INTERNATIONAL STANDARD



ISO/IEC 9945-1

IEEE Std 1003.1 First edition 1990-12-07

REFERENCE

Intermation Technology -

Portable Operating System Interface (POSIX) – Part 1: System Application Program Interface (API) [C Language]

> An Industry-Recognized Operating Systems Interface Standard based on the UNIX® Operating System, providing:



 A comprehensive applications environment
 Applications portability at the source-code level
 An applications interface to I/O, file system access, and process management facilities

Developed for:

- Software Applications Developers
 System Designers/Engineers
 - Hardware and Software Purchasers
 - End users operating a UNIX
 OS environment

UNIX is a registered trademark of UNIX System Laboratories in the U.S. and other countries

JK 486 . A8A3 ∦151-2 1990

Published by the Institute of Electrical and Electronics Engineers, Inc. *SH13680* This easy-to-use International Standard features informative footnotes, line markings, and diff markings.

Reference number ISO/IEC 9945-1 : 1990 (E) IEEE Std 1003.1-1990 This standard has been adopted for Federal Government use.

Details concerning its use within the Federal Government are contained in Federal Information Processing Standards Publication 151-2, Portable Operating System Interface (POSIX) -- system Application Program Interface [C Language]. For a complete list of the publications available in the Federal Information Processing Standards Series, write to the Standards Processing Coordinator (ADP), National Institute Standards and Technology, Gaithersburg, MD 20899

> Second Printing January 1992

ISBN 1-55937-061-0

Library of Congress Catalog Number 90-084554

Quote in 8.1.2.3 on Returns is taken from X3.159-1989, developed under the auspices of the American National Standards Accredited Committee X3 Technical Committee X3J11, Computer and Business Equipment Manufacturers Association (CBEMA), 311 First St, N.W., Suite 500, Washington, DC 20001.

©Copyright 1990 by

The Institute of Electrical and Electronics Engineers, Inc. 345 East 47th Street, New York, NY 10017, USA

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without the prior written permission of the publisher. International Standard ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990 (Revision of IEEE Std 1003.1-1988)

Information technology—Portable Operating System Interface (POSIX) Part 1: System Application Program Interface (API) [C Language]

Sponsor

Technical Committee on Operating Systems and Application Environments of the IEEE Computer Society

> Approved September 28, 1990 IEEE Standards Board

Approved 1990 by the International Organization for Standardization and by the International Electrotechnical Commission

Abstract: ISO/IEC 9945-1: 1990 (IEEE Std 1003.1-1990), Information technology—Portable Operating System Interface (POSIX)—Part 1: System Application Program Interface (API) [C Language] is part of the POSIX series of standards for applications and user interfaces to open systems. It defines the applications interface to basic system services for input/output, file system access, and process management. It also defines a format for data interchange. This standard is stated in terms of its C binding.

Keywords: API, application portability, C (programming language), data processing, information interchange, open systems, operating system, portable application, POSIX, programming language, system configuration computer interface



Adopted as an International Standard by the International Organization for Standardization and by the International Electrotechnical Commission

Published by The Institute of Electrical and Electronics Engineers, Inc.

IEEE Standards documents are developed within the Technical Committees of the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Board. Members of the committees serve voluntarily and without compensation. They are not necessarily members of the Institute. The standards developed within IEEE represent a consensus of the broad expertise on the subject within the Institute as well as those activities outside of IEEE that have expressed an interest in participating in the development of the standard.

Use of an IEEE Standard is wholly voluntary. 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.

Comments for revision of IEEE Standards are welcome from any interested party, regardless of membership affiliation with IEEE. Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments.

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 all concerned interests, it is important to ensure that any interpretation has also received the concurrence of a balance of interests. For this reason, the IEEE and the members of its technical 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 on standards and requests for interpretations should be addressed to:

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

IEEE Standards documents are adopted by the Institute of Electrical and Electronics Engineers without regard to whether their adoption may involve patents on articles, materials, or processes. Such adoption does not assume any liability to any patent owner, nor does it assume any obligation whatever to parties adopting the standards documents.

Contents

																						PAGE
Forewor	rd	•	• •	٠	•	•	•	•	•	•	•	٠	٠	٠	•	•	•	•	•	•	•	viii
Introdu	ction .	•	• •	•	•	•	•	•	•	•	•	٠	٠	٠	٠	٠	•	•	•	•	٠	ix
Section	1: Gen	eral	• •	•	•	•	•		•	•	•	•	•	•	•	•	•	•		•	•	1
1.1	Scope		• •			•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	1
1.2	Norma	tive]	Refe	renc	es	•	•	•	•	•	•	•	•	•	•	•	•	•	•		•	2
1.3	Confor	mano	e.	٠	•	٠	•	•	٠	•	•	•	•	•	•	•	•	•	•	•	•	2
Section	2: Terr	ninol	ogy a	and	Ge	nei	ral	Re	ດານi	rei	me	nts										9
2.1	Conve	ntion	s .						9 m.						•	•	•	•		•	•	9
2.2	Definit	ions														•						10
2.3	Genera	al Coi	ncep	ts														•		•	•	21
2.4	Error 1	Numł	bers					•					•			•		•			•	23
2.5	Primit	ive S	vster	n D	• ata	T	, me	s					Ţ	Ţ.	÷						•	27
2.6	Enviro	nmer	t De	n ~	inti	on	pe		•	•	•	•	•	•	•	•	•	•	•	•	•	27
2.7	C Land	סוומסב	Def	init	ion	с.		•	•	•		•	•	•	•	•	•	•	•	•	•	29
2.8	Numer	rical l	limi	te	1011	0	•	•	•	•	•	•	•	•	٠	•	•	•	•	•	٠	34
2.9	Symbo	lic Co	nsta	ante	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	37
1.0	Oy III OO		511500	11105		•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	01
Section	3: Proc	ess P	rimi	tive	s	•	•	•	•	•		•	•		•	•	•	•		•	•	41
3.1	Proces	s Cre	atio	ı an	d F	Exe	cut	tior	ı	•					•						•	41
	3.1.1	Proc	ess (Crea	atio	n	•	•		•	•	•	•	•	•	•	•	•	•	•	•	41
	3.1.2	Exec	ute	a Fi	le		•		•	•	•	•	•			•	•			•	•	42
3.2	Proces	s Ter	mina	atio	n								•								•	46
	3.2.1	Wait	t for	Pro	ces	sТ	eri	nin	ati	on			•								•	47
	3.2.2	Tern	nina	te a	Pr	oce	SS				•	•	•				•		•		•	49
3.3	Signal	s.																				51
	3.3.1	Sign	al C	once	ente	3																51
	3.3.2	Send	laS	igna	al t	- 0 a	Pr	oce	SS				•			•						56
	3.3.3	Man	inul	ate l	Sig	nal	S	ets														57
	3.3.4	Exar	nine	and	1 C	hai	nge	Si	gn	al /	Act	ion	i Č									58
	3.3.5	Exar	nine	and	- C	hai	nge	S RI	ocl	ced	Si	ion	als	•		÷	•	•				60
	336	Exar	nine	Pei	ndi	no	Sic	ma	ام ام						Ţ	Ť.						62
	337	Wait	for	a Si	on	al	016	5110	10	•	•	•	•	•	•	•	•			•		62
3.4	Timer	Oper	ation	19	6				•	•	•	•		•	•		•		Ţ	Ţ	Ţ	63
011	3 4 1	Sche	dule		• arm	•	•	•	•	•	•	•	•	•	•	•		•	•	•	•	63
	342	Sust	hend	Pro	~	e Ti	lvo	• cut	ion		•	•	•	•	•	•	•	•	•	•	•	64
	343	Dela	v Pr	110	e F	lvo		tion	101	L	•	•	•	•	•	•	•	•	•	•	•	65
	0.110	~010	J I I	0000	ע כוי	AU	cui			•	•	•	•	•	•	•	•	•	•	•	•	00
Section	4: Proc	ess E	nvir	onn	ıen	t		•		•	•	•	•					•	•		•	67
4.1	Proces	s Ide	ntific	atio	n	•	•	•	•	•		•	•	•	•			•		•	•	67
	4.1.1	Get	Proc	ess	and	l P	are	ent	Pro	oce	SS	IDs	;	•	•	•	•	•	•	•	•	67

PAGE

4.2	User I	dentification	•	•	• •	•	68
	4.2.1	Get Real User, Effective User, Real Group, and	Ef	fect	ive		
		Group IDs	•	•	• •		68
	4.2.2	Set User and Group IDs	•	•	• •		6 8
	4.2.3	Get Supplementary Group IDs	•	•	• •		70
	4.2.4	Get User Name					71
4.3	Proces	s Groups					72
-10	431	Get Process Group ID	•		• •		72
	432	Create Session and Set Process Group ID	•	•	• •	• •	72
	433	Set Process Group ID for Job Control	•	•	• •	•	73
44	System	n Identification	•	•	• •	•	7/
		Get System Name	•	•	• •	•	74
15	Timo		•	٠	• •	•	75
4.0	151	Cot System Time	•	•	• •	•	75
	4.0.1		•	•	• •	•	10
4.0	4.0.2	Get Process Times	•	•	• •	•	76
4.6	Enviro	onment variables	•	•	• •	•	77
	4.6.1	Environment Access	•	•	• •	•	111
4.7	Termi	nal Identification	•	•	• •	•	78
	4.7.1	Generate Terminal Pathname	•	•	• •	•	78
	4.7.2	Determine Terminal Device Name	•	•	• •	•	79
4.8	Config	urable System Variables	•	•	• •	•	80
	4.8.1	Get Configurable System Variables	•	•	• •	•	80
Section	5. File	s and Directories					83
5 1	Direct		•	•	• •	•	83
0.1	511	Format of Directory Entries	•	•	• •	•	63
	510	Directory Entries	•	•	• •	•	00
5.9	U.I.Z		•	•	• •	•	00
0.4	F O 1	Observe Oscillation Diverte et al.	•	•	• •	•	00
	0.2.1	Change Current working Directory	•	•	• •	•	80
F 0	5.2.2	Get Working Directory Pathname	•	•	• •	•	87
5.3	Genera		•	•	• •	•	88
	5.3.1	Open a File	•	•	• •	•	88
	5.3.2	Create a New File or Rewrite an Existing One	•	•	• •	•	91
	5.3.3	Set File Creation Mask	•	•	• •	•	91
	5.3.4	Link to a File	•	•	• •	•	92
5.4	Specia	l File Creation	•	•	• •	•	94
	5.4.1	Make a Directory	•	•	• •	•	94
	5.4.2	Make a FIFO Special File	•	•	• •	•	95
5.5	File Re	emoval	•	•		•	96
	5.5.1	Remove Directory Entries	•	•		•	96
	5.5.2	Remove a Directory	•		• •		98
	5.5.3	Rename a File					99
5.6	File Cl	haracteristics					101
0.10	5.6.1	File Characteristics: Header and Data	•	•	• •	•	-01
		Structure		•			101
	5.6.2	Get File Status		•			103
	5.6.3	Check File Accessibility		•			104
	5.6.4	Change File Modes					106
	5.6.5	Change Owner and Group of a File					107
	3.3.0	shange entitle and eroup of a fire of the	*	-	- •		-0.

	5.6.6	Set File	Access and I	Modifica	ation '	<u> Fime</u>	s	•	•	•	•	•	•	•	108
5.7	Config	urable Pa	ithname Var	riables	•	• •	•	•	•	•	•		•	•	110
	5.7.1	Get Con	fi <mark>gu</mark> rable Pa	thname	Varia	ables	5	•	•	•	•	•	•	•	110
Section	6: Inpu	it and Ou	tput Primiti	ves .		• •	•	•	•	•	•	•	•	•	113
6.1	Pipes				•	• •	•	•	•	•	•	•	•	•	113
	6.1.1	Create a	n Inter-Proc	ess Cha	annel	•	•	•	•	•	•			•	113
6.2	File D	escriptor	Manipulatio	n	•	• •	•	•	•	•	•				114
	6.2.1	Duplicat	e an Òpen F	'ile Des	cripto	r.	•	•	•	•	•			•	114
6.3	File D	escriptor	Deassignme	nt.			•	•	•						115
	6.3.1	Close a l	File												115
6.4	Input	and Outn	ut												116
	641	Read fro	maFile .												116
	642	Write to	a File	• • •	•	• •	•	•	•	•	•	•	•	•	118
65	Contro	Oporati	one on Files	• • •	• •	• •	•	•	•	•	•	•	•	•	191
0.0	651	Doto Do	finitiona for		• •	• • Onor	• •	•	•	•	•	•	•	•	101
	0.0.1	Data De		rne co.		Oper	au	JUS		•	•	•	•	•	101
	0.0.2	File Con		• • • •		• •	•	•	•	•	•	•	•	•	121
	0.0.3	Repositi	on Read/wri	te rile	Uliset	•	•	•	•	•	•	•	•	•	127
a		• 10			•										100
Section	7: Dev	ice- and (lass-Specific	e Funct	ions	• •	٠	•	•	٠	٠	•	•	•	129
7.1	Gener	al Termir	al Interface	• • •	•	• •	•	•	•	•	•	•	•	•	129
	7.1.1	Interface	e Characteri	stics	•	• •	•	٠	•	٠	٠	•	•	•	129
		7.1.1.1	Opening a '	Fermina	al Dev	rice I	File		•	•	٠	•	•	•	129
		7.1.1.2	Process Gro	oups .	•	• •	•	•	•	•	•	٠	٠	٠	129
		7.1.1.3	The Contro	lling Te	rmina	al.	•	•	•	٠	•	٠	•	٠	130
		7.1.1.4	Terminal A	ccess C	ontrol	ι.		•	•	•	•	•	•	•	130
		7.1.1.5	Input Proce	ssing a	nd Re	adin	g D)ata	ı	•	•			•	131
		7.1.1.6	Canonical N	Mode In	put P	roce	ssin	g	•	•	•	•		•	132
		7.1.1.7	Noncanonio	al Mod	e Inpi	ıt Pr	oce	ssii	ıg		•			•	132
		7.1.1.8	Writing Da	ta and (Jutpu	it Pr	oces	ssir	ng		•	•		•	133
		7.1.1.9	Special Cha	racters											133
		7.1.1.10	Modem Dis	connect											135
		7 1 1 11	Closing a T	erminal	l Devi	ce F	ile	•						•	135
	712	Paramet	ers That Ca	n Be Se	t1					Ť	Ť.	Ť	Ţ.	•	135
	•••±••	7191	termice St	ni De De	•••	•••	•	•	•	•	•	•		•	135
		7199	Input Mode	a	•	• •	•	•	•	•	•	•	•	•	136
		7199	Output Mode		• •	• •	•	•	•	•	٠	•	•	•	127
		7104	Control Mo		•	• •	•	•	•	•	•	•	•	•	107
		7.1.2.4	Control Mo	aes .	•	• •	•	•	•	•	•	•	•	•	100
		7.1.2.0	Local Mode	\mathbf{S}	• •	• •	•	٠	•	•	•	•	•	•	139
		7.1.2.6	Special Con	trol Ch	aracte	ers	•	٠	•	•	•	•	•	•	140
		7.1.2.7	Baud Kate	Values	•	• •	•	•	•	•	•	•	•	•	141
	7.1.3	Baud Ra	te Functions	5	•	• •	•	•	•	•	•	•	•	•	141
		7.1.3.1	Synopsis	• • •	•	• •	•	•	•	•	•	•	•	•	141
		7.1.3.2	Description	• • •	•	• •	•	•	•	•	•	•	•	•	142
		7.1.3.3	Returns .	• • •	•	• •	٠	•	•	٠	•	•	•	•	142
		7.1.3.4	Errors .	• • •	•	• •	•	•	•	•	•	•	•	•	142
		7.1.3.5	Cross-Refer	ences .	•	• •	•	•	•	•	•	•	•	•	142
7.2	Gener	al Termin	al Interface	Contro	l Fun	ction	S	•	•	•	•	•	•	•	143
	7.2.1	Get and	Set State .			• •	•	•	•	•			•	•	143

