 41075 is assumed to represent a non-terminal symbol. The yacc utility shall report an error for any 41076 non-terminal symbol that does not appear on the left side of at least one grammar rule. Once the type, precedence, or token number of a name is specified, it shall not be changed. If the 41077 41078 first declaration of a token does not assign a token number, yacc shall assign a token number. Once this assignment is made, the token number shall not be changed by explicit assignment. 41079 The following declarators do not follow the previous pattern. 41080 The following declares the non-terminal *name* to be the *start symbol*, which represents the largest, 41081 most general structure described by the grammar rules: 41082 41083 By default, it is the left-hand side of the first grammar rule; this default can be overridden with 41084 41085 this declaration. 41086 The following declares the *yacc* value stack to be a union of the various types of values desired: 41087 %union { body of union (in C) } By default, the values returned by actions (see below) and the lexical analyzer shall be of type 41088 41089

int. The yacc utility keeps track of types, and it shall insert corresponding union member names 41090 in order to perform strict type checking of the resulting parser.

Alternatively, given that at least one <tag> construct is used, the union can be declared in a header file (which shall be included in the declarations section by using a #include construct within %{ and %}), and a **typedef** used to define the symbol YYSTYPE to represent this union. The effect of **%union** is to provide the declaration of YYSTYPE directly from the *yacc* input.

C-language declarations and definitions can appear in the declarations section, enclosed by the 41095 following marks: 41096

41097 왕{ ... 왕}

41091

41092

41093

41094

These statements shall be copied into the code file, and have global scope within it so that they 41098 can be used in the rules and program sections. 41099

The application shall ensure that the declarations section is terminated by the token %%. 41100

Utilities yacc

Grammar Rules in yacc

The rules section defines the context-free grammar to be accepted by the function *yacc* generates, 41102 41103 and associates with those rules C-language actions and additional precedence information. The 41104 grammar is described below, and a formal definition follows.

The rules section is comprised of one or more grammar rules. A grammar rule has the form:

A : BODY ; 41106

41101

41105

41107 41108

41109

41110

41111 41112

41113

41114

41115 41116

41117

41118

41119

41121

41122

41123 41124

41125 41126

41127 41128

41129

41130

41131

41132

41133

41134

41135 41136

41137 41138

41139

41140

41141

41142 41143

41144

41145 41146 The symbol A represents a non-terminal name, and **BODY** represents a sequence of zero or more names, literals, and semantic actions that can then be followed by optional precedence rules. Only the names and literals participate in the formation of the grammar; the semantic actions and precedence rules are used in other ways. The colon and the semicolon are yacc punctuation. If there are several successive grammar rules with the same left-hand side, the vertical bar ' | ' can be used to avoid rewriting the left-hand side; in this case the semicolon appears only after the last rule. The BODY part can be empty (or empty of names and literals) to indicate that the non-terminal symbol matches the empty string.

The yacc utility assigns a unique number to each rule. Rules using the vertical bar notation are distinct rules. The number assigned to the rule appears in the description file.

The elements comprising a BODY are:

Snumber

name, literal These form the rules of the grammar: name is either a token or a non-terminal; literal stands for itself (less the lexically required quotation marks).

semantic action

With each grammar rule, the user can associate actions to be performed each time the rule is recognized in the input process. (Note that the word "action" can also refer to the actions of the parser—shift, reduce, and so on.)

These actions can return values and can obtain the values returned by previous actions. These values are kept in objects of type YYSTYPE (see %union). The result value of the action shall be kept on the parse stack with the left-hand side of the rule, to be accessed by other reductions as part of their right-hand side. By using the <tag> information provided in the declarations section, the code generated by yacc can be strictly type checked and contain arbitrary information. In addition, the lexical analyzer can provide the same kinds of values for tokens, if desired.

An action is an arbitrary C statement and as such can do input or output, call subprograms, and alter external variables. An action is one or more C statements enclosed in curly braces ' { ' and ' } '.

Certain pseudo-variables can be used in the action. These are macros for access to data structures known internally to *yacc*.

\$\$ The value of the action can be set by assigning it to \$\$. If type checking is enabled and the type of the value to be assigned cannot be determined, a diagnostic message may be generated.

> This refers to the value returned by the component specified by the token *number* in the right side of a rule, reading from left to right; number can be zero or negative. If number is zero or negative, it refers to the data associated with the name on the parser's stack preceding the leftmost symbol of the current rule. (That is, "\$0" refers to the name immediately preceding the leftmost name in the current rule to be found on the parser's stack and "\$-1" refers to the symbol to its

41120

Utilities yacc

41147 41148 41149 41150			left.) If <i>number</i> refers to an element past the current point in the rule, or beyond the bottom of the stack, the result is undefined. If type checking is enabled and the type of the value to be assigned cannot be determined, a diagnostic message may be generated.
41151		\$ <tag>numb</tag>	or .
41151		\$<\tag>11u1110	
41152			These correspond exactly to the corresponding symbols without the
41153			tag inclusion, but allow for strict type checking (and preclude
41154			unwanted type conversions). The effect is that the macro is expanded
41155			to use tag to select an element from the YYSTYPE union (using
41156			dataname.tag). This is particularly useful if number is not positive.
41157		\$ <tag>\$</tag>	This imposes on the reference the type of the union member
41158			referenced by tag. This construction is applicable when a reference
41159			to a left context value occurs in the grammar, and provides yacc with
41160			a means for selecting a type.
41161		Actions can	occur anywhere in a rule (not just at the end); an action can access
41162		values retu	rned by actions to its left, and in turn the value it returns can be
41163		accessed by	actions to its right. An action appearing in the middle of a rule shall be
41164		equivalent t	o replacing the action with a new non-terminal symbol and adding an
41165		empty rule	with that non-terminal symbol on the left-hand side. The semantic
41166			tiated with the new rule shall be equivalent to the original action. The
41167		use of actio	ns within rules might introduce conflicts that would not otherwise
41168		exist.	
41169		By default, t	the value of a rule shall be the value of the first element in it. If the first
41170		element do	es not have a type (particularly in the case of a literal) and type
41171		checking is t	turned on by %type , an error message shall result.
41172	precedence	The keywor	d %prec can be used to change the precedence level associated with a
41173		particular g	rammar rule. Examples of this are in cases where a unary and binary
41174		operator ha	ve the same symbolic representation, but need to be given different
41175		precedences	s, or where the handling of an ambiguous if-else construction is
41176		necessary. T	The reserved symbol %prec can appear immediately after the body of
41177			r rule and can be followed by a token name or a literal. It shall cause
41178			nce of the grammar rule to become that of the following token name or
41179		literal. The a	action for the rule as a whole can follow % prec .
41180	If a progran	n section follo	ws, the application shall ensure that the grammar rules are terminated
41181	by %%.		
41182	Programs Se	ection	
41183	The program	s section can	include the definition of the lexical analyzer yylex(), and any other
41184			ose used in the actions specified in the grammar rules. It is unspecified
41185			ection precedes or follows the semantic actions in the output file;
41186			on contains any macro definitions and declarations intended to apply to
41187			actions, it shall place them within " $\{ \ldots \}$ " in the declarations
41188	section.		remaining the second and administration
			