PAGE

																PAGE
	7.2.2	Line Control	Functions	•	• •			•	٠	•	•	•	•	•	•	145
	7.2.3	Get Foregrou	nd Proces	s Gr	oup	DID			•							147
	7.2.4	Set Foregrou	nd Process	s Gr	oup	ID										148
		0			1											
$\mathbf{Section}$	8: Lan	guage-Specific	Services	for t	he (C Pr	ogr	am	mi	ng						
Lang	uage .	• • • •	• • • •	•	•		٠	•			•					151
8.1	Refere	nced C Langu	age Routii	nes		• •					•					151
	8.1.1	Extensions to	Time Fu	nctio	ons	•	•	•	•							152
	8.1.2	Extensions to	setlocal	e() F	un	ctio	n						•			154
8.2	C Lang	guage Input/O	utput Fur	ictio	ns	•										155
	8.2.1	Map a Stream	n Pointer	to a	File	e De	scr	ipt	or							156
	8.2.2	Open a Strea	m on a Fil	le De	escr	ipto	r									157
	8.2.3	Interactions	of Other F	ILE	Tvi	be C	Fu	inc	tion	is						158
	8.2.4	Operations of	n Files —	the	ren	10176	?() [·]	F111	neti	on						162
8.3	Other	C Language F	unctions													162
0.0	831	Nonlocal Jun	Ins	•		• •				Ţ						162
	832	Set Time Zon		٠	•	• •	•	•	•	•	•	•	•	•	•	162
	0.0.2	Det Time 201		•	•	• •	•	•	•	•	•	•	•	•	•	102
Section	Q. Svet	om Databasos	•													165
0 1	System	n Databases	~ • • •	•	•	• •	•	•	•	•	•	•	•	•	•	165
9.1 0.9	Deteb		• • • •	٠	•	• •	•	•	٠	٠	•	٠	٠	•	•	166
9.4	Databa	ase Access	• • • •	•	•	• •	٠	٠	٠	•	•	•	•	•	•	166
	9.2.1	Group Datab	ase Access	5.	•	• •	٠	•	•	٠	•	•	٠	٠	•	100
	9.2.2	User Databa	se Access	٠	•	• •	•	•	٠	٠	•	٠	٠	٠	٠	107
Section	10. Do	ta Intorohana	o Format													169
10.1	10. Da	. A Interchang	File Format	• •	•	• •	•	•	•	•	•	•	٠	•	•	169
10.1		Enterchange	Flie Forn	al	•	• •	٠	•	•	٠	•	٠	•	•	•	160
	10.1.1	Extended ca	r rormat	•	•	• •	٠	•	•	٠	•	٠	٠	•	•	105
	10.1.2	Multiple Vel		ι.	•	• •	•	٠	•	•	•	•	•	•	•	177
	10.1.3	Multiple vol	imes .	•	•	• •	•	•	•	•	•	•	•	•	•	111
Annor	(infor	mativa) Bibli	ography						_	_						179
A 1		d Open System	ograpny na Standa	rda	•	• •	٠	•	•	•	•	•	•	•	•	179
A.I	Other	Standarda	ns Stanua	us	•	• •	•	•	•	•	•	•	•	•	•	181
A.Z	Other	Standards	· · · ·	• Trati	• nodi	• •	• ••• •	• •	• + a	•	•	•	•	•	•	182
A.3	Histor	ical Documen	tation and	Int	roa	ucto	ry.	rex	us	•	٠	•	•	•	•	102
Annow I	D (infor	mativa) Rati	nalo and	Note	96											185
D 1		and Normativ	o Roforon	2000	00	• •	•	•	•		•		Ţ			185
D.I	Scope	and Normativ	e Reference	inom	• ont	• •	•	•	•	•	•	•	•	•	•	196
B.Z	Denni	tions and Gen	eral nequ	Irem	lent	· 5	•	•	٠	•	٠	•	•	•	•	226
B.3	Proces	s Primitives	• • • •	•	•	• •	٠	•	•	•	٠	•	•	•	٠	246
B.4	Proces	s Environmer	it • • •	•	•	• •	•	•	•	•	•	•	•	•	•	240 959
B.5	Files a	and Directorie	S • • •	•	•	• •	•	•	•	•	٠	•	•	•	•	200
B.6	Input	and Output P	rimitives	•	•	• •	٠	•	•	٠	٠	•	٠	٠	•	204
B. 7	Device	- and Class-S	pecific Fui	nctio	ons	•	•	•	•	٠	•	•	•	•	•	273
B. 8	Langu	age-Specific S	ervices for	r the	CI	Prog	grai	nm	ing							000
	Langu	age · · ·	• • • •	٠	•	• •	•	٠	•	•	٠	٠	٠	٠	٠	200
B.9	Syster	n Databases		٠	•	• •	٠	•	•	•	٠	٠	٠	٠	٠	293
B.10	Data I	nterchange F	ormat .	٠	•	• •	•	•	•	٠	٠	٠	•	۰	٠	294
Annex	C (info	rmative) Head	ler Conter	nts S	Sam	ples	5	•	•	•	•	•	•	•	•	301

								PAGE
Annex D D.1 D D.2 O D.3 R D.4 R D.5 R	(info efin ptio elat elat elat	ormative) Profiles	•	•	• • • •	• • • •	•	313 313 314 315 315 315 315
Annex E (E.1 (E	(info Exar	ormative) Sample National Profile	•	•	•	•	•	317 318
Identifier	Inde	ex	٠	٠	•	•	٠	321
Alphabeti	с То	pical Index	•	•	•	٠	٠	327
TABLES								
Table 2-1	_	Primitive System Data Types	٠	•	•	•	•	27
Table 2-2	-	Reserved Header Symbols	•	٠	٠	•	٠	31
Table 2-3	_	Minimum Values	•	•	•	•	•	35
Table 2-4	-	Run-Time Increasable Values	•	•	•	•	•	35
Table 2-5 Indeter	_ min	Run-Time Invariant Values (Possibly ate)	•	•	•	•	٠	36
Table 2-6	_	Pathname Variable Values	•	•	•	٠	•	36
Table 2-7	_	Invariant Value	•	•	•	•	٠	37
Table 2-8	-	Symbolic Constants for the <i>access()</i> Function	•	•	•	•	•	38
Table 2-9	_	Symbolic Constants for the $lseek()$ Function .	•	•	•	•	•	38
Table 2-10) –	Compile-Time Symbolic Constants	•	•	•	•	•	38
Table 2-11	L —	Execution-Time Symbolic Constants	•	•	•	•	٠	39
Table 3-1	_	Required Signals	•	•	•	•	٠	52
Table 3-2	_	Job Control Signals	•	•	•	•	•	52
Table 4-1	_	uname() Structure Members	•	•	•	•	•	75
Table 4-2	_	Configurable System Variables						80

D 4 0

Table 5-1	-	<i>stat</i> Structure	•	•	•	•	•	101
Table 5-2	_	Configurable Pathname Variables	٠	٠	•	•	•	111
Table 6-1	_	cmd Values for $fcntl()$	•	•	•	•	•	122
Table 6-2	_	File Descriptor Flags Used for <i>fcntl</i> ()	•	٠	•	•	•	122
Table 6-3	_	<i>l_type</i> Values for Record Locking With <i>fcntl</i> ()	•	٠	•	•	•	122
Table 6-4	_	oflag Values for open()	•	•	•	•	•	122
Table 6-5	_	File Status Flags Used for open() and fcntl()	٠	٠	•	•	•	122
Table 6-6	_	File Access Modes Used for open() and fcntl()	•	٠	•	•	•	123
Table 6-7	-	Mask for Use With File Access Modes	•	٠	•	•	•	123
Table 6-8	_	<i>flock</i> Structure	•	•	•	•	•	125
Table 6-9	_	fcntl() Return Values	•	٠	٠	•	•	126
Table 7-1	-	termios Structure	•	•	•	•	•	136
Table 7-2	_	termios c_iflag Field	•	•	•	•	•	136
Table 7-3	-	termios c_cflag Field	•	•	•	•	•	138
Table 7-4	-	termios c_lflag Field	•	٠	٠	•	•	139
Table 7-5	-	termios c_cc Special Control Characters \cdot .	•	٠	٠	•	•	140
Table 7-6	_	termios Baud Rate Values	•	٠	٠	•	•	141
Table 9-1	_	group Structure	•	•	•	•	•	166
Table 9-2	_	passwd Structure	•	٠	٠	٠	•	167
Table 10-1	_	tar Header Block	•	•	•	•	•	170
Table 10-2	-	Byte-Oriented cpio Archive Entry	•	•	٠	٠	•	174
Table 10-3	_	Values for cpio c_mode Field	•	•	٠	٠	•	175
Table B-1	_	Suggested Feature Test Macros	•	•	•	•	•	222

Foreword

ISO (the International Organization for Standardization) and IEC (the Interna-1 tional Electrotechnical Commission) together form a system for worldwide stan-2 dardization as a whole. National bodies that are members of ISO or IEC partici-3 pate in the development of International Standards through technical committees Δ established by the respective organization to deal with particular fields of techni-5 cal activity. ISO and IEC technical committees collaborate in fields of mutual 6 interest. Other international organizations, governmental and nongovernmental, 7 in liaison with ISO and IEC, also take part in the work. 8

In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1. Draft International Standards adopted by the joint technical committee are circulated to national bodies for approval before their acceptance as International Standards. They are approved in accordance with procedures requiring at least 75% approval by the national bodies voting.

International Standard ISO/IEC 9945-1: 1990 was prepared by Joint Technical
 Committee ISO/IEC JTC 1, Information technology.

ISO/IEC 9945 consists of the following parts, under the general title Information
 technology—Portable operating system interface (POSIX):

- 18 Part 1: System application program interface (API) [C language]
- 19 Part 2: Shell and utilities (under development)
- 20 Part 3: System administration (under development)
- Annexes A to E of ISO/IEC 9945-1 are provided for information only.



International Organization for Standardization/International Electrotechnical Commission Case postale 56 • CH-1211 Genève 20 • Switzerland

Introduction

(This Introduction is not a normative part of ISO/IEC 9945-1 Information technology—Portable operating system interface (POSIX)—Part 1: System application programming interface (API) [C Language], but is included for information only.)

- The purpose of this part of ISO/IEC 9945 is to define a standard operating system interface and environment based on the UNIX¹⁾ Operating System documentation
- interface and environment based on the UNIX¹ Operating System documentation
 to support application portability at the source level. This is intended for systems
- 4 implementors and applications software developers.
- Initially,²⁾ the focus of this part of ISO/IEC 9945 is to provide standardized services via a C language interface. Future revisions are expected to contain bindings for other programming languages as well as for the C language. This will be accomplished by breaking this part of ISO/IEC 9945 into multiple portions—one defining core requirements independent of any programming language, and others composed of programming language bindings.
- The core requirements portion will define a set of required services common to 11 any programming language that can be reasonably expected to form a language 12 binding to this part of ISO/IEC 9945. These services will be described in terms of 13 functional requirements and will not define programming language-dependent 14 interfaces. Language bindings will consist of two major parts. One will contain 15 the programming language's standardized interface for accessing the core services 16 defined in the programming language-independent core requirements section of 17 this part of ISO/IEC 9945. The other will contain a standardized interface for 18 language-specific services. Any implementation claiming conformance to this part 19 of ISO/IEC 9945 with any language binding will be required to comply with both 20 sections of the language binding. 21
- Within this document, the term "POSIX.1" refers to this part of ISO/IEC 9945 itself.

24 **Organization of This Part of ISO/IEC 9945**

- 25 This part of ISO/IEC 9945 is divided into four elements:
- 26 (1) Statement of scope and list of normative references (Section 1)
- 27 (2) Definitions and global concepts (Section 2)
- 28 (3) The various interface facilities (Sections 3 through 9)
 - (4) Data interchange format (Section 10)

^{30 1)} UNIX is a registered trademark of UNIX System Laboratories in the U.S. and other countries.

The vertical rules in the right margin depict technical or significant non-editorial changes from
 IEEE Std 1003.1-1988 to IEEE Std 1003.1-1990. A vertical rule beside an empty line indicates
 deleted text.

Most of the sections describe a single service interface. The C Language binding 34 for the service interface is given in the subclause labeled Synopsis. The Descrip-35 tion subclause provides a specification of the operation performed by the service 36 interface. Some examples may be provided to illustrate the interfaces described. 37 In most cases there are also Returns and Errors subclauses specifying return 38 values and possible error conditions. References are used to direct the reader to 39 other related sections. Additional material to complement sections in this part of 40 ISO/IEC 9945 may be found in the Rationale and Notes, Annex B. This annex pro-41 vides historical perspectives into the technical choices made by the developers of 42 this part of ISO/IEC 9945. It also provides information to emphasize consequences 43 of the interfaces described in the corresponding section of this part of 44 ISO/IEC 9945. 45

Informative annexes are not part of the standard and are provided for information
only. (There is a type of annex called "normative" that is part of a standard and
imposes requirements, but there are currently no such normative annexes in this
part of ISO/IEC 9945.) They are provided for guidance and to help understanding.

In publishing this part of ISO/IEC 9945, its developers simply intend to provide a 50 yardstick against which various operating system implementations can be meas-51 ured for conformance. It is *not* the intent of the developers to measure or rate any 52 products, to reward or sanction any vendors of products for conformance or lack of 53 conformance to this part of ISO/IEC 9945, or to attempt to enforce this part of 54 ISO/IEC 9945 by these or any other means. The responsibility for determining the 55 degree of conformance or lack thereof with this part of ISO/IEC 9945 rests solely 56 with the individual who is evaluating the product claiming to be in conformance 57 with this part of ISO/IEC 9945. 58

59 Base Documents

The various interface facilities described herein are based on the 1984 /usr/group Standard derived and published by the UniForum (formerly /usr/group) Standards Committee. The 1984 /usr/group Standard and this part of ISO/IEC 9945 are largely based on UNIX Seventh Edition, UNIX System III, UNIX System V, 4.2BSD, and 4.3BSD documentation,³⁾ but wherever possible, compatibility with other systems derived from the UNIX operating system, or systems compatible with that system, has been maintained.

67 Background

The developers of POSIX.1 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. In the course of their deliberations, the developers reviewed related American and international standards, both published and in progress.

Conceptually, POSIX.1 describes a set of fundamental services needed for the
 efficient construction of application programs. Access to these services has been

I

I

³⁾ The IEEE is grateful to both AT&T and UniForum for permission to use their materials.

provided by defining an interface, using the C programming language, that establishes 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,⁴⁾ an acronym for Portable Operating System
Interface.

Although originated to refer to IEEE Std 1003.1-1988, the name POSIX more correctly refers to a *family* of related standards: IEEE 1003.*n* and the parts of International Standard 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.

87 Audience

The intended audience for ISO/IEC 9945 is all persons concerned with an industry-wide standard operating system based on the UNIX system. This includes at least four groups of people:

- 91 (1) Persons buying hardware and software systems;
- 92 (2) Persons managing companies that are deciding on future corporate com 93 puting directions;
- 94 (3) Persons implementing operating systems, and especially
 - (4) Persons developing applications where portability is an objective.

96 Purpose

95

97 Several principles guided the development of this part of ISO/IEC 9945:

98 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 part of ISO/IEC 9945 codifies the common, existing definition of the UNIX system. There was no attempt to define a new system interface.

105 Interface, Not Implementation

106This part of ISO/IEC 9945 defines an interface, not an implementation.107No distinction is made between library functions and system calls: both108are referred to as functions. No details of the implementation of any109function are given (although historical practice is sometimes indicated110in Annex B). Symbolic names are given for constants (such as signals111and error numbers) rather than numbers.

I

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.

115 Source, Not Object, Portability

This part of ISO/IEC 9945 has been written so that a program written 116 and translated for execution on one conforming implementation may 117 also be translated for execution on another conforming implementation. 118 This part of ISO/IEC 9945 does not guarantee that executable (object or 119 binary) code will execute under a different conforming implementation 120 than that for which it was translated, even if the underlying hardware 121 is identical. However, few impediments were placed in the way of 122 binary compatibility, and some remarks on this are found in Annex B. 123 See B.1.3.1.1 and B.4.8. 124

125 The C Language

126This part of ISO/IEC 9945 is written in terms of the standard C language127as specified in the C Standard {2}.⁵⁾ See B.1.3 and B.1.1.1.

128 No Super-User, No System Administration

There was no intention to specify all aspects of an operating system. System administration facilities and functions are excluded from POSIX.1, and functions usable only by the super-user have not been included. Annex B notes several such instances. Still, an implementation of the standard interface may also implement features not in this part of ISO/IEC 9945; see 1.3.1.1. This part of ISO/IEC 9945 is also not concerned with hardware constraints or system maintenance.

136 Minimal Interface, Minimally Defined

- In keeping with the historical design principles of the UNIX system, 137 POSIX.1 is as minimal as possible. For example, it usually specifies only 138 one set of functions to implement a capability. Exceptions were made in 139 some cases where long tradition and many existing applications 140 included certain functions, such as creat(). In such cases, as throughout 141 POSIX.1, redundant definitions were avoided: *creat()* is defined as a spe-142 cial case of open(). Redundant functions or implementations with less 143 tradition were excluded. 144
- 145 Broadly Implementable
- 146The developers of POSIX.1 endeavored to make all specified functions147implementable across a wide range of existing and potential systems,148including:
- (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

1

I

The number in braces corresponds to those of the references in 1.2 (or the bibliographic entry in
 Annex A if the number is preceded by the letter B).

155	(3) Emulations hosted on entirely different operating systems
156	(4) Networked systems
157	(5) Distributed systems
158	(6) Systems running on a broad range of hardware
159 160	No direct references to this goal appear in this part of ISO/IEC 9945, but some results of it are mentioned in Annex B.
161	Minimal Changes to Historical Implementations
162	There are no known historical implementations that will not have to
163	change in some area to conform to this part of ISO/IEC 9945, and in a
164	few areas POSIX.1 does not exactly match any existing system interface
165	(for example, see the discussion of O_NONBLOCK in B.6). Nonetheless,
166	there is a set of functions, types, definitions, and concepts that form an
167	interface that is common to most historical implementations. POSIX.1
168	specifies that common interface and extends it in areas where there has
169	historically been no consensus, preferably
170	(1) By standardizing an interface like one in an historical implemen-
171	tation; e.g., directories, or;
172	(2) By specifying an interface that is readily implementable in terms
173	of, and backwards compatible with, historical implementations,
174	such as the extended tar format in 10.1.1, or;
175	(3) By specifying an interface that, when added to an historical imple-
176	mentation, will not conflict with it, like B.6.
177	Required changes to historical implementations have been kept to a
178	minimum, but they do exist, and Annex B points out some of them.
179	POSIX.1 is specifically not a codification of a particular vendor's product.
180	It is similar to the UNIX system, but it is not identical to it.
181	It should be noted that implementations will have different kinds of
182	extensions. Some will reflect "historical usage" and will be preserved for
183	execution of pre-existing applications. These functions should be con-
184	sidered "obsolescent" and the standard functions used for new applica-
185	tions. Some extensions will represent functions beyond the scope of POSIX 1. These need to be used with coreful management to be able to
180	adapt to future POSIX 1 extensions and/or port to implementations that
188	provide these services in a different manner.
100	Minimal Changes to Existing Application Code
199	
190	A goal of POSIX.1 was to minimize additional work for the developers of
191	applications. However, because every known historical implementation
192	will nave to change at least slightly to conform, some applications will have to change. Appear B points out the major places where DOSIV 1
193	implies such changes
134	mipnes such changes.