yacc Utilities

```
41189
            Input Grammar
41190
             The following input to yacc yields a parser for the input to yacc. This formal syntax takes
41191
            precedence over the preceding text syntax description.
41192
            The lexical structure is defined less precisely; Lexical Structure of the Grammar (on page 1066)
41193
            defines most terms. The correspondence between the previous terms and the tokens below is as
            follows.
41194
            IDENTIFIER
41195
                             This corresponds to the concept of name, given previously. It also includes
41196
                             literals as defined previously.
41197
            C_IDENTIFIER
                             This is a name, and additionally it is known to be followed by a colon. A literal
41198
                             cannot yield this token.
            NUMBER
                             A string of digits (a non-negative decimal integer).
41199
            TYPE, LEFT, MARK, LCURL, RCURL
41200
                             These correspond directly to %type, %left, %%, %{, and %}.
41201
            {...}
                             This indicates C-language source code, with the possible inclusion of '$'
41202
41203
                             macros as discussed previously.
             /* Grammar for the input to yacc. */
41204
             /* Basic entries. */
41205
41206
             /* The following are recognized by the lexical analyzer. */
41207
             %token
                         IDENTIFIER
                                            /* Includes identifiers and literals */
                                            /* identifier (but not literal)
41208
             %token
                         C IDENTIFIER
41209
                                                followed by a :. */
             %token
                         NUMBER
                                            /* [0-9] [0-9] * */
41210
             /* Reserved words : %type=>TYPE %left=>LEFT, and so on */
41211
41212
                         LEFT RIGHT NONASSOC TOKEN PREC TYPE START UNION
             %token
41213
             %token
                         MARK
                                            /* The %% mark. */
             %token
                         LCURL
                                            /* The %{ mark. */
41214
                                            /* The %} mark. */
41215
             %token
                         RCURL
41216
             /* 8-bit character literals stand for themselves; */
             /* tokens have to be defined for multi-byte characters. */
41217
41218
             %start
                         spec
41219
                    : defs MARK rules tail
41220
             spec
41221
             tail
41222
                    : MARK
41223
41224
                      /* In this action, set up the rest of the file. */
41225
41226
                      /* Empty; the second MARK is optional. */
41227
            defs
41228
                      /* Empty. */
                          defs def
41229
41230
41231
             def
                      START IDENTIFIER
41232
                          UNION
```

Utilities yacc

```
41233
41234
                     /* Copy union definition to output. */
41235
41236
                        LCURL
41237
41238
                     /* Copy C code to output file. */
41239
                     RCURL
41240
                        rword tag nlist
41241
41242
41243
            rword : TOKEN
41244
                   LEFT
41245
                   RIGHT
41246
                   NONASSOC
41247
                    TYPE
41248
                   : /* Empty: union tag ID optional. */
41249
            tag
41250
                    '<' IDENTIFIER '>'
41251
            nlist : nmno
41252
                   nlist nmno
41253
41254
                  : IDENTIFIER
                                         /* Note: literal invalid with % type. */
41255
                   | IDENTIFIER NUMBER /* Note: invalid with % type. */
41256
41257
            /* Rule section */
41258
            rules : C_IDENTIFIER rbody prec
41259
41260
                   | rules rule
41261
41262
            rule : C IDENTIFIER rbody prec
                   | '|' rbody prec
41263
41264
            rbody : /* empty */
41265
                   | rbody IDENTIFIER
41266
41267
                   rbody act
41268
                   : '{'
41269
            act
41270
                       /* Copy action, translate $$, and so on. */
41271
41272
                     '}'
41273
41274
            prec : /* Empty */
41275
                   PREC IDENTIFIER
41276
                   | PREC IDENTIFIER act
41277
41278
                    prec ';'
41279
```

yacc Utilities

Conflicts

The parser produced for an input grammar may contain states in which conflicts occur. The conflicts occur because the grammar is not LALR(1). An ambiguous grammar always contains at least one LALR(1) conflict. The *yacc* utility shall resolve all conflicts, using either default rules or user-specified precedence rules.

Conflicts are either shift/reduce conflicts or reduce/reduce conflicts. A shift/reduce conflict is where, for a given state and lookahead symbol, both a shift action and a reduce action are possible. A reduce/reduce conflict is where, for a given state and lookahead symbol, reductions by two different rules are possible.

The rules below describe how to specify what actions to take when a conflict occurs. Not all shift/reduce conflicts can be successfully resolved this way because the conflict may be due to something other than ambiguity, so incautious use of these facilities can cause the language accepted by the parser to be much different from that which was intended. The description file shall contain sufficient information to understand the cause of the conflict. Where ambiguity is the reason either the default or explicit rules should be adequate to produce a working parser.

The declared precedences and associativities (see **Declarations Section** (on page 1066)) are used to resolve parsing conflicts as follows:

- 1. A precedence and associativity is associated with each grammar rule; it is the precedence and associativity of the last token or literal in the body of the rule. If the **%prec** keyword is used, it overrides this default. Some grammar rules might not have both precedence and associativity.
- 2. If there is a shift/reduce conflict, and both the grammar rule and the input symbol have precedence and associativity associated with them, then the conflict is resolved in favor of the action (shift or reduce) associated with the higher precedence. If the precedences are the same, then the associativity is used; left associative implies reduce, right associative implies shift, and non-associative implies an error in the string being parsed.
- 3. When there is a shift/reduce conflict that cannot be resolved by rule 2, the shift is done. Conflicts resolved this way are counted in the diagnostic output described in **Error Handling**.
- 4. When there is a reduce/reduce conflict, a reduction is done by the grammar rule that occurs earlier in the input sequence. Conflicts resolved this way are counted in the diagnostic output described in **Error Handling**.