195 Related Standards Activities

Activities to extend this part of ISO/IEC 9945 to address additional requirements are in progress, and similar efforts can be anticipated in the future.

The following areas are under active consideration at this time, or are expected to become active in the near future:⁶⁾

- 200 (1) Language-independent service descriptions of this part of ISO/IEC 9945
- 201 (2) C, Ada, and FORTRAN Language bindings to (1)
- 202 (3) Shell and Utility facilities
- 203 (4) Verification testing methods
- 204 (5) Realtime facilities
- 205 (6) Secure/Trusted System considerations
- 206 (7) Network interface facilities
- 207 (8) System Administration
- 208 (9) Graphical User Interfaces
- (10) Profiles describing application- or user-specific combinations of Open Systems standards for: supercomputing, multiprocessor, and batch extensions; transaction processing; realtime systems; and multiuser systems
 based on historical models
- (11) An overall guide to POSIX-based or related Open Systems standards and
 profiles
- Extensions are approved as "amendments" or "revisions" to this document, follow ing the IEEE and ISO/IEC Procedures.
- Approved amendments are published separately until the full document is reprinted and such amendments are incorporated in their proper positions.

If you have interest in participating in the TCOS working groups addressing these issues, please send your name, address, and phone number to the Secretary, IEEE Standards Board, Institute of Electrical and Electronics Engineers, Inc., P.O. Box 1331, 445 Hoes Lane, Piscataway, NJ 08855-1331, and ask to have this forwarded to the chairperson of the appropriate TCOS working group. If you have interest in participating in this work at the international level, contact your ISO/IEC national body.

⁶⁾ A Standards Status Report that lists all current IEEE Computer Society standards projects is available from the IEEE Computer Society, 1730 Massachusetts Avenue NW, Washington, DC 20036-1903; Telephone: +1 202 371-0101; FAX: +1 202 728-9614. Working drafts of POSIX standards under development are also available from this office.

IEEE Std 1003.1-1990 was prepared by the 1003.1 Working Group, sponsored by the Technical Committee on Operating Systems and Application Environments of the IEEE Computer Society. At the time this standard was approved, the membership of the 1003.1 Working Group was as follows:

Technical Committee on Operating Systems and Application Environments (TCOS)

Chair: Luis-Felipe Cabrera

Standards Subcommittee for TCOS

Chair:	Jim Isaak
Treasurer:	Quin Hahn
Secretary:	Shane McCarron

1003.1 Working Group Officials

Chair:	Donn Terry
Vice Chair:	Keith Stuck
Editor:	Hal Jespersen
Secretary:	Keith Stuck

Working Group

Greg Goddard Andrew Griffith Rand Hoven Randall Howard Mike Karels Jeff Kimmel David Korn Bob Lenk Shane McCarron John Meyer Martha Nalebuff Neguine Navab Paul Rabin Seth Rosenthal Lorne Schachter Steve Schwarm Paul Shaughnessy Steve Sommars Ravi Tavakley Jeff Tofano David Willcox John Wu

Steve Bartels Robert Bismuth James Bohem Kathy Bohrer Keith Bostic Jonathan Brown Tim Carter Myles Connors Landon Curt Noll Dave Decot Mark Doran Glenn Fowler The following persons were members of the 1003.1 Balloting Group that approved the standard for submission to the IEEE Standards Board:

	David Chinn Michael Lambert	Open Software Foundation Instituti X/Open Institutional Representativ	ional Representative e
	Shane McCarron	UNIX International Institutional Re	tive presentative
Helene Arm	itage	William Henderson	Martha Nalebuff
David Ather	sych	Lee A. Hollaar	Barry Needham
Timothy Bal	ker	Terrence Holm	Alan F. Nugent
Geoff Baldw	in	Randall Howard	Jim Oldroyd
Steven E. Ba	arber	Irene Hu	Craig Partridge
Robert Barn	ned	Andrew Huber	John Peace
John Barr		Richard Hughes-Rowlands	John C. Penney
James Bohe	m	Judith Hurwitz	P. Plauger
Kathryn Bol	hrer	Jim Isaak	Gerald Powell
Robert Boro	choff	Dan Iuster	Scott E. Preece
Keith Bostic	2	Richard James	Joseph Ramus
James P. Bo	ound	Hal Jespersen	Wendy Rauch
Joseph Boyk	tin	Michael J. Karels	Carol Rave
Kevin Brady	V	Sol Kavy	Wavne B. Reed
Phyllis Eve	Bregman	Lorraine C. Kevra	Christopher J. Riddick
Fred Lee Br	own. Jr.	Jeffrey S. Kimmel	Andrew K. Roach
A. Winsor B	rown	M. J. Kirk	Robert Sarr
Luis-Felipe	Cabrera	Dale Kirkland	Lorne H. Schachter
Nicholas A.	Camillone	John T. Kline	Norman Schneidewind
Clvde Camp)	Kenneth Klingman	Stephen Schwarm
John Carsor	ı	Joshua Knight	Richard Scott
Steven Cart	- er	Andrew R. Knipp	Leonard Seagren
Jerry Cashi	n	David Korn	Glen Seeds
Kilnam Cho	n	Don Kretsch	Karen Sheaffer
Anthony Cir	ncotta	Takahiko Kuki	Charles Smith
Mark Colbu	rn	Thomas Kwan	Steven Sommars
Donald W. (Cragun	Robin B. Lake	Douglas H. Steves
Ana Maria I	DeAlvare	Mark Lamonds	James Tanner
Dave Decot		Doris Lebovits	Ravi Tavaklev
Steven Delle	er	Maggie Lee	Marc Teitelbaum
Terence Dov	vling	Greger Leijonhufvud	Donn S. Terry
Stephen A	Dum	Robert Lenk	Gary F Tom
John D. Ear	ls	David Lennert	Andrew Twigger
Ron Elliott	10	Donald Lewine	Mark-Rene Uchida
David Emer	v	Kevin Lewis	L David Umbaugh
Philip H. Er	s Islow	F C Lim	Michael W. Vannier
Ken Faubel		James Loniers	David John Wallace
Kester Fong	r	Warren E. Loner	Stephen Walli
Kenneth R	, Gibb	Roger Martin	Larry Wehr
Michel Gien		Martin J. McGowan	Bruce Weiner
Gregory W.	Goddard	Marshall McKusick	Robert Weissensee
Dave Grinde	eland	Robert McWhirter	P.J. Weyman
Judy Guist	orana	Paul Merry	Andrew Wheeler Jr
James Hall		Doug Michels	David Willcox
Charles Har	mmons	Gary W Miller	Randall F Wright
Allen Hanki	inson	James Moo	Oren Vien
Steve Head		James W Moore	Jason Zions
Barry Hedq	uist	Sames W. MOULE	045011 210115

When the IEEE Standards Board approved this standard on September 28, 1990, it had the following membership:

Marco W. Migliaro, Chairman Andrew G. Salem, Secretary James M. Daly, Vice Chairman

Dennis Bodson Paul L. Borrill Fletcher J. Buckley Allen L. Clapp Stephen R. Dillon Donald C. Fleckenstein Jay Forster* Thomas L. Hannan

Kenneth D. Hendrix John W. Horch Joseph L. Koepfinger* Irving Kolodny Michael A. Lawler Donald J. Loughry John E. May, Jr. Lawrence V. McCall L. Bruce McClung Donald T. Michael* Stig Nilsson Roy T. Oishi Gary S. Robinson Terrance R. Whittemore Donald W. Zipse

*Member Emeritus



Information technology—Portable operating system interface (POSIX)—Part 1: System application programming interface (API) [C Language]

Section 1: General

1 **1.1 Scope**

6

7

9

12

2 This part of ISO/IEC 9945 defines a standard operating system interface and 3 environment to support application portability at the source-code level. It is 4 intended to be used by both application developers and system implementors.

- 5 This part of ISO/IEC 9945 comprises four major components:
 - (1) Terminology, concepts, and definitions and specifications that govern structures, headers, environment variables, and related requirements
- 8 (2) Definitions for system service interfaces and subroutines
 - (3) Language-specific system services for the C programming language
- 10 (4) Interface issues, including portability, error handling, and error recovery
- 11 The following areas are outside of the scope of this part of ISO/IEC 9945:
 - (1) User interface (shell) and associated commands
- 13 (2) Networking protocols and system call interfaces to those protocols
- 14 (3) Graphics interfaces
- 15 (4) Database management system interfaces
- 16 (5) Record I/O considerations

18

19

20

- (6) Object or binary code portability
 - (7) System configuration and resource availability
 - (8) The behavior of system services on systems supporting concurrency within a single process

This part of ISO/IEC 9945 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.

This part of ISO/IEC 9945 has been defined exclusively at the source-code level. The objective is that a Strictly Conforming POSIX.1 Application source program can be translated to execute on a conforming implementation.

29 **1.2 Normative References**

The following standards contain provisions which, through references in this text, constitute provisions of this part of ISO/IEC 9945. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this part of this International Standard are encouraged to investigate the possibility of applying the most recent editions of the standards listed below. Members of IEC and ISO maintain registers of currently valid International Standards.

- ISO/IEC 646: 1983,¹⁾ Information processing—ISO 7-bit coded character set
 for information interchange.
- ³⁹ {2} ISO/IEC 9899: ...,²⁾ Information technology—Programming languages—C.
- 40 **1.3 Conformance**
- 41 **1.3.1 Implementation Conformance**
- 42 **1.3.1.1 Requirements**
- ⁴³ A *conforming implementation* shall meet all of the following criteria:

 ^{44 1)} Under revision. (This notation is meant to explicitly reference the 1990 Draft International
 45 Standard version of ISO/IEC 646.)

ISO/IEC documents can be obtained from the ISO office, 1, rue de Varembé, Case Postale 56, CH 1211, Genève 20, Switzerland/Suisse.

^{48 2)} To be approved and published.

1

- (1) The system shall support all required interfaces defined within this part
 of ISO/IEC 9945. These interfaces shall support the functional behavior
 described herein.
- (2) The system may provide additional functions or facilities not required by 52this part of ISO/IEC 9945. Nonstandard extensions should be identified 53 as such in the system documentation. Nonstandard extensions, when 54 used, may change the behavior of functions or facilities defined by this 55 part of ISO/IEC 9945. The conformance document shall define an 56 environment in which an application can be run with the behavior 57 specified by the standard. In no case shall such an environment require 58 modification of a Strictly Conforming POSIX.1 Application. 59

60 **1.3.1.2 Documentation**

A conformance document with the following information shall be available for an implementation claiming conformance to this part of ISO/IEC 9945. The conformance document shall have the same structure as this part of ISO/IEC 9945, with the information presented in the appropriately numbered sections, clauses, and subclauses. The conformance document shall not contain information about extended facilities or capabilities outside the scope of this part of ISO/IEC 9945.

The conformance document shall contain a statement that indicates the full name, number, and date of the standard that applies. The conformance document may also list international software standards that are available for use by a Conforming POSIX.1 Application. Applicable characteristics where documentation is required by one of these standards, or by standards of government bodies, may also be included.

The conformance document shall describe the behavior of the implementation for all implementation-defined features defined in this part of ISO/IEC 9945. This requirement shall be met by listing these features and providing either a specific reference to the system documentation or providing full syntax and semantics of these features. The conformance document may specify the behavior of the implementation for those features where this part of ISO/IEC 9945 states that implementations may vary or where features are identified as undefined or unspecified.

- No specifications other than those described in this part of ISO/IEC 9945 shall be
 present in the conformance document.
- The phrases "shall document" or "shall be documented" in this part of ISO/IEC 9945 mean that documentation of the feature shall appear in the conformance document, as described previously, unless the system documentation is explicitly mentioned.
- The system documentation should also contain the information found in the conformance document.

91 **1.3.1.3 Conforming Implementation Options**

The following symbolic constants, described in the subclauses indicated, reflect implementation options for this part of ISO/IEC 9945 that could warrant requirement by Conforming POSIX.1 Applications, or in specifications of conforming systems, or both:

96	{NGROUPS_MAX}	Multiple groups option (in 2.8.3)
97	{_POSIX_JOB_CONTROL}	Job control option (in 2.9.3)
98	{_POSIX_CHOWN_RESTRICTED}	Administrative/security option (in 2.9.4)

The remaining symbolic constants in 2.9.3 and 2.9.4 are useful for testing purposes and as a guide to applications on the types of behaviors they need to be able to accommodate. They do not reflect sufficient functional difference to warrant requirement by Conforming POSIX.1 Applications or to distinguish between conforming implementations.

In the cases where omission of an option would cause functions described by this part of ISO/IEC 9945 to not be defined, an implementation shall provide a function that is callable with the syntax defined in this part of ISO/IEC 9945, even though in an instance of the implementation the function may always do nothing but return an error.

109 **1.3.2 Application Conformance**

All applications claiming conformance to this part of ISO/IEC 9945 shall use only language-dependent services for the C programming language described in 1.3.3

and shall fall within one of the following categories:

113 **1.3.2.1 Strictly Conforming POSIX.1 Application**

A Strictly Conforming POSIX.1 Application is an application that requires only the facilities described in this part of ISO/IEC 9945 and the applicable language standards. Such an application shall accept any behavior described in this part of ISO/IEC 9945 as *unspecified* or *implementation-defined*, and for symbolic constants, shall accept any value in the range permitted by this part of ISO/IEC 9945. Such applications are permitted to adapt to the availability of facilities whose availability is indicated by the constants in 2.8 and 2.9.

121 **1.3.2.2 Conforming POSIX.1 Application**

122 **1.3.2.2.1 ISO/IEC Conforming POSIX.1 Application**

An ISO/IEC Conforming POSIX.1 Application is an application that uses only the facilities described in this part of ISO/IEC 9945 and approved Conforming Language bindings for any ISO or IEC standard. Such an application shall include a statement of conformance that documents all options and limit dependencies, and all other ISO or IEC standards used.

128 1.3.2.2.2 <National Body> Conforming POSIX.1 Application

A <National Body> Conforming POSIX.1 Application differs from an ISO/IEC Conforming POSIX.1 Application in that it also may use specific standards of a single ISO/IEC member body referred to here as "<National Body>." Such an application shall include a statement of conformance that documents all options and limit dependencies, and all other <National Body> standards used.

134 **1.3.2.3 Conforming POSIX.1 Application Using Extensions**

A Conforming POSIX.1 Application Using Extensions is an application that differs 135 from a Conforming POSIX.1 Application only in that it uses nonstandard facilities 136 that are consistent with this part of ISO/IEC 9945. Such an application shall fully 137 document its requirements for these extended facilities, in addition to the docu-138 mentation required of a Conforming POSIX.1 Application. A Conforming POSIX.1 139 Application Using Extensions shall be either an ISO/IEC Conforming POSIX.1 140 Application Using Extensions or a <National Body> Conforming POSIX.1 Applica-141 tion Using Extensions (see 1.3.2.2.1 and 1.3.2.2.2). 142

143 **1.3.3 Language-Dependent Services for the C Programming Language**

Parts of ISO/IEC 9899 [2] (hereinafter referred to as the "C Standard [2]") will be 144 referenced to describe requirements also mandated by this part of ISO/IEC 9945. 145 The sections of the C Standard {2} referenced to describe requirements for this 146 part of ISO/IEC 9945 are specified in Section 8. Section 8 also sets forth additions 147 and amplifications to the referenced sections of the C Standard {2}. Any imple-148 mentation claiming conformance to this part of ISO/IEC 9945 with the C Language 149 Binding shall provide the facilities referenced in Section 8, along with any addi-150 tions and amplifications Section 8 requires. 151

Although this part of ISO/IEC 9945 references parts of the C Standard {2} to describe some of its own requirements, conformance to the C Standard {2} is unnecessary for conformance to this part of ISO/IEC 9945. Any C language implementation providing the facilities stipulated in Section 8 may claim conformance; however, it shall clearly state that its C language does not conform to the C Standard {2}.

158 **1.3.3.1 Types of Conformance**

Implementations claiming conformance to this part of ISO/IEC 9945 with the C Language Binding shall claim one of two types of conformance—conformance to POSIX.1, C Language Binding (C Standard Language-Dependent System Support), or to POSIX.1, C Language Binding (Common-Usage C Language-Dependent System Support).

164 **1.3.3.2 C Standard Language-Dependent System Support**

Implementors shall meet the requirements of Section 8 using for reference the
 C Standard {2}. Implementors shall clearly document the version of the
 C Standard {2} referenced in fulfilling the requirements of Section 8.

1

1

I

I

Implementors seeking to claim conformance using the C Standard {2} shall claim
 conformance to POSIX.1, C Language Binding (C Standard Language-Dependent
 System Support).