Conflicts resolved by precedence or associativity shall not be counted in the shift/reduce and reduce/reduce conflicts reported by *yacc* on either standard error or in the description file.

Error Handling

The token **error** shall be reserved for error handling. The name **error** can be used in grammar rules. It indicates places where the parser can recover from a syntax error. The default value of **error** shall be 256. Its value can be changed using a **%token** declaration. The lexical analyzer should not return the value of **error**.

The parser shall detect a syntax error when it is in a state where the action associated with the lookahead symbol is **error**. A semantic action can cause the parser to initiate error handling by executing the macro YYERROR. When YYERROR is executed, the semantic action passes control back to the parser. YYERROR cannot be used outside of semantic actions.

When the parser detects a syntax error, it normally calls *yyerror()* with the character string "syntax error" as its argument. The call shall not be made if the parser is still recovering

Utilities yacc

from a previous error when the error is detected. The parser is considered to be recovering from a previous error until the parser has shifted over at least three normal input symbols since the last error was detected or a semantic action has executed the macro *yyerrok*. The parser shall not call *yyerror*() when YYERROR is executed.

The macro function YYRECOVERING shall return 1 if a syntax error has been detected and the parser has not yet fully recovered from it. Otherwise, zero shall be returned.

When a syntax error is detected by the parser, the parser shall check if a previous syntax error has been detected. If a previous error was detected, and if no normal input symbols have been shifted since the preceding error was detected, the parser checks if the lookahead symbol is an endmarker (see **Interface to the Lexical Analyzer**). If it is, the parser shall return with a non-zero value. Otherwise, the lookahead symbol shall be discarded and normal parsing shall resume.

When YYERROR is executed or when the parser detects a syntax error and no previous error has been detected, or at least one normal input symbol has been shifted since the previous error was detected, the parser shall pop back one state at a time until the parse stack is empty or the current state allows a shift over **error**. If the parser empties the parse stack, it shall return with a non-zero value. Otherwise, it shall shift over **error** and then resume normal parsing. If the parser reads a lookahead symbol before the error was detected, that symbol shall still be the lookahead symbol when parsing is resumed.

The macro *yyerrok* in a semantic action shall cause the parser to act as if it has fully recovered from any previous errors. The macro *yyclearin* shall cause the parser to discard the current lookahead token. If the current lookahead token has not yet been read, *yyclearin* shall have no effect.

The macro YYACCEPT shall cause the parser to return with the value zero. The macro YYABORT shall cause the parser to return with a non-zero value.

Interface to the Lexical Analyzer

 The *yylex()* function is an integer-valued function that returns a *token number* representing the kind of token read. If there is a value associated with the token returned by *yylex()* (see the discussion of *tag* above), it shall be assigned to the external variable *yylval*.

If the parser and yylex() do not agree on these token numbers, reliable communication between them cannot occur. For (single-byte character) literals, the token is simply the numeric value of the character in the current character set. The numbers for other tokens can either be chosen by yacc, or chosen by the user. In either case, the #define construct of C is used to allow yylex() to return these numbers symbolically. The #define statements are put into the code file, and the header file if that file is requested. The set of characters permitted by yacc in an identifier is larger than that permitted by C. Token names found to contain such characters shall not be included in the #define declarations.

If the token numbers are chosen by *yacc*, the tokens other than literals shall be assigned numbers greater than 256, although no order is implied. A token can be explicitly assigned a number by following its first appearance in the declarations section with a number. Names and literals not defined this way retain their default definition. All token numbers assigned by *yacc* shall be unique and distinct from the token numbers used for literals and user-assigned tokens. If duplicate token numbers cause conflicts in parser generation, *yacc* shall report an error; otherwise, it is unspecified whether the token assignment is accepted or an error is reported.

The end of the input is marked by a special token called the *endmarker*, which has a token number that is zero or negative. (These values are invalid for any other token.) All lexical analyzers shall return zero or negative as a token number upon reaching the end of their input. If

yacc Utilities

the tokens up to, but excluding, the endmarker form a structure that matches the start symbol, the parser shall accept the input. If the endmarker is seen in any other context, it shall be considered an error.

Completing the Program

In addition to *yyparse()* and *yylex()*, the functions *yyerror()* and *main()* are required to make a complete program. The application can supply *main()* and *yyerror()*, or those routines can be obtained from the *yacc* library.

Yacc Library

The following functions shall appear only in the *yacc* library accessible through the –**l y** operand to *c99*; they can therefore be redefined by a conforming application:

int main(void)

This function shall call *yyparse()* and exit with an unspecified value. Other actions within this function are unspecified.

int yyerror(const char *s)

This function shall write the NUL-terminated argument to standard error, followed by a <newline>.

The order of the $-\mathbf{l}$ y and $-\mathbf{l}$ l operands given to c99 is significant; the application shall either provide its own main() function or ensure that $-\mathbf{l}$ y precedes $-\mathbf{l}$ l.

Debugging the Parser

The parser generated by *yacc* shall have diagnostic facilities in it that can be optionally enabled at either compile time or at runtime (if enabled at compile time). The compilation of the runtime debugging code is under the control of YYDEBUG, a preprocessor symbol. If YYDEBUG has a non-zero value, the debugging code shall be included. If its value is zero, the code shall not be included.

In parsers where the debugging code has been included, the external **int** *yydebug* can be used to turn debugging on (with a non-zero value) and off (zero value) at runtime. The initial value of *yydebug* shall be zero.

When –t is specified, the code file shall be built such that, if YYDEBUG is not already defined at compilation time (using the *c99* –**D** YYDEBUG option, for example), YYDEBUG shall be set explicitly to 1. When –t is not specified, the code file shall be built such that, if YYDEBUG is not already defined, it shall be set explicitly to zero.

The format of the debugging output is unspecified but includes at least enough information to determine the shift and reduce actions, and the input symbols. It also provides information about error recovery.

Algorithms

The parser constructed by *yacc* implements an LALR(1) parsing algorithm as documented in the literature. It is unspecified whether the parser is table-driven or direct-coded.

A parser generated by *yacc* shall never request an input symbol from *yylex*() while in a state where the only actions other than the error action are reductions by a single rule.

The literature of parsing theory defines these concepts.

Utilities yacc

Limits 41412

41413 41414

41415

41416

41417

41418

41419

41442

41443

41444

41446

41447

41448

41449

41450

41451

41452

41453

41454

41455

41456 41457

41458

The yacc utility may have several internal tables. The minimum maximums for these tables are shown in the following table. The exact meaning of these values is implementation-defined. The implementation shall define the relationship between these values and between them and any error messages that the implementation may generate should it run out of space for any internal structure. An implementation may combine groups of these resources into a single pool as long as the total available to the user does not fall below the sum of the sizes specified by this section.

Table 4-22 Internal Limits in yacc

41420 41421	Limit	Minimum Maximum	Description
41422	{NTERMS}	126	Number of tokens.
41423	{NNONTERM}	200	Number of non-terminals.
41424	{NPROD}	300	Number of rules.
41425	{NSTATES}	600	Number of states.
41426	{MEMSIZE}	5 200	Length of rules. The total length, in names
41427			(tokens and non-terminals), of all the rules of the
41428			grammar. The left-hand side is counted for each
41429			rule, even if it is not explicitly repeated, as
41430			specified in Grammar Rules in yacc (on page
41431			1068).
41432	{ACTSIZE}	4 000	Number of actions. "Actions" here (and in the
41433			description file) refer to parser actions (shift,
41434			reduce, and so on) not to semantic actions
41435			defined in Grammar Rules in yacc (on page
41436			1068).

41437 EXIT STATUS

The following exit values shall be returned: 41438

- Successful completion. 41439
- >0 An error occurred. 41440

41441 CONSEQUENCES OF ERRORS

If any errors are encountered, the run is aborted and yacc exits with a non-zero status. Partial code files and header files may be produced. The summary information in the description file shall always be produced if the –v flag is present.

41445 APPLICATION USAGE

Historical implementations experience name conflicts on the names yacc.tmp, yacc.acts, yacc.debug, y.tab.c, y.tab.h, and y.output if more than one copy of yacc is running in a single directory at one time. The -b option was added to overcome this problem. The related problem of allowing multiple yacc parsers to be placed in the same file was addressed by adding a $-\mathbf{p}$ option to override the previously hard-coded yy variable prefix.

The description of the $-\mathbf{p}$ option specifies the minimal set of function and variable names that cause conflict when multiple parsers are linked together. YYSTYPE does not need to be changed. Instead, the programmer can use -b to give the header files for different parsers different names, and then the file with the yylex() for a given parser can include the header for that parser. Names such as *yyclearerr* do not need to be changed because they are used only in the actions; they do not have linkage. It is possible that an implementation has other names, either internal ones for implementing things such as yyclearerr, or providing non-standard features that it wants to change with $-\mathbf{p}$.

yacc Utilities

Unary operators that are the same token as a binary operator in general need their precedence adjusted. This is handled by the **%prec** advisory symbol associated with the particular grammar rule defining that unary operator. (See **Grammar Rules in yacc** (on page 1068).) Applications are not required to use this operator for unary operators, but the grammars that do not require it are rare.

41464 EXAMPLES

41459

41460

41461

41462 41463

41470

41471 41479

Access to the *yacc* library is obtained with library search operands to *c99*. To use the *yacc* library main():

```
41467 c99 y.tab.c -l y
```

Both the *lex* library and the *yacc* library contain *main*(). To access the *yacc main*():

```
41469 c99 y.tab.c lex.yy.c -l y -l l
```

This ensures that the *yacc* library is searched first, so that its *main()* is used.

The historical *yacc* libraries have contained two simple functions that are normally coded by the application programmer. These functions are similar to the following code:

```
41473
            #include <locale.h>
41474
            int main(void)
41475
                extern int yyparse();
41476
                setlocale(LC ALL, "");
41477
                /* If the following parser is one created by lex, the
41478
41479
                    application must be careful to ensure that LC CTYPE
                    and LC COLLATE are set to the POSIX locale. */
41480
41481
                 (void) yyparse();
41482
                return (0);
            }
41483
41484
            #include <stdio.h>
41485
            int yyerror(const char *msg)
            {
41486
                 (void) fprintf(stderr, "%s\n", msg);
41487
                return (0);
41488
41489
```

41490 RATIONALE

41491

41492

41493 41494

41495 41496

41497 41498

41499

41500 41501

41502

41503 41504 The references in **Referenced Documents** (on page xxxi) may be helpful in constructing the parser generator. The referenced DeRemer and Pennello article (along with the works it references) describes a technique to generate parsers that conform to this volume of IEEE Std 1003.1-2001. Work in this area continues to be done, so implementors should consult current literature before doing any new implementations. The original Knuth article is the theoretical basis for this kind of parser, but the tables it generates are impractically large for reasonable grammars and should not be used. The "equivalent to" wording is intentional to assure that the best tables that are LALR(1) can be generated.

There has been confusion between the class of grammars, the algorithms needed to generate parsers, and the algorithms needed to parse the languages. They are all reasonably orthogonal. In particular, a parser generator that accepts the full range of LR(1) grammars need not generate a table any more complex than one that accepts SLR(1) (a relatively weak class of LR grammars) for a grammar that happens to be SLR(1). Such an implementation need not recognize the case, either; table compression can yield the SLR(1) table (or one even smaller than that) without

Utilities yacc

recognizing that the grammar is SLR(1). The speed of an LR(1) parser for any class is dependent more upon the table representation and compression (or the code generation if a direct parser is generated) than upon the class of grammar that the table generator handles.

The speed of the parser generator is somewhat dependent upon the class of grammar it handles. However, the original Knuth article algorithms for constructing LR parsers were judged by its author to be impractically slow at that time. Although full LR is more complex than LALR(1), as computer speeds and algorithms improve, the difference (in terms of acceptable wall-clock execution time) is becoming less significant.

Potential authors are cautioned that the referenced DeRemer and Pennello article previously cited identifies a bug (an over-simplification of the computation of LALR(1) lookahead sets) in some of the LALR(1) algorithm statements that preceded it to publication. They should take the time to seek out that paper, as well as current relevant work, particularly Aho's.

The **-b** option was added to provide a portable method for permitting *yacc* to work on multiple separate parsers in the same directory. If a directory contains more than one *yacc* grammar, and both grammars are constructed at the same time (by, for example, a parallel *make* program), conflict results. While the solution is not historical practice, it corrects a known deficiency in historical implementations. Corresponding changes were made to all sections that referenced the filenames **y.tab.c** (now "the code file"), **y.tab.h** (now "the header file"), and **y.output** (now "the description file").