171 **1.3.3.3 Common-Usage C Language-Dependent System Support**

Implementors, instead of referencing the C Standard {2}, shall provide the rou-172 tines and support required in Section 8 using common usage as guidance. Imple-173 mentors shall meet all the requirements of Section 8 except where references are 174 made to the C Standard {2}. In places where the C Standard {2} is referenced, 175 implementors shall provide equivalent support in a manner consistent with com-176 mon usage of the C programming language. Implementors shall document, in 177 Section 8 of the conformance document, all differences between the interface pro-178 vided and the interface that would have been provided had the C Standard {2} 179 been implemented instead of common usage. Implementors shall clearly docu-180 ment the version of the C Standard {2} referenced in documenting interface differ-181 ences and should issue updates on differences for all new versions of the 182 C Standard {2}. 183

184

Where a function has been introduced by the C Standard {2}, and thus there is no common-usage referent for it, if the function is implemented, it shall be implemented as described in the C Standard {2}. If the function is not implemented, it shall be documented as a difference from the C Standard {2} as required above.

189 **1.3.4 Other C Language-Related Specifications**

The following rules apply to the usage of C language library functions; each of the statements in this subclause applies to the detailed function descriptions in Sections 3 through 9, unless explicitly stated otherwise:

- (1) If an argument to a function has an invalid value (such as a value outside the domain of the function, or a pointer outside the address space of the program, or a NULL pointer when that is not explicitly permitted), the behavior is undefined.
- (2) Any function may also be implemented as a macro in a header. Applica tions should use #undef to remove any macro definition and ensure that
 an actual function is referenced. Applications should also use #undef
 prior to declaring any function in this part of ISO/IEC 9945.
- (3) Any invocation of a library function that is implemented as a macro shall
 expand to code that evaluates each of its arguments only once, fully pro tected by parentheses where necessary, so it is generally safe to use arbi trary expressions as arguments.
- (4) Provided that a library function can be declared without reference to any
 type defined in a header, it is also permissible to declare the function,
 either explicitly or implicitly, and use it without including its associated
 header.

1

1

(5) If a function that accepts a variable number of arguments is not declared
 (explicitly or by including its associated header), the behavior is
 undefined.

1.3.5 Other Language-Related Specifications

This part of ISO/IEC 9945 is currently specified in terms of the language defined by the C Standard {2}. Bindings to other programming languages are being developed.

If conformance to this part of ISO/IEC 9945 is claimed for implementation of any 216 programming language, the implementation of that language shall support the 217 use of external symbols distinct to at least 31 bytes in length in the source pro-218 gram text. (That is, identifiers that differ at or before the thirty-first byte shall be 219 distinct.) If a national or international standard governing a language defines a 220 maximum length that is less than this value, the language-defined maximum 221 222 shall be supported. External symbols that differ only by case shall be distinct when the character set in use distinguishes upper- and lowercase characters and 223 the language permits (or requires) upper- and lowercase characters to be distinct 224 in external symbols. 225

Subsequent sections of this part of ISO/IEC 9945 refer only to the C Language.



I

I

Section 2: Terminology and General Requirements

1	2.1 Co	nventions
2	This par	rt of ISO/IEC 9945 uses the following typographic conventions:
3	(1)	The <i>italic</i> font is used for:
4 5 6		- Cross references to defined terms within 1.3, 2.2.1, and 2.2.2; symbolic parameters that are generally substituted with real values by the application
7 8		 C language data types and function names (except in function Synopsis subclauses)
9		— Global external variable names
10	(2)	The bold font is used with a word in all capital letters, such as
11		PATH
12 13		to represent an environment variable, as described in 2.6. It is also used for the term "NULL pointer."
14	(3)	The constant-width (Courier) font is used:
15 16		 For C language data types and function names within function Synopsis subclauses
17 18		 To illustrate examples of system input or output where exact usage is depicted
19		— For references to utility names and C language headers
20 21	(4)	Symbolic constants returned by many functions as error numbers are represented as:
22		[ERRNO]
23		See 2.4.
24 25	(5)	Symbolic constants or limits defined in certain headers are represented as:
26		{LIMIT}
27		See 2.8 and 2.9.
28	In some	cases tabular information is presented "inline"; in others it is presented

In some cases tabular information is presented "inline"; in others it is presented in a separately labeled table. This arrangement was employed purely for ease of typesetting and there is no normative difference between these two cases.

1

- The conventions listed previously are for ease of reading only. Editorial inconsistencies in the use of typography are unintentional and have no normative meaning in this part of ISO/IEC 9945.
- NOTEs provided as parts of labeled tables and figures are integral parts of this part of ISO/IEC 9945 (normative). Footnotes and notes within the body of the text are for information only (informative).
- Numerical quantities are presented in international style: comma is used as a decimal sign and units are from the International System (SI).

39 **2.2 Definitions**

40 **2.2.1 Terminology**

- 41 For the purposes of this part of ISO/IEC 9945, the following definitions apply:
- 2.2.1.1 conformance document: A document provided by an implementor that
 contains implementation details as described in 1.3.1.2.

2.2.1.2 implementation defined: An indication that the implementation shall
 define and document the requirements for correct program constructs and correct
 data of a value or behavior.

47 **2.2.1.3 may:** An indication of an optional feature.

With respect to implementations, the word *may* is to be interpreted as an optional feature that is not required in this part of ISO/IEC 9945, but can be provided. With respect to Strictly Conforming POSIX.1 Applications, the word *may* means that the optional feature shall not be used.

- 52 **2.2.1.4 obsolescent:** An indication that a certain feature may be considered for 53 withdrawal in future revisions of this part of ISO/IEC 9945.
- 54 Obsolescent features are retained in this version because of their widespread use. 55 Their use in new applications is discouraged.
- 2.2.1.5 shall: An indication of a requirement on the implementation or on
 Strictly Conforming POSIX.1 Applications, where appropriate.

58 **2.2.1.6 should:**

- (1) With respect to implementations, an indication of an implementation recommendation, but not a requirement.
- (2) With respect to applications, an indication of a recommended program ming practice for applications and a requirement for Strictly Conforming
 POSIX.1 Applications.

59

1

1

l

I

I

64 **2.2.1.7 supported:** A condition regarding optional functionality.

65 Certain functionality in this part of ISO/IEC 9945 is optional, but the interfaces to 66 that functionality are always required. If the functionality is *supported*, the 67 interfaces work as specified by this part of ISO/IEC 9945 (except that they do not 68 return the error condition indicated for the unsupported case). If the functional-69 ity is not *supported*, the interface shall always return the indication specified for 67 this situation.

- 71 **2.2.1.8 system documentation:** All documentation provided with an imple-72 mentation, except the conformance document.
- Electronically distributed documents for an implementation are considered part ofthe system documentation.

2.2.1.9 undefined: An indication that this part of ISO/IEC 9945 imposes no portability requirements on an application's use of an indeterminate value or its behavior with erroneous program constructs or erroneous data.

- Implementations (or other standards) may specify the result of using that value or
 causing that behavior. An application using such behaviors is using extensions,
 as defined in 1.3.2.3.
- 81 **2.2.1.10 unspecified:** An indication that this part of ISO/IEC 9945 imposes no 82 portability requirements on applications for correct program constructs or correct 83 data regarding a value or behavior.
- Implementations (or other standards) may specify the result of using that value or
 causing that behavior. An application requiring a specific behavior, rather than
 tolerating any behavior when using that functionality, is using extensions, as
 defined in 1.3.2.3.
- 88 2.2.2 General Terms
- 89 For the purposes of this part of ISO/IEC 9945, the following definitions apply:
- 90 **2.2.2.1 absolute pathname:** See *pathname resolution* in 2.3.6.
- 91 **2.2.2.2 access mode:** A form of access permitted to a file.
- 92 2.2.2.3 address space: The memory locations that can be referenced by a
 93 process.
- 2.2.2.4 appropriate privileges: An implementation-defined means of associat ing privileges with a process with regard to the function calls and function call
 options defined in this part of ISO/IEC 9945 that need special privileges.
- 97 There may be zero or more such means.

2.2.2.5 background process: A process that is a member of a background process group.

2.2.2.6 background process group: Any process group, other than a fore ground process group, that is a member of a session that has established a con nection with a controlling terminal.

103 **2.2.2.7 block special file:** A file that refers to a device.

104 A block special file is normally distinguished from a character special file by pro-105 viding access to the device in a manner such that the hardware characteristics of 106 the device are not visible.

107 **2.2.2.8 character:** A sequence of one or more bytes representing a single 108 graphic symbol.

109 NOTE: This term corresponds in the C Standard {2} to the term *multibyte character*, noting that a

single-byte character is a special case of multibyte character. Unlike the usage in the C Standard (2), character here has no necessary relationship with storage space, and byte is used when storage

- 112 space is discussed.
- 113 **2.2.2.9 character special file:** A file that refers to a device.

114 One specific type of character special file is a terminal device file, whose access is

defined in 7.1. Other character special files have no structure defined by this part

of ISO/IEC 9945, and their use is unspecified by this part of ISO/IEC 9945.

- 117 **2.2.2.10 child process:** See *process* in 2.2.2.62.
- 118 **2.2.2.11 clock tick:** An interval of time.

119 A number of these occur each second. Clock ticks are one of the units that may be 120 used to express a value found in type $clock_t$.

121 **2.2.2.12 controlling process:** The session leader that established the connec-122 tion to the controlling terminal.

123 Should the terminal subsequently cease to be a controlling terminal for this ses-124 sion, the session leader shall cease to be the controlling process.

125 **2.2.2.13 controlling terminal:** A terminal that is associated with a session.

Each session may have at most one controlling terminal associated with it, and a controlling terminal is associated with exactly one session. Certain input sequences from the controlling terminal (see 7.1) cause signals to be sent to all processes in the process group associated with the controlling terminal.

130 **2.2.2.14 current working directory:** See *working directory* in 2.2.2.89.

- 131 2.2.2.15 device: A computer peripheral or an object that appears to the applica 132 tion as such.
- 133 **2.2.2.16 directory:** A file that contains directory entries.
- 134 No two directory entries in the same directory shall have the same name.
- 135 **2.2.2.17 directory entry [link]:** An object that associates a filename with a file.
- 136 Several directory entries can associate names with the same file.
- 137 **2.2.2.18 dot:** The filename consisting of a single dot character (.).
- 138 See *pathname resolution* in 2.3.6.
- 139 **2.2.2.19 dot-dot:** The filename consisting solely of two dot characters (...).
- 140 See *pathname resolution* in 2.3.6.
- 2.2.2.20 effective group ID: An attribute of a process that is used in determin ing various permissions, including file access permissions, described in 2.3.2.
- See group ID. This value is subject to change during the process lifetime, as
 described in 3.1.2 and 4.2.2.
- 2.2.2.21 effective user ID: An attribute of a process that is used in determining
 various permissions, including file access permissions.
- See user ID. This value is subject to change during the process lifetime, as
 described in 3.1.2 and 4.2.2.
- **2.2.2.22 empty directory:** A directory that contains, at most, directory entries
 for dot and dot-dot.
- 151 2.2.2.3 empty string [null string]: A character array whose first element is a
 152 null character.
- 2.2.2.24 Epoch: The time 0 hours, 0 minutes, 0 seconds, January 1, 1970, Coor dinated Universal Time.
- 155 See seconds since the Epoch.
- 2.2.2.25 feature test macro: A #defined symbol used to determine whether a
 particular set of features will be included from a header.
- 158 See 2.7.1.

2.2.2.26 FIFO special file [FIFO]: A type of file with the property that data
 written to such a file is read on a first-in-first-out basis.

- Other characteristics of *FIFO*s are described in 5.3.1, 6.4.1, 6.4.2, and 6.5.3.
- 162 **2.2.2.27 file:** An object that can be written to, or read from, or both.

A file has certain attributes, including access permissions and type. File types include regular file, character special file, block special file, FIFO special file, and directory. Other types of files may be defined by the implementation.

166 **2.2.2.28 file description:** See open file description in 2.2.2.51.

167 2.2.2.29 file descriptor: A per-process unique, nonnegative integer used to
 168 identify an open file for the purpose of file access.

169 2.2.2.30 file group class: The property of a file indicating access permissions
 170 for a process related to the process's group identification.

A process is in the file group class of a file if the process is not in the file owner class and if the effective group ID or one of the supplementary group IDs of the process matches the group ID associated with the file. Other members of the class may be implementation defined.

175 **2.2.2.31 file mode:** An object containing the file permission bits and other 176 characteristics of a file, as described in 5.6.1.

177 2.2.2.32 filename: A name consisting of 1 to {NAME_MAX} bytes used to name a178 file.

The characters composing the name may be selected from the set of all character values excluding the slash character and the null character. The filenames dot and dot-dot have special meaning; see *pathname resolution* in 2.3.6. A filename is sometimes referred to as a pathname component.

183 2.2.2.33 file offset: The byte position in the file where the next I/O operation
 184 begins.

Each open file description associated with a regular file, block special file, or directory has a file offset. A character special file that does not refer to a terminal device may have a file offset. There is no file offset specified for a pipe or FIFO.

- 2.2.2.34 file other class: The property of a file indicating access permissions for
 a process related to the process's user and group identification.
- A process is in the file other class of a file if the process is not in the file owner class or file group class.
- 2.2.2.35 file owner class: The property of a file indicating access permissions
 for a process related to the process's user identification.
- A process is in the file owner class of a file if the effective user ID of the process
 matches the user ID of the file.
2.2.2.36 file permission bits: Information about a file that is used, along with
 other information, to determine if a process has read, write, or execute/search per mission to a file.

The bits are divided into three parts: owner, group, and other. Each part is used with the corresponding file class of processes. These bits are contained in the file mode, as described in 5.6.1. The detailed usage of the file permission bits in access decisions is described in *file access permissions* in 2.3.2.

- 203 **2.2.2.37 file serial number:** A per-file system unique identifier for a file.
- File serial numbers are unique throughout a file system.
- 205 **2.2.2.38 file system:** A collection of files and certain of their attributes.
- 206 It provides a name space for file serial numbers referring to those files.
- 207 **2.2.2.39 foreground process:** A process that is a member of a foreground pro-208 cess group.

209 **2.2.2.40 foreground process group:** A process group whose member processes 210 have certain privileges, denied to processes in background process groups, when 211 accessing their controlling terminal.

- Each session that has established a connection with a controlling terminal has exactly one process group of the session as the foreground process group of that controlling terminal. See 7.1.1.4.
- 215 2.2.2.41 foreground process group ID: The process group ID of the foreground
 216 process group.
- 217 **2.2.2.42 group ID:** A nonnegative integer, which can be contained in an object of 218 type *gid_t*, that is used to identify a group of system users.
- Each system user is a member of at least one group. When the identity of a group is associated with a process, a group ID value is referred to as a real group ID, an effective group ID, one of the (optional) supplementary group IDs, or an (optional) saved set-group-ID.
- 223 **2.2.2.43 job control:** A facility that allows users to selectively stop (suspend) 224 the execution of processes and continue (resume) their execution at a later point.
- The user typically employs this facility via the interactive interface jointly supplied by the terminal I/O driver and a command interpreter. Conforming implementations may optionally support job control facilities; the presence of this option is indicated to the application at compile time or run time by the definition of the {_POSIX_JOB_CONTROL} symbol; see 2.9.
- 230 **2.2.2.44 link:** See *directory entry* in 2.2.2.17.

I

- 231 **2.2.2.45 link count:** The number of directory entries that refer to a particular 232 file.
- 233 **2.2.2.46 login:** The unspecified activity by which a user gains access to the 234 system.
- Each login shall be associated with exactly one login name.
- 236 **2.2.2.47 login name:** A user name that is associated with a login.
- 237 2.2.2.48 mode: A collection of attributes that specifies a file's type and its access
 238 permissions.
- 239 See *file access permissions* in 2.3.2.
- 240 **2.2.2.49 null string:** See *empty string* in 2.2.2.23.
- 241 **2.2.2.50 open file:** A file that is currently associated with a file descriptor.
- 242 2.2.51 open file description: A record of how a process or group of processes
 243 are accessing a file.
- Each file descriptor shall refer to exactly one open file description, but an open file description may be referred to by more than one file descriptor. A file offset, file status (see Table 6-5), and file access modes (see Table 6-6) are attributes of an open file description.
- 248 **2.2.2.52 orphaned process group:** A process group in which the parent of 249 every member is either itself a member of the group or is not a member of the 250 group's session.
- 251 **2.2.2.53 parent directory:**
- (1) When discussing a given directory, the directory that both contains a
 directory entry for the given directory and is represented by the path name dot-dot in the given directory.
- (2) When discussing other types of files, a directory containing a directory entry for the file under discussion.
- ²⁵⁷ This concept does not apply to dot and dot-dot.
- 258 **2.2.2.54 parent process:** See *process* in 2.2.2.62.
- 259 **2.2.2.55 parent process ID:** An attribute of a new process after it is created by 260 a currently active process.

The parent process ID of a process is the process ID of its creator, for the lifetime of the creator. After the creator's lifetime has ended, the parent process ID is the process ID of an implementation-defined system process.

- 264 **2.2.2.56 path prefix:** A pathname, with an optional ending slash, that refers to 265 a directory.
- 266 **2.2.2.57 pathname:** A string that is used to identify a file.