The grammar for *yacc* input is based on System V documentation. The textual description shows there that the ';' is required at the end of the rule. The grammar and the implementation do not require this. (The use of **C_IDENTIFIER** causes a reduce to occur in the right place.)

Also, in that implementation, the constructs such as **%token** can be terminated by a semicolon, but this is not permitted by the grammar. The keywords such as **%token** can also appear in uppercase, which is again not discussed. In most places where '\$' is used, '\' can be substituted, and there are alternate spellings for some of the symbols (for example, **%LEFT** can be "\$<" or even "\<").

Historically, <*tag*> can contain any characters except '>', including white space, in the implementation. However, since the *tag* must reference an ISO C standard union member, in practice conforming implementations need to support only the set of characters for ISO C standard identifiers in this context.

Some historical implementations are known to accept actions that are terminated by a period. Historical implementations often allow ' $$^{\circ}$'$ in names. A conforming implementation does not need to support either of these behaviors.

Deciding when to use **%prec** illustrates the difficulty in specifying the behavior of *yacc*. There may be situations in which the *grammar* is not, strictly speaking, in error, and yet *yacc* cannot interpret it unambiguously. The resolution of ambiguities in the grammar can in many instances be resolved by providing additional information, such as using **%type** or **%union** declarations. It is often easier and it usually yields a smaller parser to take this alternative when it is appropriate.

The size and execution time of a program produced without the runtime debugging code is usually smaller and slightly faster in historical implementations.

Statistics messages from several historical implementations include the following types of information:

n/512 terminals, n/300 non-terminals 41550 n/600 grammar rules, n/1500 states

n shift/reduce, n reduce/reduce conflicts reported

yacc Utilities

```
41552
             n/350 working sets used
41553
             Memory: states, etc. n/15000, parser n/15000
41554
             n/600 distinct lookahead sets
             n extra closures
41555
41556
             n shift entries, n exceptions
             n goto entries
41557
41558
             n entries saved by goto default
             Optimizer space used: input n/15000, output n/15000
41559
41560
             n table entries, n zero
41561
             Maximum spread: n, Maximum offset: n
             The report of internal tables in the description file is left implementation-defined because all
41562
             aspects of these limits are also implementation-defined. Some implementations may use
41563
             dynamic allocation techniques and have no specific limit values to report.
41564
             The format of the y.output file is not given because specification of the format was not seen to
41565
             enhance applications portability. The listing is primarily intended to help human users
41566
             understand and debug the parser; use of y.output by a conforming application script would be
41567
             unusual. Furthermore, implementations have not produced consistent output and no popular
41568
             format was apparent. The format selected by the implementation should be human-readable, in
41569
41570
             addition to the requirement that it be a text file.
             Standard error reports are not specifically described because they are seldom of use to
41571
             conforming applications and there was no reason to restrict implementations.
41572
             Some implementations recognize "={" as equivalent to '{' because it appears in historical
41573
41574
             documentation. This construction was recognized and documented as obsolete as long ago as
             1978, in the referenced Yacc: Yet Another Compiler-Compiler. This volume of IEEE Std 1003.1-2001
41575
             chose to leave it as obsolete and omit it.
41576
41577
             Multi-byte characters should be recognized by the lexical analyzer and returned as tokens. They
             should not be returned as multi-byte character literals. The token error that is used for error
41578
41579
             recovery is normally assigned the value 256 in the historical implementation. Thus, the token
             value 256, which is used in many multi-byte character sets, is not available for use as the value
41580
41581
             of a user-defined token.
41582 FUTURE DIRECTIONS
41583
             None.
41584 SEE ALSO
41585
             c99, lex
41586 CHANGE HISTORY
             First released in Issue 2.
41588 Issue 5
             The FUTURE DIRECTIONS section is added.
41589
41590 Issue 6
             This utility is marked as part of the C-Language Development Utilities option.
41591
             Minor changes have been added to align with the IEEE P1003.2b draft standard.
41592
             The normative text is reworded to avoid use of the term "must" for application requirements.
41593
41594
             IEEE PASC Interpretation 1003.2 #177 is applied, changing the comment on RCURL from the \}%
```

token to the %}.

Utilities zcat

41596 **NAME** zcat — expand and concatenate data 41597 41598 SYNOPSIS zcat [file...] 41599 XSI 41600 41601 **DESCRIPTION** The zcat utility shall write to standard output the uncompressed form of files that have been 41602 compressed using the *compress* utility. It is the equivalent of *uncompress* –c. Input files are not 41603 affected. 41604 41605 OPTIONS None. 41606 41607 OPERANDS The following operand shall be supported: 41608 file The pathname of a file previously processed by the *compress* utility. If *file* already 41609 has the .Z suffix specified, it is used as submitted. Otherwise, the .Z suffix is 41610 41611 appended to the filename prior to processing. 41612 **STDIN** 41613 The standard input shall be used only if no *file* operands are specified, or if a *file* operand is '-'. 41614 INPUT FILES 41615 Input files shall be compressed files that are in the format produced by the *compress* utility. 41616 ENVIRONMENT VARIABLES The following environment variables shall affect the execution of zcat: 41617 LANG Provide a default value for the internationalization variables that are unset or null. 41618 (See the Base Definitions volume of IEEE Std 1003.1-2001, Section 8.2, 41619 41620 Internationalization Variables for the precedence of internationalization variables used to determine the values of locale categories.) 41621 LC ALL If set to a non-empty string value, override the values of all the other 41622 internationalization variables. 41623 Determine the locale for the interpretation of sequences of bytes of text data as 41624 LC_CTYPE characters (for example, single-byte as opposed to multi-byte characters in 41625 arguments). 41626 LC_MESSAGES 41627 Determine the locale that should be used to affect the format and contents of 41628 diagnostic messages written to standard error. 41629 Determine the location of message catalogs for the processing of *LC_MESSAGES*. **NLSPATH** 41630 41631 ASYNCHRONOUS EVENTS Default. 41632 41633 **STDOUT** The compressed files given as input shall be written on standard output in their uncompressed 41634 form. 41635

The standard error shall be used only for diagnostic messages.

41636 STDERR

41637

zcat **Utilities**

41638 OUTPUT FILES 41639 None. 41640 EXTENDED DESCRIPTION None. 41641 41642 EXIT STATUS 41643 The following exit values shall be returned: 0 Successful completion. 41644 >0 An error occurred. 41645 41646 CONSEQUENCES OF ERRORS 41647 Default. 41648 APPLICATION USAGE

41649 None.

41650 EXAMPLES None. 41651

41652 RATIONALE 41653 None.

41654 FUTURE DIRECTIONS

None. 41655

41656 SEE ALSO

41657 compress, uncompress

41658 CHANGE HISTORY

First released in Issue 4. 41659

Index

<control>-V</control>	368	string functions	167
<control>-W</control>	368	user-defined functions	
_CFLAGS		variables and special variables	
_ _LDFLAGS		background work	
_ _LIBS		at	144
POSIX VDISABLE		batch	
Account_Name		bg	
actions equivalent to functions		crontab	
admin		fg	
ADV		jobs	
AIO		nice	
alias2		nohup	
alias substitution		renice	
AND lists		BAR	
AND-OR list		basename	
appending redirected output		batch	
ar		batch administration	
archives		batch authorization	
ar command	134	batch client-server interaction	
ARG_MAX		batch environment	
arithmetic expansion		services	101
arithmetic language		utilities	
bc	193	Batch Job Abort	
arithmetic precision and operations		batch job creation	•
array identifiers		Batch Job Execution	
asa		Batch Job Exit	
asynchronous lists		batch job identifier	
at		Batch Job Message Request	
at-job		Batch Job Routing	
automatic storage class		batch job states	
awk		Batch Job Status Request	
actions		batch job tracking	
arithmetic functions		batch notification	
escape sequences		batch queue	
expression patterns		Batch Queue Status Request	
expressions		Batch Server Restart	
functions		batch services	
grammar		bc	
input/output and general functions		grammar	
lexical conventions		lexical conventions	
output statements		operations	
overall program structure		operators	
pattern ranges		bcc (mailer blind carbon copy)	608
patterns		BC_BASE_MAX	
regular expressions		BC_DIM_MAX	
special patterns		BC SCALE MAX	

BC_STRING_MAX18, 196	uux	980
BE 10	write	1054
bg28, 48, 208	compilers	
binary primaries908	c99	211
break 65	fort77	464
built-in utilities28	yacc	1063
builtin263	compound commands	52
c99 211	compound-list	
external symbols215	compress	
standard libraries214	compression	
cal 220	compress	264
can1	uncompress	
carriage-control characters141	zcat	1079
case conditional construct53	concurrent execution of processes	
cat	configuration values	
cc (mailer carbon copy)608	conforming application	
CD 10	consequences of shell errors	
cd	continue	
cflow	control characters	
changing the current working directory7	controlling terminal	
character counting1045	Coordinated Universal Time (UTC)	
charmap	copy files commands	965
with localedef	dd	
writing names with locale553		
charmap file557, 892	ln	
CHAR_BIT348	mv	
Checkpoint108	pax	
checksums	cp	
cksum248	cpio format	
chgrp	CPT	
chmod237	CPU time	
grammar240	cron daemon	
chown244	crontab	
cksum248	CS	
cmp 253	csplit	
codeset conversion506	ctags	
tr926	current working directory	3
COLL_WEIGHTS_MAX18	cut	288
colon 67 , 69	CX	
comm 256	cxref	292
command28, 48, 259	data keywords	747
command mode336	date	295
command search and execution48	conversion specifications	295
command substitution40	modified conversion specifications	
communications commands	dd	
mailx586	default queue	
talk900	deferred batch services	
uucp965	Delete Batch Job Request	
uudecode969	delta	
uuencode	destination	
uustat	df	

diff	319	read command	
binary output format	321	regular expressions	
default output format	321	shell escape command	346
directory comparison format	320	substitute command	
-c or -C output format	322	undo command	345
–e output format	322	write command	346
-f output format		edit buffer	355, 987
directory commands		edit line	
cd	226	editors	
pwd		ed	336
directory lister		ex	355
dirname		sed	
disk space commands		vi	
df	315	ED FILE MAX	
du		ED_LINE_MAX	
ulimit		effective group ID	
documentation		effective user ID	
dot		Eighth Edition UNIX	
double-quotes		env	
du		EPERM	
duplicating an input file descriptor		Error_Path	
duplicating an output file descriptor		escape character (backslash)	
echo		•	
		escape sequences	106
ed		awk	
addresses		gencat	
append command		lex	
change command		establish the locale	
commands		eval	
copy command		ex	
delete command		<backslash></backslash>	
edit command		<control>-D command</control>	
edit without checking command		<newline></newline>	
filename command		abbreviate command	
global command		addressing	
global non-matched command	346	adjust window command	
help command		append command	
help-mode command		args command	
insert command		autoindent option	
interactive global command	342	autoprint option	393
interactive global not-matched comm	and346	autowrite option	
join command	343	beautify option	393
line number command	346	change command	371
list command	343	chdir command	372
mark command	343	command descriptions	368
move command	343	copy command	372
null command	347	delete command	
number command		directory option	
print command		edcompatible option	
prompt command		edit command	
quit command		edit options	
quit without checking command		errorbells option	
		1	

escape command		tag command	
execute command	391	taglength option	397
exrc command	394	tags command	398
file command	373	term command	398
global command	374	terse command	398
ignorecase option	394	unabbrev command	385
initialization		undo command	385
input editing	366	unmap command	385
insert command		version command	
join command		visual command	
list command		warn command	
magic command		window command	
map command		wrapmargin option	
mark command		wrapscan option	
mesg command		write command	
move command		write line number command	
next command		writeany option	
number command		xit command	
number option		yank command	
open command		exec	
paragraphs option		exec family28	
paragrapus option preserve		Execution_Time	
		EXINIT	
preserve command		exit	
print command			
prompt command		exit status and errors	
put command		exit status for commands	
quit command		expand	
read command		export	
readonly command		expr	
recover command		matching expression	
redraw command		string operand	
regular expressions		expression argument	
remap command		expression list	
replacement strings		EXPR_NEST_MAX	
report command		extended regular expression	
rewind command		455, 495, 539, 657	', 710, 824, 1059
scroll command		extension	
sections command		CX	10
set command		OH	
shell command		XSI	
shell option	397	false	28, 48, 434
shift left command	390	fc	28, 48, 436
shift right command	390	FD	
shiftwidth option	397	fg	28, 48, 442
showmatch option	397	field splitting	
showmode command		FIFO special files	
slowopen command		file	
source command		file access permissions	
substitute command		file comparisons	
suspend command		cmp	253
tabstop option		comm	256
1 1			