A pathname consists of, at most, {PATH_MAX} bytes, including the terminating 267null character. It has an optional beginning slash, followed by zero or more 268 filenames separated by slashes. If the pathname refers to a directory, it may also 269 have one or more trailing slashes. Multiple successive slashes are considered to 270 be the same as one slash. A pathname that begins with two successive slashes 271 may be interpreted in an implementation-defined manner, although more than 272two leading slashes shall be treated as a single slash. The interpretation of the 273 pathname is described in 2.3.6. 274

- 275 **2.2.2.58 pathname component:** See *filename* in 2.2.2.32.
- 276 **2.2.2.59 pipe:** An object accessed by one of the pair of file descriptors created by the *pipe()* function.
- Once created, the file descriptors can be used to manipulate it, and it behaves identically to a FIFO special file when accessed in this way. It has no name in the file hierarchy.
- 281 **2.2.2.60 portable filename character set:** The set of characters from which 282 portable filenames are constructed.
- For a filename to be portable across conforming implementations of this part of ISO/IEC 9945, it shall consist only of the following characters:
- 285
 A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

 286
 a b c d e f g h i j k l m n o p q r s t u v w x y z

 287
 0 1 2 3 4 5 6 7 8 9 . _

The last three characters are the period, underscore, and hyphen characters, respectively. The hyphen shall not be used as the first character of a portable filename. Upper- and lowercase letters shall retain their unique identities between conforming implementations. In the case of a portable pathname, the slash character may also be used.

- 293 **2.2.2.61 privilege:** See appropriate privileges in 2.2.2.4.
- 294 **2.2.2.62 process:** An address space and single thread of control that executes 295 within that address space, and its required system resources.
- A process is created by another process issuing the fork() function. The process that issues fork() is known as the parent process, and the new process created by the fork() is known as the child process.
- 299 **2.2.2.63 process group:** A collection of processes that permits the signaling of 300 related processes.

I

- Each process in the system is a member of a process group that is identified by a process group ID. A newly created process joins the process group of its creator.
- 2.2.2.64 process group ID: The unique identifier representing a process group
 during its lifetime.
- A process group ID is a positive integer that can be contained in a pid_t . It shall | not be reused by the system until the process group lifetime ends.
- 307 2.2.2.65 process group leader: A process whose process ID is the same as its
 308 process group ID.

2.2.2.66 process group lifetime: A period of time that begins when a process group is created and ends when the last remaining process in the group leaves the group, due either to the end of the last process's process lifetime or to the last remaining process calling the *setsid()* or *setpgid()* functions.

313 **2.2.2.67 process ID:** The unique identifier representing a process.

A process ID is a positive integer that can be contained in a pid_t . A process ID shall not be reused by the system until the process lifetime ends. In addition, if there exists a process group whose process group ID is equal to that process ID, the process ID shall not be reused by the system until the process group lifetime ends. A process that is not a system process shall not have a process ID of 1.

2.2.2.68 process lifetime: The period of time that begins when a process is created and ends when its process ID is returned to the system.

After a process is created with a fork() function, it is considered active. Its thread of control and address space exist until it terminates. It then enters an inactive state where certain resources may be returned to the system, although some resources, such as the process ID, are still in use. When another process executes a wait() or waitpid() function for an inactive process, the remaining resources are returned to the system. The last resource to be returned to the system is the process ID. At this time, the lifetime of the process ends.

- **2.2.2.69 read-only file system:** A file system that has implementation-defined characteristics restricting modifications.
- 2.2.2.70 real group ID: The attribute of a process that, at the time of process
 creation, identifies the group of the user who created the process.
- See group ID in 2.2.2.42. This value is subject to change during the process lifetime, as described in 4.2.2.
- 2.2.2.71 real user ID: The attribute of a process that, at the time of process
 creation, identifies the user who created the process.

See *user ID* in 2.2.2.87. This value is subject to change during the process lifetime, as described in 4.2.2.

- 2.2.2.72 regular file: A file that is a randomly accessible sequence of bytes, with
 no further structure imposed by the system.
- **2.2.2.73 relative pathname:** See *pathname resolution* in 2.3.6.
- 341 2.2.2.74 root directory: A directory, associated with a process, that is used in
 pathname resolution for pathnames that begin with a slash.
- 2.2.2.75 saved set-group-ID: An attribute of a process that allows some flexibility in the assignment of the effective group ID attribute, when the saved set-userID option is implemented, as described in 3.1.2 and 4.2.2.
- 2.2.2.76 saved set-user-ID: An attribute of a process that allows some flexibility in the assignment of the effective user ID attribute, when the saved set-user-ID
 option is implemented, as described in 3.1.2 and 4.2.2.
- 349 2.2.2.77 seconds since the Epoch: A value to be interpreted as the number of
 350 seconds between a specified time and the Epoch.
- A Coordinated Universal Time name (specified in terms of seconds (tm_sec) , minutes (tm_min) , hours (tm_hour) , days since January 1 of the year (tm_yday) , and calendar year minus 1900 (tm_year) is related to a time represented as seconds since the Epoch, according to the expression below.
- If the year < 1970 or the value is negative, the relationship is undefined. If the year ≥ 1970 and the value is nonnegative, the value is related to a Coordinated Universal Time name according to the expression:
- 2.2.2.78 session: A collection of process groups established for job control
 purposes.
- Each process group is a member of a session. A process is considered to be a member of the session of which its process group is a member. A newly created process joins the session of its creator. A process can alter its session membership (see 4.3.2). Implementations that support the *setpgid()* function (see 4.3.3) can have multiple process groups in the same session.
- **2.2.2.79 session leader:** A process that has created a session (see 4.3.2).
- **2.2.2.80 session lifetime:** The period between when a session is created and the end of the lifetime of all the process groups that remain as members of the session.
- 371 2.2.2.81 signal: A mechanism by which a process may be notified of, or affected
 372 by, an event occurring in the system.

- Examples of such events include hardware exceptions and specific actions by processes. The term *signal* is also used to refer to the event itself.
- 375 **2.2.2.82 slash:** The literal character "/".
- This character is also known as *solidus* in ISO 8859-1 (B34).
- 2.2.2.83 supplementary group ID: An attribute of a process used in determining file access permissions.
- A process has up to {NGROUPS_MAX} supplementary group IDs in addition to the effective group ID. The supplementary group IDs of a process are set to the supplementary group IDs of the parent process when the process is created. Whether a process's effective group ID is included in or omitted from its list of supplementary group IDs is unspecified.
- 384 **2.2.2.84 system:** An implementation of this part of ISO/IEC 9945.
- **2.2.2.85 system process:** An object, other than a process executing an application, that is defined by the system and has a process ID.
- 387 2.2.2.86 terminal [terminal device]: A character special file that obeys the
 388 specifications of 7.1.
- 2.2.2.87 user ID: A nonnegative integer, which can be contained in an object of
 type uid_t, that is used to identify a system user.
- When the identity of a user is associated with a process, a user ID value is referred to as a real user ID, an effective user ID, or an (optional) saved set-user-ID.
- 394 **2.2.2.88 user name:** A string that is used to identify a user, as described in 9.1.
- 2.2.2.89 working directory [current working directory]: A directory, asso ciated with a process, that is used in pathname resolution for pathnames that do
 not begin with a slash.
- 398 2.2.3 Abbreviations
- ³⁹⁹ For the purposes of this part of ISO/IEC 9945, the following abbreviations apply:
- 400 2.2.3.1 C Standard: ISO/IEC 9899, Information technology—Programming
 401 languages—C {2}.
- 402 2.2.3.2 IRV: The International Reference Version coded character set described
 403 in ISO/IEC 646 {1}.

404 **2.2.3.3 POSIX.1:** This part of ISO/IEC 9945.

405 **2.3 General Concepts**

2.3.1 extended security controls: The access control (see file access permissions) and privilege (see appropriate privileges in 2.2.2.4) mechanisms have been defined to allow implementation-defined extended security controls. These permit an implementation to provide security mechanisms to implement different security policies than described in this part of ISO/IEC 9945. These mechanisms shall not alter or override the defined semantics of any of the functions in this part of ISO/IEC 9945.

2.3.2 file access permissions: The standard file access control mechanism uses
the file permission bits, as described below. These bits are set at file creation by
open(), creat(), mkdir(), and mkfifo() and are changed by chmod(). These bits are
read by stat() or fstat().

Implementations may provide additional or alternate file access control mechanisms, or both. An additional access control mechanism shall only further restrict
the access permissions defined by the file permission bits. An alternate access
control mechanism shall:

- (1) Specify file permission bits for the file owner class, file group class, and
 file other class of the file, corresponding to the access permissions, to be
 returned by stat() or fstat().
- 424 (2) Be enabled only by explicit user action, on a per-file basis by the file
 425 owner or a user with the appropriate privilege.
- 426 (3) Be disabled for a file after the file permission bits are changed for that
 427 file with *chmod()*. The disabling of the alternate mechanism need not
 428 disable any additional mechanisms defined by an implementation.

429 Whenever a process requests file access permission for read, write, or 430 execute/search, if no additional mechanism denies access, access is determined as 431 follows:

- 432 (1) If a process has the appropriate privilege:
- (a) If read, write, or directory search permission is requested, access is
 granted.
- (b) If execute permission is requested, access is granted if execute permission is granted to at least one user by the file permission bits or
 by an alternate access control mechanism; otherwise, access is denied.
- 439 (2) Otherwise:
- (a) The file permission bits of a file contain read, write, and
 execute/search permissions for the file owner class, file group class,
 and file other class.
- (b) Access is granted if an alternate access control mechanism is not
 enabled and the requested access permission bit is set for the class

l

۱

1

- (file owner class, file group class, or file other class) to which the
 process belongs, or if an alternate access control mechanism is
 enabled and it allows the requested access; otherwise, access is
 denied.
- **2.3.3 file hierarchy:** Files in the system are organized in a hierarchical structure in which all of the nonterminal nodes are directories and all of the terminal nodes are any other type of file. Because multiple directory entries may refer to the same file, the hierarchy is properly described as a "directed graph."
- 2.3.4 filename portability: Filenames should be constructed from the portable
 filename character set because the use of other characters can be confusing or
 ambiguous in certain contexts.

2.3.5 file times update: Each file has three distinct associated time values: $st_atime, st_mtime, and st_ctime$. The st_atime field is associated with the times that the file data is accessed; st_mtime is associated with the times that the file data is modified; and st_ctime is associated with the times that file status is changed. These values are returned in the file characteristics structure, as described in 5.6.1.

- Any function in this part of ISO/IEC 9945 that is required to read or write file data 462 or change the file status indicates which of the appropriate time-related fields are 463 to be "marked for update." If an implementation of such a function marks for 464 update a time-related field not specified by this part of ISO/IEC 9945, this shall be 465 documented, except that any changes caused by pathname resolution need not be 466 documented. For the other functions in this part of ISO/IEC 9945 (those that are 467 not explicitly required to read or write file data or change file status, but that in 468 some implementations happen to do so), the effect is unspecified. 469
- An implementation may update fields that are marked for update immediately, or it may update such fields periodically. When the fields are updated, they are set to the current time and the update marks are cleared. All fields that are marked for update shall be updated when the file is no longer open by any process, or when a *stat()* or *fstat()* is performed on the file. Other times at which updates are done are unspecified. Updates are not done for files on read-only file systems.
- **2.3.6 pathname resolution:** Pathname resolution is performed for a process to
 resolve a pathname to a particular file in a file hierarchy. There may be multiple
 pathnames that resolve to the same file.
- Each filename in the pathname is located in the directory specified by its prede-479 cessor (for example, in the pathname fragment "a/b", file "b" is located in direc-480 tory "a"). Pathname resolution fails if this cannot be accomplished. If the path-481 name begins with a slash, the predecessor of the first filename in the pathname is 482 taken to be the root directory of the process (such pathnames are referred to as 483 absolute pathnames). If the pathname does not begin with a slash, the predeces-484 sor of the first filename of the pathname is taken to be the current working direc-485 tory of the process (such pathnames are referred to as "relative pathnames"). 486
- The interpretation of a pathname component is dependent on the values of {NAME_MAX} and {_POSIX_NO_TRUNC} associated with the path prefix of that component. If any pathname component is longer than {NAME_MAX}, and

- 490 {_POSIX_NO_TRUNC} is in effect for the path prefix of that component (see 5.7.1),
 491 the implementation shall consider this an error condition. Otherwise, the imple492 mentation shall use the first {NAME_MAX} bytes of the pathname component.
- The special filename, dot, refers to the directory specified by its predecessor. The special filename, dot-dot, refers to the parent directory of its predecessor directory. As a special case, in the root directory, dot-dot may refer to the root directory itself.
- A pathname consisting of a single slash resolves to the root directory of the process. A null pathname is invalid.

499 **2.4 Error Numbers**

- 500 Most functions provide an error number in the external variable *errno*, which is 501 defined as:
- 502 extern int errno;

The value of this variable shall be defined only after a call to a function for which it is explicitly stated to be set and until it is changed by the next function call. The variable *errno* should only be examined when it is indicated to be valid by a function's return value. No function defined in this part of ISO/IEC 9945 sets *errno* to zero to indicate an error.

- If more than one error occurs in processing a function call, this part of ISO/IEC 9945 does not define in what order the errors are detected; therefore, any one of the possible errors may be returned.
- Implementations may support additional errors not included in this clause, may 511 generate errors included in this clause under circumstances other than those 512 described in this clause, or may contain extensions or limitations that prevent 513 some errors from occurring. The Errors subclause in each function description 514 specifies which error conditions shall be detected by all implementations and 515 which may be optionally detected by an implementation. Each implementation 516 shall document, in the conformance document, situations in which each of the 517 optional conditions are detected. If no error condition is detected, the action 518 requested shall be successful. Implementations may contain extensions or limita-519 tions that prevent some specified errors from occurring. 520
- Implementations may generate error numbers listed in this clause under circumstances other than those described, if and only if all those error conditions can always be treated identically to the error conditions as described in this part of ISO/IEC 9945. Implementations may support additional errors not listed in this clause, but shall not generate a different error number from one required by this part of ISO/IEC 9945 for an error condition described in this part of ISO/IEC 9945.

The following symbolic names identify the possible error numbers, in the context of functions specifically defined in this part of ISO/IEC 9945; these general descriptions are more precisely defined in the Errors subclauses of functions that return them. Only these symbolic names should be used in programs, since the actual value of an error number is unspecified. All values listed in this clause shall be unique. The values for these names shall be found in the header <erro.h>.

L

533	The actual valu	es are unspecified by this part of ISO/IEC 9945.
534 535 536 537	[E2BIG]	Arg list too long The sum of the number of bytes used by the new process image's argument list and environment list was greater than the system- imposed limit of {ARG_MAX} bytes.
538 539 540	[EACCES]	Permission denied An attempt was made to access a file in a way forbidden by its file access permissions.
541 542 543	[EAGAIN]	Resource temporarily unavailable This is a temporary condition, and later calls to the same routine may complete normally.
544 545 546 547	[EBADF]	Bad file descriptor A file descriptor argument was out of range, referred to no open file, or a read (write) request was made to a file that was only open for writing (reading).
548 549 550 551 552	[EBUSY]	Resource busy An attempt was made to use a system resource that was not available at the time because it was being used by a process in a manner that would have conflicted with the request being made by this process.
553 554 555	[ECHILD]	No child processes A <i>wait()</i> or <i>waitpid()</i> function was executed by a process that had no existing or unwaited-for child processes.
556 557 558	[EDEADLK]	Resource deadlock avoided An attempt was made to lock a system resource that would have resulted in a deadlock situation.
559 560 561	[EDOM]	Domain error Defined in the C Standard {2}; an input argument was outside the defined domain of the mathematical function.
562 563 564	[EEXIST]	File exists An existing file was specified in an inappropriate context; for instance, as the new link name in a $link()$ function.
565 566 567 568 569	[EFAULT]	Bad address The system detected an invalid address in attempting to use an argument of a call. The reliable detection of this error is imple- mentation defined; however, implementations that do detect this condition shall use this value.
570 571 572	[EFBIG]	File too large The size of a file would exceed an implementation-defined max- imum file size.
573 574 575 576	[EINTR]	Interrupted function call An asynchronous signal (such as SIGINT or SIGQUIT; see the description of header $< signal.h > in 3.3.1$) was caught by the process during the execution of an interruptible function. If the

577 578		signal handler performs a normal return, the interrupted func- tion call may return this error condition.
579 580 581	[EINVAL]	Invalid argument Some invalid argument was supplied. [For example, specifying an undefined signal to a <i>signal()</i> or <i>kill()</i> function].
582 583 584 585 586	[EIO]	Input/output error Some physical input or output error occurred. This error may be reported on a subsequent operation on the same file descriptor. Any other error-causing operation on the same file descriptor may cause the [EIO] error indication to be lost.
587 588 589	[EISDIR]	Is a directory An attempt was made to open a directory with write mode specified.
590 591 592	[EMFILE]	Too many open files An attempt was made to open more than the maximum number of {OPEN_MAX} file descriptors allowed in this process.
593 594 595	[EMLINK]	Too many links An attempt was made to have the link count of a single file exceed {LINK_MAX}.
596 597 598 599	[ENAMETOC	DLONG] Filename too long The size of a pathname string exceeded {PATH_MAX}, or a path- name component was longer than {NAME_MAX} and {_POSIX_NO_TRUNC} was in effect for that file.
600 601 602 603	[ENFILE]	Too many open files in system Too many files are currently open in the system. The system reached its predefined limit for simultaneously open files and temporarily could not accept requests to open another one.
604 605 606 607	[ENODEV]	No such device An attempt was made to apply an inappropriate function to a dev- ice; for example, trying to read a write-only device such as a printer.
608 609 610	[ENOENT]	No such file or directory A component of a specified pathname did not exist, or the path- name was an empty string.
611 612 613 614	[ENOEXEC]	Exec format error A request was made to execute a file that, although it had the appropriate permissions, was not in the format required by the implementation for executable files.
615 616 617 618	[ENOLCK]	No locks available A system-imposed limit on the number of simultaneous file and record locks was reached, and no more were available at that time.
619 620	[ENOMEM]	Not enough space The new process image required more memory than was allowed

621 622		by the hardware or by system-imposed memory management constraints.
623 624 625	[ENOSPC]	No space left on device During a <i>write()</i> function on a regular file, or when extending a directory, there was no free space left on the device.
626 627 628	[ENOSYS]	Function not implemented An attempt was made to use a function that is not available in this implementation.
629 630 631	[ENOTDIR]	Not a directory A component of the specified pathname existed, but it was not a directory, when a directory was expected.
632 633 634	[ENOTEMPI	[Y] Directory not empty A directory with entries other than dot and dot-dot was supplied when an empty directory was expected.
635 636 637	[ENOTTY]	Inappropriate I/O control operation A control function was attempted for a file or a special file for which the operation was inappropriate.
638 639 640 641 642	[ENXIO]	No such device or address Input or output on a special file referred to a device that did not exist, or made a request beyond the limits of the device. This error may also occur when, for example, a tape drive is not online or a disk pack is not loaded on a drive.
643 644 645 646	[EPERM]	Operation not permitted An attempt was made to perform an operation limited to processes with appropriate privileges or to the owner of a file or other resource.
647 648 649	[EPIPE]	Broken pipe A write was attempted on a pipe or FIFO for which there was no process to read the data.
650 651 652	[ERANGE]	Result too large Defined in the C Standard {2}; the result of the function was too large to fit in the available space.
653 654 655	[EROFS]	Read-only file system An attempt was made to modify a file or directory on a file system that was read-only at that time.
656 657	[ESPIPE]	Invalid seek An <i>lseek()</i> function was issued on a pipe or FIFO.
658 659 660	[ESRCH]	No such process No process could be found corresponding to that specified by the given process ID.
661 662	[EXDEV]	Improper link A link to a file on another file system was attempted.