diff		more	
uniq	959	nl	
file contents	6	paste	
file conversion		pax	
cut	288	pr	736
dd	302	read	816
expand	426	sed	841
fold	461	tail	897
head	503	tee	904
join	525	tr	926
od	679	uncompress	951
paste	687	unexpand	
patch		zcat	
sort		find	
strings		fold	
tail		for loop	
tr		fort77	
tsort		external symbols	
unexpand		standard libraries	
uniq		FR	
uudecode		FSC	
uuencode		function definition command	
file creation		function identifiers	
file descriptor		fuser	
file mode creation mask		g-file	
file permission commands		gencat	
chgrp	934	escape sequences	
chmod		generated file	
chown		get	
umask		getconf	
file read		getopts	
file removal		global storage class	
file searching	0	GNU make	
grep	405	grep	
file time values		grouping commands	
file tree commands	0	hash	
diff	210	head	
C J	470		
ls		11010 40041110111	44
mkdir		history command fc	196
rmdir		Hold Batch Job Request	
file write		<u> </u>	
	4	Hold_Types HOME	
filters	1.41		
asa		hunk	
awk		iconv	
compress		id	
dd		if conditional construct	
expand		implementation-defined	
fold		inference rule	
head		input field separator	
iconv	506	input mode	336

IP6	11	delete aliases	
ipcrm	513	delete messages	597
ipcs	515	delete messages and display	598
I_ISVTX	239	direct messages to mbox	600
jobs	28, 48, 521	discard header fields	597
Job_Owner	110	display beginning of messages	603
join	525	display current message number	
Join_Path	110	display header summary	
Keep_Files		display list of folders	
keyword-value pairs		display message	
kill		display message size	
legacy	1	echo a string	
lex		edit message	
actions		execute commands conditionally	
definitions		exit	
escape sequences		follow up specified messages	
regular expressions		help	
rules		hold messages	
table sizes		internal variables	
user subroutines		invoke a shell	
lex, translation table		invoke shell command	
libraries		list available commands	
ar command	134	mail a message	
LIMIT		null command	
line counting		pipe message	
LINE_MAX18, 154, 348-349, 356		process next specified message	
link		quit	60
lists		read mailx commands from a file	
AND-OR		receive mode	
compound-list		reply to a message	
ln		reply to a message list	
locale		retain header fields	
localedef		save messages	
Locate Batch Job Request		scroll header display	
locking file		send mode	
logger		set variables	
logname		start-up	
lp		touch messages	60°
LR(1) grammars		undelete messages	
ls		unset variables	
m4		write messages to a file	
macro processor		Mail_Points	
magic file		Mail_Users	
mailx		make	
change current directory		default rules	
· ·		inference rules	
change foldercommand escapes		internal macros	
commands		libraries	
copy messages		macros	
declare aliases		makefile execution	
declare allasesdeclare alternatives			
ueciare anternatives	397	makefile syntax	013

target rules	614	MX	
make, GNU version		NAME_MAX	-
man	631	newgrp	
mathematical functions		NGROUPS_MAX	
may	2	nice	
MC1	11	Ninth Edition UNIX	205, 335, 745
MC2	11	nl	
MC3	11	nm	672
mesg	635	noclobber option	700
message catalog generation	473	nohup	
MF	11	non-printable	348, 848, 903
MIL-STD-1753	468	ОВ	
Minimum_Cpu_Interval	108	object files	672
mkdir		od	679
mkfifo	641	OF	12
ML	11	OH	13
MLR	12	open file descriptors for rea	ding and writing46
Modify Batch Job Request	117	open mode	
MON		option	
more		ADV	
discard and refresh		AIO	
display position		BAR	
examine new file		BE	
examine next file		CD	
examine previous file		CPT	
go to beginning of file		CS	
go to end-of-file		FD	
go to tag		FR	
help		FSC	
invoke editor		IP6	
mark position		MC1	
quit		MC2	
refresh the screen		MC3	
repeat search		MF	
repeat search in reverse		ML	
return to mark		MLR	
return to previous position		MON	
scroll backward one half screenful		MPR	
scroll backward one line		MSG	
scroll backward one screenful		MX	
scroll forward one half screenful		PIO	
scroll forward one line		PS	
scroll forward one screenful		RS	
search backward for pattern		RTS	
search forward for pattern		SD	
skip forward one line		SEM	
motion command		SHM	
Move Batch Job Request		SIO	
MPR		SPI	
MSG		SPN	
mv		SS	
111 7		00	17

TCT14	cpio file data	723
TEF14	cpio filename	723
THR14	cpio header	721
TMO15	cpio interchange format	721
TMR15	cpio special entries	723
TPI15	extended header	
TPP15	extended header file times	717
TPS15	extended header keyword precedenc	e717
TRC15	list mode format specifications	
TRI15	ustar format	
TRL15	ustar interchange format	
TSA16	PIO	
TSF16	pipelines	
TSH16	portable character set	
TSP	positional parameters	
TSS	POSIX2_BC_BASE_MAX	
TYM	POSIX2_BC_DIM_MAX	
UP	POSIX2_BC_SCALE_MAX	
XSR	POSIX2_BC_STRING_MAX	
OR lists51	POSIX2_COLL_WEIGHTS_MAX	
ordinary identifiers198	POSIX2_EXPR_NEST_MAX	
Output_Path111	POSIX2_LINE_MAX	
paginators	POSIX2_RE_DUP_MAX	
more643	POSIX2_SYMLINKS	
parameter expansion37	pr	
parameter expansion	print-related commands	<i>1</i> JU
paste687	fold	161
patch	lp	
filename determination	pr	
patch application694	printf	
patch file format693	Priority	
•		
pathchk	privileges	
pathname expansion	process attributesprocess group ID	
pathname manipulation basename187	process group ID	
	process ID	
dirname	process status report	
pathchk697	prs PS	
pathname resolution7	1 0	
PATH_MAX18, 733, 826	ps	
pattern matching452, 710, 966, 982	public locale	
definition62	pwd	
in case statements53	qalter	
in shell variables39	qdel	
pattern matching notation62, 458, 726	qhold	
pattern scanning and processing language	qmove	
at153	qmsg	
patterns matching a single character62	qrerun	
patterns matching multiple characters	qrls	
patterns used for filename expansion63	qselect	
pax	qsig	
archive character set encoding/decoding732	qstat	798