663 2.5 Primitive System Data Types

Some data types used by the various system functions are not defined as part of this part of ISO/IEC 9945, but are defined by the implementation. These types are then defined in the header <sys/types.h>, which contains definitions for at least the types shown in Table 2-1.

Defined Type	Description
dev_t	Used for device numbers.
gid_t	Used for group IDs.
ino_t	Used for file serial numbers.
mode_t	Used for some file attributes, for example file type, file access permissions.
nlink_t	Used for link counts.
off_t	Used for file sizes.
pid_t	Used for process IDs and process group IDs.
size_t	As defined in the C Standard {2}.
ssize_t	Used by functions that return a count of bytes (memory space) or an error indication.
uid_t	Used for user IDs.

Table 2-1 - Primitive System Data Types

All of the types listed in Table 2-1 shall be arithmetic types; pid_t , $ssize_t$, and off_t shall be signed arithmetic types. The type $ssize_t$ shall be capable of storing values in the range from -1 to (SSIZE_MAX), inclusive. The types $size_t$ and $ssize_t$ shall also be defined in the header <unistd.h>.

687 Additional unspecified type symbols ending in _t may be defined in any header 688 specified by POSIX.1. The visibility of such symbols need not be controlled by any 689 feature test macro other than _POSIX_SOURCE.

690 **2.6 Environment Description**

An array of strings called the *environment* is made available when a process begins. This array is pointed to by the external variable *environ*, which is defined as:

694 extern char **environ;

These strings have the form "name=value"; names shall not contain the character '='. There is no meaning associated with the order of the strings in the environment. If more than one string in a process's environment has the same name, the consequences are undefined. The following names may be defined and have the indicated meaning if they are defined:

700HOMEThe name of the user's initial working directory from the
user database (see the description of the header <pwd.h>701in 9.2.2).

703 704 705	LANG	The name of the locale to use for locale categories when both LC_ALL and the corresponding environment variable (beginning with "LC_") do not specify a locale.
706 707 708	LC_ALL	The name of the locale to be used to override any values for locale categories specified by the setting of LANG or any environment variables beginning with "LC_".
709	LC_COLLATE	The name of the locale for collation information.
710	LC_CTYPE	The name of the locale for character classification.
711 712	LC_MONETARY	The name of the locale containing monetary-related numeric editing information.
713 714	LC_NUMERIC	The name of the locale containing numeric editing (i.e., radix character) information.
715 716	LC_TIME	The name of the locale for date/time formatting information.
717 718 719	LOGNAME	The login name associated with the current process. The value shall be composed of characters from the portable filename character set.
720 721 722 723 724 725 726 727		NOTE: An application that requires, or an installation that actually uses, characters outside the portable filename character set would not strictly conform to this part of ISO/IEC 9945. However, it is reasonable to expect that such characters would be used in many countries (recog- nizing the reduced level of interchange implied by this), and applica- tions or installations should permit such usage where possible. No error is defined by this part of ISO/IEC 9945 for violation of this condition.
728 729 730 731 732 733 734 735 736 737 738 737 738 739 740 741	PATH	The sequence of path prefixes that certain functions apply in searching for an executable file known only by a filename (a pathname that does not contain a slash). The prefixes are separated by a colon (:). When a nonzero- length prefix is applied to this filename, a slash is inserted between the prefix and the filename. A zero- length prefix is a special prefix that indicates the current working directory. It appears as two adjacent colons (::), as an initial colon preceding the rest of the list, or as a trailing colon following the rest of the list. The list is searched from beginning to end until an executable pro- gram by the specified name is found. If the pathname being sought contains a slash, the search through the path prefixes is not performed.
742 743 744 745	TERM	The terminal type for which output is to be prepared. This information is used by commands and application programs wishing to exploit special capabilities specific to a terminal.

746TZTime zone information. The format of this string is747defined in 8.1.1.

Environment variable *names* used or created by an application should consist solely of characters from the portable filename character set. Other characters may be permitted by an implementation; applications shall tolerate the presence of such names. Upper- and lowercase letters retain their unique identities and are not folded together. System-defined environment variable names should begin with a capital letter or underscore and be composed of only capital letters, underscores, and numbers.

The values that the environment variables may be assigned are not restricted except that they are considered to end with a null byte, and the total space used to store the environment and the arguments to the process is limited to (ARG_MAX) bytes.

759 Other *name=value* pairs may be placed in the environment by manipulating the 760 *environ* variable or by using *envp* arguments when creating a process (see 3.1.2).

761 **2.7 C Language Definitions**

762 2.7.1 Symbols From the C Standard

The following terms and symbols used in this part of ISO/IEC 9945 are defined in the C Standard {2}: NULL, byte, array of char, clock_t, header, null character, string, time_t. The type clock_t shall be capable of representing all integer values from zero to the number of clock ticks in 24 h.

The term NULL pointer in this part of ISO/IEC 9945 is equivalent to the term null pointer used in the C Standard {2}. The symbol NULL shall be declared in <unistd.h> with the same value as required by the C Standard {2}, in addition to several headers already required by the C Standard {2}.

- Additionally, the reservation of symbols that begin with an underscore applies:
- (1) All external identifiers that begin with an underscore are reserved.
- (2) All other identifiers that begin with an underscore and either an upper case letter or another underscore are reserved.
- (3) If the program defines an external identifier with the same name as a
 reserved external identifier, even in a semantically equivalent form, the
 behavior is undefined.

Certain other namespaces are reserved by the C Standard {2}. These reservations
apply to this part of ISO/IEC 9945 as well. Additionally, the C Standard {2}
requires that it be possible to include a header more than once and that a symbol
may be defined in more than one header. This requirement is also made of
headers for this part of ISO/IEC 9945.

783 **2.7.2 POSIX.1 Symbols**

Certain symbols in this part of ISO/IEC 9945 are defined in headers. Some of those headers could also define other symbols than those defined by this part of ISO/IEC 9945, potentially conflicting with symbols used by the application. Also, this part of ISO/IEC 9945 defines symbols that are not permitted by other standards to appear in those headers without some control on the visibility of those symbols.

Symbols called *feature test macros* are used to control the visibility of symbols that might be included in a header. Implementations, future versions of this part of ISO/IEC 9945, and other standards may define additional feature test macros. Feature test macros shall be defined in the compilation of an application before an #include of any header where a symbol should be visible to some, but not all, applications. If the definition of the macro does not precede the #include, the result is undefined.

⁷⁹⁷ Feature test macros shall begin with the underscore character (_).

Implementations may add symbols to the headers shown in Table 2-2, provided the identifiers for those symbols begin with the corresponding reserved prefixes in Table 2-2. Similarly, implementations may add symbols to the headers in Table 2-2 that end in the string indicated as a reserved suffix as long as the reserved suffix is in that part of the name considered significant by the implementation. This shall be in addition to any reservations made in the C Standard {2}.

- If any header defined by this part of ISO/IEC 9945 is included, all symbols with the suffix _t are reserved for use by the implementation, both before and after the #include directive.
- After the last inclusion of a given header, an application may use any of the symbol classes reserved in Table 2-2 for its own purposes, as long as the requirements in the note to Table 2-2 are satisfied, noting that the symbol declared in the header may become inaccessible.
- Future revisions of this part of ISO/IEC 9945, and other POSIX standards, are likely to use symbols in these same reserved spaces.
- In addition, implementations may add members to a structure or union without controlling the visibility of those members with a feature test macro, as long as a user-defined macro with the same name cannot interfere with the correct interpretation of the program.
- The header <fcntl.h> may contain the following symbols in addition to those specifically required elsewhere in POSIX.1:

819	SEEK_CUR	S_IRUSR	S_ISCHR	S_ISREG	S_IWUSR
820	SEEK_END	S_IRWXG	S_ISDIR	S_ISUID	S_IXGRP
821	SEEK_SET	S_IRWXO	S_ISFIFO	S_IWGRP	S_IXOTH
822	S_IRGRP	S_IRWXU	S_ISGID	S_IWOTH	S_IXUSR
823	S_IROTH	S_ISBLK			

In addition, an implementation may define the symbols "cuserid" in <unistd.h> and "L_cuserid" in <stdio.h>.

828	Header	Kev	Reserved	Reserved
829			Prefix	Suffix
830	<dirent.h></dirent.h>	1	d	
831	<fcntl.h></fcntl.h>	1	1	
832		2	F	
833		2	0	
834		2	s	
835	<grp.h></grp.h>	1	gr	
836	<limits.h></limits.h>	1		MAX
837	<locale.h></locale.h>	2	LC [A-Z]	—
838	<pwd.h></pwd.h>	1	pw_	
839	<signal.h></signal.h>	1	sa	
840		2	SIG	
841		2	SA	
842	<sys stat.h=""></sys>	1	st	
843		2	s_	
844	<sys times.h=""></sys>	1	tms_	
845	<termios.h></termios.h>	1	c	
846		2	V	
847		2	I	
848		2	0	
849		2	TC	
850		2	B[0-9]	
851	any POSIX.1 header included	1		t
852	-			_

Table 2-2 - Reserved Header Symbols

NOTE: The notation "[0-9]" indicates any digit and "[A-Z]" any uppercase character in the portable
filename character set. The Key values are:

855 (1) Prefixes and suffixes of symbols that shall not be declared or #defined by the application.

856 (2) Prefixes and suffixes of symbols that shall be preceded in the application with a #undef of
 857 that symbol before any other use.

858 The following feature test macro is defined:

859	Name		Description
860	_POSIX_SOURCE	When a	n application includes a header described by
861		POSIX.1	, and when this feature test macro is defined
862		accordin	ng to the preceding rules:
863		(1)	All symbols required by POSIX.1 to appear when
864			the header is included shall be made visible.
865		(2)	Symbols that are explicitly permitted, but not
866			required, by POSIX.1 to appear in that header
867			(including those in reserved namespaces) may be
868			made visible.
869		(3)	Additional symbols not required or explicitly per-
870			mitted by POSIX.1 to be in that header shall not
871			be made visible.

I

The exact meaning of feature test macros depends on the type of C language support chosen: C Standard Language-Dependent Support and Common-Usage-Dependent Support, described in the following two subclauses.

875 2.7.2.1 C Standard Language-Dependent Support

If there are no feature test macros present in a program, the implementation shall make visible only those identifiers specified as reserved identifiers in the C Standard (2), permitting the reservation of symbols and namespace defined in 2.7.1. For each feature test macro present, only the symbols specified by that feature test macro plus those of the C Standard (2) shall be defined when a header is included.

882 2.7.2.2 Common-Usage-Dependent Support

If the feature test macro _POSIX_SOURCE is not defined in a program, the set of
symbols defined in each header that are beyond the requirements of this part of
ISO/IEC 9945 is unspecified.

If _POSIX_SOURCE is defined before any header is included, no symbols other than those from the C Standard {2} and those made visible by feature test macros defined for the program (including _POSIX_SOURCE) will be visible. Symbols from the namespace reserved for the implementation, as defined by the C Standard {2}, are also permitted. The symbols beginning with two underscores are examples of this.

⁸⁹² If _POSIX_SOURCE is not defined before any header is included, the behavior is ⁸⁹³ undefined.

894 2.7.3 Headers and Function Prototypes

Implementations claiming C Standard {2} Language-Dependent Support shall
 declare function prototypes for all functions.

Implementations claiming Common-Usage C Language-Dependent Support shall
 declare the result type for all functions not returning a "plain" *int*.

For functions described in the C Standard {2} and included by reference in Section 899 8 (whether or not they are further described in this part of ISO/IEC 9945), these 900 prototypes or declarations (if required) shall appear in the headers defined for 901 them in the C Standard {2}. For other functions in this part of ISO/IEC 9945, the 902 prototypes or declarations shall appear in the headers listed below. If a function 903 is defined by this part of ISO/IEC 9945, is not described in the C Standard {2}, and 904 is not listed below, it shall have its prototype or declaration (if required) appear in 905 <unistd.h>, which shall be #include-ed by the application before using any 906 function declared in it, whether or not it is mentioned in the Synopsis subclause 907 for that function. The requirements about the visibility of symbols in 2.7.2 shall 908 be honored. 909

910	<dirent.h></dirent.h>	<pre>opendir(), readdir(), rewinddir(), closedir()</pre>
911	<fcntl.h></fcntl.h>	open(), creat(), fcntl()
912	<grp.h></grp.h>	getgrgid(), getgrnam()
913	<pwd.h></pwd.h>	getpwuid(), getpwnam()
914	<setjmp.h></setjmp.h>	sigsetjmp(), siglongjmp()
915 916 917	<signal.h></signal.h>	<pre>kill(), sigemptyset(), sigfillset(), sigaddset(), sigdelset(), sigismember(), sigaction(), sigprocmask(), sigpending(), sigsuspend()</pre>
918	<stdio.h></stdio.h>	ctermid(), fileno(), fdopen()
919	<sys stat.h=""></sys>	umask(), mkdir(), mkfifo(), stat(), fstat(), chmod()
920	<sys times.h=""></sys>	times()
921	<sys utsname.h=""></sys>	uname()
922	<sys wait.h=""></sys>	<pre>wait(), waitpid()</pre>
923 924 925	<termios.h></termios.h>	<pre>cfgetospeed(), cfsetospeed(), cfgetispeed(), cfsetispeed(), tcgetattr(), tcsetattr(), tcsendbreak(), tcdrain(), tcflush(), tcflow()</pre>
926	<time.h></time.h>	time(), tzset()
927	<utime.h></utime.h>	utime()

The declarations in the headers shall follow the proper form for the C language 928 option chosen by the implementation. Additionally, pointer arguments that refer 929 to objects not modified by the function being described are declared with const 930 qualifying the type to which it points. Implementations claiming Common-Usage 931 C conformance to this part of ISO/IEC 9945 may ignore the presence of this key-932 word and need not include it in any function declarations. Implementations 933 claiming conformance using the C Standard {2} shall use the const modifier as 934 indicated in the prototypes they provide. 935

Implementations claiming conformance using Common-Usage C may use
equivalent implementation-defined constructs when void is used as a result type
for a function prototype. They may also use *int* when a function result is declared
ssize_t.

Neither the names of the formal parameters nor their types, as they appear in an 940 implementation, are specified by this part of ISO/IEC 9945. The names are used 941 within this part of ISO/IEC 9945 as a notational mechanism. However, any 942 declaration provided by an implementation shall accept all actual parameter 943 944 types that a declaration lexically identical to one in this part of ISO/IEC 9945 shall accept, including the effects of both type conversion and checking for the number 945 of arguments implied by the presence of a filled-out prototype. 946 The implementation's declaration shall not cause a syntax error if an application pro-947 vides a prototype lexically identical to one in this part of ISO/IEC 9945. It is not a 948 requirement that nonconforming parameters to functions that may be used by an 949 application be diagnosed by an implementation, except as specifically required by 950 this part of ISO/IEC 9945 or the C Standard {2}, as applicable. Where the 951

I

L

L

C Standard {2} has a more restrictive requirement for a function defined by that standard, that requirement shall be honored, and this exception does not apply.

954 **2.8 Numerical Limits**

The following subclauses list magnitude limitations imposed by a specific implementation. The braces notation, {LIMIT}, is used in this part of ISO/IEC 9945 to indicate these values, but the braces are not part of the name.