qsub		addresses	
Queue Batch Job Request		editing commands	
quote removal	42	regular expressions	
quoting	30	Select Batch Jobs Request	120
read	28, 48, 816	SEM	13
readonly	82	sequential lists	51
real group ID		Server Shutdown Request	120
real user ID		Server Status Request	
redirecting input		session membership	
redirecting output		set	
redirection		set-group-ID	
regular expressions161		set-user-ID	
391, 431, 455, 495, 539		set-user-ID scripts	
707, 824, 843, 10		shsh	
related to shell patterns		command history list	
relational database operator		command line editing	
Release Batch Job Request		vi line editing command mode	
remove directories		vi line editing insert mode	
remove files		vi-mode command line editing	
renice		shall	
requested batch services		shell command language	
Rerun Batch Job Request		alias substitution	
Rerunable		appending redirected output	
reserved words		arithmetic expansion	
Resource_List		command substitution	
return		compound commands	
RE_DUP_MAX		consequences of shell errors	
rm	823	double-quotes	
rmdel	828	duplicating an input file descriptor	45
rmdir	830	duplicating an output file descriptor	45
root directory	3	escape character (backslash)	30
RS	13	exit status and errors	46
RTS	13	exit status for commands	46
sact	833	field splitting	42
saved set-group-ID	3	function definition command	
saved set-user-ID		grammar	
SCCS		here-document	
SCCS commands		introduction	0.0
admin	126	lists	
delta		open file descriptors for reading and	
get		parameter expansion	
prs		parameters and variables	
rmdel		pathname expansion	
sact		pattern matching notation	
SCCS			
		patterns matching a single character	
unget		patterns matching multiple characte	
val		patterns used for filename expansion	
what		pipelines	
SD		positional parameters	
search pattern		quote removal	
sed	841	quoting	30

redirecting input44	export79
redirecting output44	readonly82
redirection43	return84
reserved words33	set86
shell commands47	shift92
shell execution environment61	times94
shell grammar lexical conventions55	trap96
shell grammar rules56	unset99
shell variables34	special parameters34
signals and error handling61	special targets615
simple commands47	SPI 1 4
single-quotes30	split 87 6
special built-in utilities64	split files
special parameters34	csplit279
tilde expansion37	split876
token recognition31	SPN 1 4
word expansions36	spoofing83
shell commands47	SS14
shell execution environment61, 132, 818, 944	standard error43
shell grammar55	standard input43
shell grammar lexical conventions55	standard output43
shell grammar rules56	strings 87 9
shell introduction29	strip 88 2
shell variables34	stty 88 4
Shell_Path_List112	combination modes889
shift 92	control modes884
SHM 13	input modes885
should2	local modes887
SIGCONT358	output modes886
SIGHUP337, 358, 987, 1028	special control character assignments888
SIGINT313, 337, 357, 1040	st_gid139
Signal Batch Job Request121	st_mode139
signal processes530	st_mtime139
signals and error handling61	st_size139
SIGQUIT337	st_uid139
SIGTERM358	superuser440, 577, 725
simple commands47	supplementary group IDs3
single-quotes30	system configuration values484
SIO14	system name948
sleep 867	tabs 89 3
SLR(1) grammars1077	tag file creation283
sort 870	tail 89 7
special built-in263, 666, 678, 759, 864, 917, 933	talk900
special built-in utilities64	tar format717
break65, 69	target queue118
characteristics64	target rule610
colon67	TCT14
dot71	tee 90 4
eval73	TEF14
exec75	terminal characteristics
exit77	stty884
	-

tabs	893	user identity	
tput	923	id	509
tty		logname	564
terminate processes		newgrp	
terminology		who	
test		User_List	
THR		ustar format	
tilde expansion		utility option parsing	
time		uucp	
times	94	uudecode	
TMO		uuencode	
TMPDIR		uustat	
TMR		uux	
token recognition		val	
touch		Variable_List	
TPI		vi	
TPP		<esc></esc>	
TPS		append	
tput		change	
tr		change to end-of-line	
Track Batch Job Request		clear and redisplay	
trap		command descriptions	
TRC		control-D	
TRI		control-H	
TRL		control-T	
trojan horse		control-U	
true		control-V	
TSA	, ,	current and line above	
TSF		deletedelete	
TSH		delete character	
tsort		delete to end-of-line	
TSP		display information	
TSS		edit the alternate file	
tty		enter ex mode	
TYM		execute	
type		execute an ex command	
ulimit		exit	
umask		find character	
unalias	· ·	find regular expression	
uname		Initialization	
unary primaries		input mode commands	
uncompress		insert	
undefined		insert empty line	
unexpand		join	
unget		mark position	
uniq		move back	
unlink		move cursor	
unset		move down	
unspecified		move down	
until loop		move to bigword	
UP		move to bigwordmove to bottom of screen	
O1	10	move to bottom of screen	1013

move to first character in line	1005
move to first non- <blank></blank>	1001
move to line	1012
move to matching character	
move to middle of screen	1014
move to next section	1001
move to specific column	1003
move to top of screen	1012
move to word101	1, 1018
move up	995
newline	1025
nul102	4, 1027
page backwards	993
page forward	994
put from buffer	
redraw screen	996
redraw window	1021
regular expression	1007
repeat	
repeat find100	
repeat substitution	
replace character101	
replace text with command	
return to previous context99	
return to previous section	
reverse case	
reverse find character	
scroll backward	
scroll backward by line	
scroll forward	
scroll forward by line	
search for tagstring	
shift left	
shift right	
substitute character	
substitute lines	
terminate command or input mode	
undo	
undo current line	
yank	
yank current line	
visual mode	
wait28, 4	
warning	,
OB	12
OF	
wc	
what	
while loop	
who	
word counting	

word expansions	36
write	1054
kargs	1057
XSI	
XSR	17
yacc	1063
algorithms	1074
code file	1065
completing the program	1074
conflicts	
debugging the parser	1074
declarations section	1066
description file	1065
error handling	1072
grammar rules	1068
header file	1065
input grammar	1070
input language	1065
interface to the lexical analyzer	1073
lexical structure of the grammar	
library	1074
limits	1075
programs section	1069
zcat	