958 2.8.1 C Language Limits

The following limits used in this part of ISO/IEC 9945 are defined in the C Standard {2}: {CHAR_BIT}, {CHAR_MAX}, {CHAR_MIN}, {INT_MAX}, {INT_MIN}, {LONG_MAX}, {LONG_MIN}, {MB_LEN_MAX}, {SCHAR_MAX}, {SCHAR_MIN}, SHRT_MAX}, {SHRT_MIN}, {UCHAR_MAX}, {UINT_MAX}, {ULONG_MAX}, {USHRT_MAX}.

964 **2.8.2 Minimum Values**

The symbols in Table 2-3 shall be defined in <limits.h> with the values shown. These are symbolic names for the most restrictive value for certain features on a system conforming to this part of ISO/IEC 9945. Related symbols are defined elsewhere in this part of ISO/IEC 9945, which reflect the actual implementation and which need not be as restrictive. A conforming implementation shall provide values at least this large. A portable application shall not require a larger value for correct operation.

972 **2.8.3 Run-Time Increasable Values**

⁹⁷³ The magnitude limitations in Table 2-4 shall be fixed by specific implementations.

A Strictly Conforming POSIX.1 Application shall assume that the value supplied by <limits.h> in a specific implementation is the minimum value that pertains whenever the Strictly Conforming POSIX.1 Application is run under that implementation.³⁾ A specific instance of a specific implementation may increase the value relative to that supplied by <limits.h> for that implementation. The actual value supported by a specific instance shall be provided by the *sysconf()* function.

1010

35	Name	Description	Value
36 37	[_POSIX_ARG_MAX}	The length of the arguments for one of the <i>exec</i> func- tions, in bytes, including environment data.	4096
38 39	{_POSIX_CHILD_MAX}	The number of simultaneous processes per real user ID.	6
90	{_POSIX_LINK_MAX}	The value of a file's link count.	8
)1)2	{_POSIX_MAX_CANON}	The number of bytes in a terminal canonical input queue.	255
)3)4	{_POSIX_MAX_INPUT}	The number of bytes for which space will be available in a terminal input queue.	255
95	{_POSIX_NAME_MAX}	The number of bytes in a filename.	14
96 97	{_POSIX_NGROUPS_MAX}	The number of simultaneous supplementary group IDs per process.	0
)8)9	{_POSIX_OPEN_MAX}	The number of files that one process can have open at one time.	16
000	{_POSIX_PATH_MAX}	The number of bytes in a pathname.	255
001 002	{_POSIX_PIPE_BUF}	The number of bytes that can be written atomically when writing to a pipe.	512
03 04	[_POSIX_SSIZE_MAX}	The value that can be stored in an object of type <i>ssize_t</i> .	32767
05 06	[_POSIX_STREAM_MAX}	The number of streams that one process can have open at one time.	8
07 08	{_POSIX_TZNAME_MAX}	The maximum number of bytes supported for the name of a time zone (not of the TZ variable).	3

Table 2-3 – Minimum Values

Table 2-4 - Run-Time Increasable Values

Name	Description	Minimum Value
{NGROUPS_MAX}	Maximum number of simultaneous supple- mentary group IDs per process.	{_POSIX_NGROUPS_MAX}

1016 **2.8.4 Run-Time Invariant Values (Possibly Indeterminate)**

1017 A definition of one of the values in Table 2-5 shall be omitted from the 1018 <limits.h> on specific implementations where the corresponding value is equal 1019 to or greater than the stated minimum, but is indeterminate.

1020 This might depend on the amount of available memory space on a specific 1021 instance of a specific implementation. The actual value supported by a specific 1022 instance shall be provided by the *sysconf()* function.

Name	Description	Minimum Value
{ARG_MAX}	Maximum length of arguments for the <i>exec</i> functions, in bytes, including environment data.	{_POSIX_ARG_MAX}
{CHILD_MAX}	Maximum number of simultaneous processes per real user ID.	{_POSIX_CHILD_MAX}
{OPEN_MAX}	Maximum number of files that one process can have open at any given time.	{_POSIX_OPEN_MAX}
{STREAM_MAX}	The number of streams that one process can have open at one time. If defined, it shall have the same value as (FOPEN_MAX) from the C Standard (2).	{_POSIX_STREAM_MAX}
{TZNAME_MAX}	The maximum number of bytes supported for the name of a time zone (not of the TZ variable).	{_POSIX_TZNAME_MAX}

Table 2-5 – Run-Time Invariant Values (Possibly Indeterminate)

1041 2.8.5 Pathname Variable Values

The values in Table 2-6 may be constants within an implementation or may vary from one pathname to another.

Name	Description	Minimum Value
{LINK_MAX}	Maximum value of a file's link count.	{_POSIX_LINK_MAX}
{MAX_CANON}	Maximum number of bytes in a terminal canonical input line. (See 7.1.1.6.)	{_POSIX_MAX_CANON}
{MAX_INPUT}	Minimum number of bytes for which space will be available in a terminal input queue; therefore, the maximum number of bytes a portable application may require to be typed as input before reading them.	{_POSIX_MAX_INPUT}
{NAME_MAX}	Maximum number of bytes in a file name (not a string length; count excludes a ter- minating null).	{_POSIX_NAME_MAX}
{PATH_MAX}	Maximum number of bytes in a pathname (not a string length; count excludes a ter- minating null).	{_POSIX_PATH_MAX}
{PIPE_BUF}	Maximum number of bytes that can be written atomically when writing to a pipe.	{_POSIX_PIPE_BUF}

Table 2-6 - Pathname Variable Values

1064 For example, file systems or directories may have different characteristics.

1065 A definition of one of the values from Table 2-6 shall be omitted from 1066 <limits.h> on specific implementations where the corresponding value is equal 1067 to or greater than the stated minimum, but where the value can vary depending

on the file to which it is applied. The actual value supported for a specific pathname shall be provided by the pathconf() function.

1070 2.8.6 Invariant Values

1071 The value in Table 2-7 shall not vary in a given implementation. The value in 1072 that table shall appear in <limits.h>.

1073		Table 2-7 – Invariant Value	
1075	Name	Description	Value
1076 1077	{SSIZE_MAX}	The maximum value that can be stored in an object of type <i>ssize_t</i> .	{_POSIX_SSIZE_MAX}
1078			

1079 2.9 Symbolic Constants

A conforming implementation shall have the header <unistd.h>. This header defines the symbolic constants and structures referenced elsewhere in this part of ISO/IEC 9945. The constants defined by this header are shown in the following subclauses. The actual values of the constants are implementation defined.

1084 2.9.1 Symbolic Constants for the access() Function

The constants used by the access() function are shown in Table 2-8. The constants F_OK, R_OK, W_OK, and X_OK, and the expressions

- 1087 **R_OK | W_OK**
- 1088 (where the | represents the bitwise inclusive OR operator),
- 1089 R_OK | X_OK

1090 and

- 1091 R_OK | W_OK | X_OK
- 1092 shall all have distinct values.

1093 2.9.2 Symbolic Constant for the lseek() Function

1094 The constants used by the *lseek()* function are shown in Table 2-9.

Constant	Description
R_OK	Test for read permission.
W_OK	Test for write permission.
X_ok	Test for execute or search permission.
F_OK	Test for existence of file.

L03	Table 2-9 – Symbolic Constants for the lseek() Function		
105	Constant	Description	
106	SEEK_SET	Set file offset to offset.	
107	SEEK_CUR	Set file offset to current plus offset.	
.08	SEEK_END	Set file offset to EOF plus offset.	

1110 2.9.3 Compile-Time Symbolic Constants for Portability Specifications

1111 The constants in Table 2-10 may be used by the application, at compile time, to 1112 determine which optional facilities are present and what actions shall be taken by 1113 the implementation.

2-10 – Compile-Time Symbolic Constants
Description
If this symbol is defined, it indicates that the implementation supports job control.
If defined, each process has a saved set-user-ID and a saved set- group-ID.
The integer value 199009L. This value shall be used for systems that conform to this part of ISO/IEC 9945.

Although a Strictly Conforming POSIX.1 Application can rely on the values compiled from the <unistd.h> header to afford it portability on all instances of an implementation, it may choose to interrogate a value at run-time to take advantage of the current configuration. See 4.8.1.

1128 2.9.4 Execution-Time Symbolic Constants for Portability Specifications

The constants in Table 2-11 may be used by the application, at execution time, to determine which optional facilities are present and what actions shall be taken by the implementation in some circumstances described by this part of ISO/IEC 9945

1132 as implementation defined.

Name	Description
{_POSIX_CHOWN_RESTRICTED}	The use of the <i>chown</i> () function is restricted to a process with appropriate privileges, and to changing the group ID of a file only to the effective group ID of the process or to one of its supplementary group IDs.
(_POSIX_NO_TRUNC)	Pathname components longer than {NAME_MAX} generate an error.
(_POSIX_VDISABLE)	Terminal special characters defined in 7.1.1.9 can be dis- abled using this character value, if it is defined. See <i>tcgetattr()</i> and <i>tcsetattr()</i> .

Table 2-11 – Execution-Time Symbolic Constants

1146 If any of the constants in Table 2-11 are not defined in the header <unistd.h>, 1147 the value varies depending on the file to which it is applied. See 5.7.1.

1148 If any of the constants in Table 2-11 are defined to have value -1 in the header 1149 <unistd.h>, the implementation shall not provide the option on any file; if any 1150 are defined to have a value other than -1 in the header <unistd.h>, the imple-1151 mentation shall provide the option on all applicable files.

All of the constants in Table 2-11, whether defined in <unistd.h> or not, may be queried with respect to a specific file using the *pathconf()* or *fpathconf()* functions.



Section 3: Process Primitives

The functions described in this section perform the most primitive operating system services dealing with processes, interprocess signals, and timers. All attributes of a process that are specified in this part of ISO/IEC 9945 shall remain unchanged by a process primitive unless the description of that process primitive states explicitly that the attribute is changed.

6 **3.1 Process Creation and Execution**

- 7 3.1.1 Process Creation
- 8 Function: fork()

9 **3.1.1.1 Synopsis**

- 10 #include <sys/types.h>
- 11 pid_t fork(void);

12 **3.1.1.2 Description**

15

16

The *fork()* function creates a new process. The new process (child process) shall be an exact copy of the calling process (parent process) except for the following:

- (1) The child process has a unique process ID. The child process ID also does not match any active process group ID.
- 17 (2) The child process has a different parent process ID (which is the process
 18 ID of the parent process).
- (3) The child process has its own copy of the parent's file descriptors. Each of the child's file descriptors refers to the same open file description with the corresponding file descriptor of the parent.
- (4) The child process has its own copy of the parent's open directory streams
 (see 5.1.2). Each open directory stream in the child process may share
 directory stream positioning with the corresponding directory stream of
 the parent.
- (5) The child process's values of tms_utime, tms_stime, tms_cutime, and
 tms_cstime are set to zero (see 4.5.2).

- (6) File locks previously set by the parent are not inherited by the child.
 (See 6.5.2.)
- 30 (7) Pending alarms are cleared for the child process. (See 3.4.1.)
 - (8) The set of signals pending for the child process is initialized to the empty set. (See 3.3.1.)

All other process characteristics defined by this part of ISO/IEC 9945 shall be the same in the parent and the child processes. The inheritance of process characteristics not defined by this part of ISO/IEC 9945 is unspecified by this part of ISO/IEC 9945, but should be documented in the system documentation.

After *fork()*, both the parent and the child processes shall be capable of executing independently before either terminates.

39 **3.1.1.3 Returns**

31

32

⁴⁰ Upon successful completion, fork() shall return a value of zero to the child process ⁴¹ and shall return the process ID of the child process to the parent process. Both ⁴² processes shall continue to execute from the fork() function. Otherwise, a value of ⁴³ -1 shall be returned to the parent process, no child process shall be created, and ⁴⁴ errno shall be set to indicate the error.

45 **3.1.1.4 Errors**

- If any of the following conditions occur, the fork() function shall return -1 and set *errno* to the corresponding value:
- 48 [EAGAIN] The system lacked the necessary resources to create another 49 process, or the system-imposed limit on the total number of 50 processes under execution by a single user would be exceeded.
- For each of the following conditions, if the condition is detected, the fork() function shall return -1 and set *errno* to the corresponding value:
- 53 [ENOMEM] The process requires more space than the system is able to 54 supply.
- 55 3.1.1.5 Cross-References
- 56 alarm(), 3.4.1; exec, 3.1.2; fcntl(), 6.5.2; kill(), 3.3.2; times(), 4.5.2; wait, 3.2.1.

57 3.1.2 Execute a File

58 Functions: execl(), execv(), execle(), execve(), execvp(), execvp().

59 **3.1.2.1 Synopsis**

- 60 int execl(const char *path, const char *arg, ...);
- 61 int execv(const char *path, char *const argv[]);

L

62 int execle(const char *path, const char *arg, ...); 63 int execve(const char *path, char *const argv[], char *const envp[]); 64 int execlp(const char *file, const char *arg, ...); 65 classical const char *file, const char *arg, ...); 66 classical const char *file, const char *arg, ...); 67 classical const char *file, const char *arg, ...); 68 classical const char *file, const char *arg, ...); 69 classical const char *file, const char *arg, ...); 60 classical const char *file, const char *arg, ...); 61 classical const char *file, const char *arg, ...); 62 classical const char *file, const char *arg, ...); 63 classical const char *file, const char *arg, ...); 64 classical const char *file, const char *arg, ...); 65 classical const char *file, const char *arg, ...); 66 classical const char *file, const char *arg, ...); 67 classical const char *file, const char *arg, ...); 68 classical const char *file, const char *arg, ...); 69 classical const char *file, const char *arg, ...); 60 classical const char *file, const char *arg, ...); 61 classical const char *file, const char *arg, ...); 62 classical const char *file, const char *arg, ...); 63 classical const char *file, const char *arg, ...); 64 classical const char *file, const char *arg, ...); 64 classical const char *file, const char *arg, ...); 65 classical const char *file, const char *arg, ...); 66 classical const char *file, const char *arg, ...); 67 classical const char *file, const char *file, const char *arg, ...); 68 classical const char *file, const char

65 int execvp(const char *file, char *const argv[]);

66 **3.1.2.2 Description**

The exec family of functions shall replace the current process image with a new process image. The new image is constructed from a regular, executable file called the *new process image file*. There shall be no return from a successful *exec* because the calling process image is overlaid by the new process image.

- 71 When a C program is executed as a result of this call, it shall be entered as a C 72 language function call as follows:
- 73 int main(int argc, char *argv[]);
- where *argc* is the argument count and *argv* is an array of character pointers to the arguments themselves. In addition, the following variable:
- 76 extern char **environ;
- is initialized as a pointer to an array of character pointers to the environment
 strings. The *argv* and *environ* arrays are each terminated by a NULL pointer.
 The NULL pointer terminating the *argv* array is not counted in *argc*.
- The arguments specified by a program with one of the *exec* functions shall be passed on to the new process image in the corresponding *main()* arguments.
- The argument *path* points to a pathname that identifies the new process image file.
- The argument *file* is used to construct a pathname that identifies the new process image file. If the *file* argument contains a slash character, the *file* argument shall be used as the pathname for this file. Otherwise, the path prefix for this file is obtained by a search of the directories passed as the environment variable **PATH** (see 2.6). If this environment variable is not present, the results of the search are implementation defined.
- 90

The argument *argv* is an array of character pointers to null-terminated strings. The last member of this array shall be a NULL pointer. These strings constitute the argument list available to the new process image. The value in *argv*[0] should point to a filename that is associated with the process being started by one of the *exec* functions.

The const char *arg and subsequent ellipses in the execl(), execlp(), and exe-96 cle() functions can be thought of as arg0, arg1, ..., argn. Together they describe 97 a list of one or more pointers to null-terminated character strings that represent 98 the argument list available to the new program. The first argument should point 99 to a filename that is associated with the process being started by one of the exec 100 functions, and the last argument shall be a NULL pointer. For the *execle()* func-101 tion, the environment is provided by following the NULL pointer that shall ter-102 minate the list of arguments in the parameter list to execle() with an additional 103

ł

I

104 parameter, as if it were declared as

105 char *const envp[]

The argument *envp* to *execve()* and the final argument to *execle()* name an array of character pointers to null-terminated strings. These strings constitute the environment for the new process image. The environment array is terminated by a NULL pointer.

For those forms not containing an *envp* pointer [execl(), execv(), execlp(), and*execup()*], the environment for the new process image is taken from the externalvariable*environ*in the calling process.

The number of bytes available for the new process's combined argument and environment lists is {ARG_MAX}. The implementation shall specify in the system documentation (see 1.3.1.2) whether any combination of null terminators, pointers, or alignment bytes are included in this total.

- File descriptors open in the calling process image remain open in the new process image, except for those whose close-on-exec flag FD_CLOEXEC is set (see 6.5.2 and 6.5.1). For those file descriptors that remain open, all attributes of the open file description, including file locks (see 6.5.2), remain unchanged by this function call.
- Directory streams open in the calling process image shall be closed in the new process image.

Signals set to the default action (SIG_DFL) in the calling process image shall be set to the default action in the new process image. Signals set to be ignored (SIG_IGN) by the calling process image shall be set to be ignored by the new process image. Signals set to be caught by the calling process image shall be set to the default action in the new process image (see 3.3.1).

If the set-user-ID mode bit of the new process image file is set (see 5.6.4), the effec-129 tive user ID of the new process image is set to the owner ID of the new process 130 image file. Similarly, if the set-group-ID mode bit of the new process image file is 131 set, the effective group ID of the new process image is set to the group ID of the 132 new process image file. The real user ID, real group ID, and supplementary group 133 IDs of the new process image remain the same as those of the calling process 134 image. If {_POSIX_SAVED_IDS} is defined, the effective user ID and effective 135group ID of the new process image shall be saved (as the saved set-user-ID and the 136 saved set-group-ID) for use by the setuid() function. 137

- The new process image also inherits the following attributes from the calling process image:
- 140 (1) Process ID
- 141 (2) Parent process ID
- 142 (3) Process group ID
- (4) Session membership
- 144 (5) Real user ID
- 145 (6) Real group ID

- 146 (7) Supplementary group IDs
- 147 (8) Time left until an alarm clock signal (see 3.4.1)
- 148 (9) Current working directory
- 149 (10) Root directory
- 150 (11) File mode creation mask (see 5.3.3)
- (12) Process signal mask (see 3.3.5)
- 152 (13) Pending signals (see 3.3.6)
- 153 (14) tms_utime, tms_stime, tms_cutime, and tms_cstime (see 4.5.2)

All process attributes defined by this part of ISO/IEC 9945 and not specified in this subclause (3.1.2) shall be the same in the new and old process images. The inheritance of process characteristics not defined by this part of ISO/IEC 9945 is unspecified by this part of ISO/IEC 9945, but should be documented in the system documentation.

Upon successful completion, the *exec* functions shall mark for update the *st_atime* field of the file. If the *exec* function failed, but was able to locate the *process image file*, whether the *st_atime* field is marked for update is unspecified. Should the *exec* function succeed, the process image file shall be considered to have been *open()-ed.* The corresponding *close()* shall be considered to occur at a time after this open, but before process termination or successful completion of a subsequent call to one of the *exec* functions.

The argv[] and envp[] arrays of pointers and the strings to which those arrays
point shall not be modified by a call to one of the exec functions, except as a consequence of replacing the process image.

169 **3.1.2.3 Returns**

170 If one of the *exec* functions returns to the calling process image, an error has 171 occurred; the return value shall be -1, and *errno* shall be set to indicate the error.

172 **3.1.2.4 Errors**

173 If any of the following conditions occur, the *exec* functions shall return -1 and set 174 *errno* to the corresponding value:

- 175[E2BIG]The number of bytes used by the argument list and the environ-176ment list of the new process image is greater than the system-177imposed limit of {ARG_MAX} bytes.178[EACCES]Search permission is denied for a directory listed in the path
- 178[EACCES]Search permission is denied for a directory listed in the path179prefix of the new process image file, or the new process image180file denies execution permission, or the new process image file181is not a regular file and the implementation does not support182execution of files of its type.
- 183 [ENAMETOOLONG]

184

The length of the *path* or *file* arguments, or an element of the

185 186 187		environment variable PATH prefixed to a file, exceeds {PATH_MAX}, or a pathname component is longer than {NAME_MAX} and {_POSIX_NO_TRUNC} is in effect for that file.
188 189 190	[ENOENT]	One or more components of the pathname of the new process image file do not exist, or the <i>path</i> or <i>file</i> argument points to an empty string.

- 191 [ENOTDIR] A component of the path prefix of the new process image file is 192 not a directory.
- If any of the following conditions occur, the execl(), execv(), execle(), and execve()
 functions shall return -1 and set errno to the corresponding value:
- 195 [ENOEXEC] The new process image file has the appropriate access permis-196 sion, but is not in the proper format.

For each of the following conditions, if the condition is detected, the *exec* functions
shall return -1 and return the corresponding value in *errno*:

199[ENOMEM]The new process image requires more memory than is allowed200by the hardware or system-imposed memory management con-201straints.

202 3.1.2.5 Cross-References

203 alarm(), 3.4.1; chmod(), 5.6.4; _exit(), 3.2.2; fcntl(), 6.5.2; fork(), 3.1.1; setuid(), 204 4.2.2; <signal.h>, 3.3.1; sigprocmask(), 3.3.5; sigpending(), 3.3.6; stat(), 5.6.2; 205 <sys/stat.h>, 5.6.1; times(), 4.5.2; umask(), 5.3.3; 2.6.

206 **3.2 Process Termination**

- 207 There are two kinds of process termination:
- 208 (1) Normal termination occurs by a return from main() or when requested
 209 with the exit() or _exit() functions.
- (2) Abnormal termination occurs when requested by the abort() function or
 when some signals are received (see 3.3.1).

The *exit()* and *abort()* functions shall be as described in the C Standard {2}. Both | *exit()* and *abort()* shall terminate a process with the consequences specified in 3.2.2, except that the status made available to *wait()* or *waitpid()* by *abort()* shall be that of a process terminated by the SIGABRT signal.

A parent process can suspend its execution to wait for termination of a child process with the *wait()* or *waitpid()* functions.

218 3.2.1 Wait for Process Termination

- Functions: wait(), waitpid()
- 220 **3.2.1.1 Synopsis**
- 221 #include <sys/types.h>
- 222 #include <sys/wait.h>
- 223 pid_t wait(int *stat_loc);

224 pid_t waitpid(pid_t pid, int *stat_loc, int options);

225 **3.2.1.2 Description**

The wait() and waitpid() functions allow the calling process to obtain status information pertaining to one of its child processes. Various options permit status information to be obtained for child processes that have terminated or stopped. If status information is available for two or more child processes, the order in which their status is reported is unspecified.

The *wait()* function shall suspend execution of the calling process until status information for one of its terminated child processes is available, or until a signal whose action is either to execute a signal-catching function or to terminate the process is delivered. If status information is available prior to the call to *wait()*, return shall be immediate.

- The waitpid() function shall behave identically to the wait() function if the pid argument has a value of -1 and the options argument has a value of zero. Otherwise, its behavior shall be modified by the values of the pid and options arguments.
- The *pid* argument specifies a set of child processes for which status is requested. The *waitpid()* function shall only return the status of a child process from this set.
- (1) If pid is equal to -1, status is requested for any child process. In this respect, waitpid() is then equivalent to wait().
- (2) If *pid* is greater than zero, it specifies the process ID of a single child process for which status is requested.
- (3) If *pid* is equal to zero, status is requested for any child process whose
 process group ID is equal to that of the calling process.
- (4) If *pid* is less than -1, status is requested for any child process whose process group ID is equal to the absolute value of *pid*.
- The options argument is constructed from the bitwise inclusive OR of zero or more of the following flags, defined in the header <sys/wait.h>:
- 253WNOHANGThe waitpid() function shall not suspend execution of the cal-254ling process if status is not immediately available for one of the255child processes specified by pid.
- WUNTRACED If the implementation supports job control, the status of any child processes specified by *pid* that are stopped, and whose

status has not yet been reported since they stopped, shall also be reported to the requesting process.

If wait() or waitpid() return because the status of a child process is available, 260 these functions shall return a value equal to the process ID of the child process. 261 In this case, if the value of the argument *stat_loc* is not NULL, information shall 262 be stored in the location pointed to by stat_loc. If and only if the status returned 263 is from a terminated child process that returned a value of zero from main() or 264 passed a value of zero as the status argument to _exit() or exit(), the value stored 265 at the location pointed to by stat loc shall be zero. Regardless of its value, this 266 information may be interpreted using the following macros, which are defined in 267 268 <sys/wait.h> and evaluate to integral expressions; the stat_val argument is the integer value pointed to by *stat_loc*. 269

WIFEXITED(stat_val) 270

271 272

283

284

285

287

288

This macro evaluates to a nonzero value if status was returned for a child process that terminated normally.

WEXITSTATUS(stat_val) 273

If the value of WIFEXITED(stat_val) is nonzero, this macro 274 evaluates to the low-order 8 bits of the status argument that 275the child process passed to _exit() or exit(), or the value the 276 child process returned from *main()*. 277

WIFSIGNALED(stat val) 278

This macro evaluates to a nonzero value if status was returned 279 for a child process that terminated due to the receipt of a signal 280 that was not caught (see 3.3.1). 281

WTERMSIG(stat val) 282

If the value of WIFSIGNALED(stat_val) is nonzero, this macro evaluates to the number of the signal that caused the termination of the child process.

WIFSTOPPED(*stat_val*) 286

This macro evaluates to a nonzero value if status was returned for a child process that is currently stopped.

WSTOPSIG(stat_val) 289

If the value of WIFSTOPPED(stat_val) is nonzero, this macro 290 evaluates to the number of the signal that caused the child pro-291 cess to stop. 292

If the information stored at the location pointed to by stat loc was stored there by 293 294 a call to the *waitpid()* function that specified the WUNTRACED flag, exactly of the macros WIFEXITED(*stat_loc), WIFSIGNALED(*stat_loc), one or 295 WIFSTOPPED(*stat_loc) shall evaluate to a nonzero value. If the information 296 stored at the location pointed to by *stat_loc* was stored there by a call to the *wait*-297 pid() function that did not specify the WUNTRACED flag or by a call to the 298 function, exactly one of the macros WIFEXITED(*stat_loc) or wait() 299WIFSIGNALED(*stat loc) shall evaluate to a nonzero value. 300

An implementation may define additional circumstances under which wait() or waitpid() reports status. This shall not occur unless the calling process or one of its child processes explicitly makes use of a nonstandard extension. In these cases, the interpretation of the reported status is implementation defined.

305

306 **3.2.1.3 Returns**

If the *wait()* or *waitpid()* functions return because the status of a child process is 307 available, these functions shall return a value equal to the process ID of the child 308 process for which status is reported. If the wait() or waitpid() functions return 309 due to the delivery of a signal to the calling process, a value of -1 shall be 310 returned and errno shall be set to [EINTR]. If the waitpid() function was invoked 311 with WNOHANG set in options, has at least one child process specified by pid for 312 which status is not available, and status is not available for any process specified 313 by pid, a value of zero shall be returned. Otherwise, a value of -1 shall be 314 returned, and *errno* shall be set to indicate the error. 315

316 **3.2.1.4 Errors**

If any of the following conditions occur, the wait() function shall return -1 and set
 errno to the corresponding value:

- 319 [ECHILD] The calling process has no existing unwaited-for child 320 processes.
- 321 [EINTR] The function was interrupted by a signal. The value of the 322 location pointed to by *stat_loc* is undefined.
- If any of the following conditions occur, the waitpid() function shall return -1 and set *errno* to the corresponding value:
- 325 [ECHILD] The process or process group specified by *pid* does not exist or 326 is not a child of the calling process.
- 327[EINTR]The function was interrupted by a signal. The value of the
location pointed to by stat_loc is undefined.
- 329 [EINVAL] The value of the options argument is not valid.

330 3.2.1.5 Cross-References

- 331 _exit(), 3.2.2; fork(), 3.1.1; pause(), 3.4.2; times(), 4.5.2; <signal.h>, 3.3.1.
- 332 **3.2.2 Terminate a Process**
- 333 Function: _exit()

334 **3.2.2.1 Synopsis**

335 void _exit(int status);

336 **3.2.2.2 Description**

- The _exit() function shall terminate the calling process with the following consequences:
- (1) All open file descriptors and directory streams in the calling process are closed.
- (2) If the parent process of the calling process is executing a wait() or wait *pid()*, it is notified of the termination of the calling process and the low
 order 8 bits of status are made available to it; see 3.2.1.
- (3) If the parent process of the calling process is not executing a *wait()* or
 waitpid() function, the exit *status* code is saved for return to the parent
 process whenever the parent process executes an appropriate subsequent
 wait() or *waitpid()*.
- (4) Termination of a process does not directly terminate its children. The sending of a SIGHUP signal as described below indirectly terminates children in some circumstances. Children of a terminated process shall be assigned a new parent process ID, corresponding to an implementation-defined system process.
- (5) If the implementation supports the SIGCHLD signal, a SIGCHLD signal
 shall be sent to the parent process.
- (6) If the process is a controlling process, the SIGHUP signal shall be sent to
 each process in the foreground process group of the controlling terminal
 belonging to the calling process.
- (7) If the process is a controlling process, the controlling terminal associated
 with the session is disassociated from the session, allowing it to be
 acquired by a new controlling process.
- (8) If the implementation supports job control, and if the exit of the process causes a process group to become orphaned, and if any member of the newly orphaned process group is stopped, then a SIGHUP signal followed by a SIGCONT signal shall be sent to each process in the newly orphaned process group.
- 366 These consequences shall occur on process termination for any reason.

367 **3.2.2.3 Returns**

- 368 The _exit() function cannot return to its caller.
- 369 **3.2.2.4 Cross-References**
- 370 *close()*, 6.3.1; *sigaction()*, 3.3.4; *wait*, 3.2.1.
371 **3.3 Signals**

372 3.3.1 Signal Concepts

373 **3.3.1.1 Signal Names**

The <signal.h> header declares the *sigset_t* type and the *sigaction* structure. It also defines the following symbolic constants, each of which expands to a distinct constant expression of the type void(*)(), whose value matches no declarable function.

78 Symbolic 79 <u>Constant</u>		Description	
380	SIG_DFL	Request for default signal handling	
381	SIG_IGN	Request that signal be ignored	

The type *sigset_t* is used to represent sets of signals. It is always an integral or structure type. Several functions used to manipulate objects of type *sigset_t* are defined in 3.3.3.

The <signal.h> header also declares the constants that are used to refer to the signals that occur in the system. Each of the signals defined by this part of ISO/IEC 9945 and supported by the implementation shall have distinct, positive integral values. The value zero is reserved for use as the null signal (see 3.3.2). An implementation may define additional signals that may occur in the system.

390 The constants shown in Table 3-1 shall be supported by all implementations.

The constants shown in Table 3-2 shall be defined by all implementations. However, implementations that do not support job control are not required to support these signals. If these signals are supported by the implementation, they shall behave in accordance with this part of ISO/IEC 9945. Otherwise, the implementation shall not generate these signals, and attempts to send these signals or to examine or specify their actions shall return an error condition. See 3.3.2 and 3.3.4.

398 3.3.1.2 Signal Generation and Delivery

A signal is said to be *generated* for (or sent to) a process when the event that causes the signal first occurs. Examples of such events include detection of hardware faults, timer expiration, and terminal activity, as well as the invocation of the *kill*() function. In some circumstances, the same event generates signals for multiple processes.

Each process has an action to be taken in response to each signal defined by the system (see 3.3.1.3). A signal is said to be *delivered* to a process when the appropriate action for the process and signal is taken.

During the time between the generation of a signal and its delivery, the signal is said to be *pending*. Ordinarily, this interval cannot be detected by an application. However, a signal can be *blocked* from delivery to a process. If the action associated with a blocked signal is anything other than to ignore the signal, and if that

SymbolicDefaultConstantAction		Description		
SIGABRT	1	Abnormal termination signal, such as is initiated by the <i>abort()</i> function (as defined in the C Standard {2}).		
SIGALRM	1	Timeout signal, such as initiated by the <i>alarm()</i> function (see 3.4.1).		
SIGFPE	1	Erroneous arithmetic operation, such as division by zero or an operation resulting in overflow.		
SIGHUP	1	Hangup detected on controlling terminal (see 7.1.1.10) or death of controlling process (see 3.2.2).		
SIGILL	1	Detection of an invalid hardware instruction.		
SIGINT	1	Interactive attention signal (see 7.1.1.9).		
SIGKILL	1	Termination signal (cannot be caught or ignored).		
SIGPIPE	1	Write on a pipe with no readers (see 6.4.2).		
SIGQUIT	1	Interactive termination signal (see 7.1.1.9).		
SIGSEGV	1	Detection of an invalid memory reference.		
SIGTERM	1	Termination signal.		
SIGUSR1	1	Reserved as application-defined signal 1.		
SIGUSR2	1	Reserved as application-defined signal 2.		

Table 3-1 - Required Signals

NOTE: The default actions are 432

433

434

Abnormal termination of the process. 1

Symbolic Constant	Default Action	Description
SIGCHLD	2	Child process terminated or stopped.
SIGCONT	4	Continue if stopped.
SIGSTOP	3	Stop signal (cannot be caught or ignored).
SIGTSTP	3	Interactive stop signal (see 7.1.1.9).
SIGTTIN	3	Read from control terminal attempted by a member of a background process group (see 7.1.1.4).
SIGTTOU	3	Write to control terminal attempted by a member of a background proces group (see $7.1.1.4$).

Table 3-2 - Job Control Signals

- 44
- Ignore the signal. 448 2
- Stop the process. 449 3
- Continue the process if it is currently stopped; otherwise, ignore the signal. 450 4

signal is generated for the process, the signal shall remain pending until either it
is unblocked or the action associated with it is set to ignore the signal. If the
action associated with a blocked signal is to ignore the signal, and if that signal is
generated for the process, it is unspecified whether the signal is discarded
immediately upon generation or remains pending.

Each process has a *signal mask* that defines the set of signals currently blocked from delivery to it. The signal mask for a process is initialized from that of its parent. The *sigaction()*, *sigprocmask()*, and *sigsuspend()* functions control the manipulation of the signal mask.

- The determination of which action is taken in response to a signal is made at the time the signal is delivered, allowing for any changes since the time of generation. This determination is independent of the means by which the signal was originally generated. If a subsequent occurrence of a pending signal is generated, it is implementation defined as to whether the signal is delivered more than once. The order in which multiple, simultaneously pending signals are delivered to a process is unspecified.
- When any stop signal (SIGSTOP, SIGTSTP, SIGTTIN, SIGTTOU) is generated for a process, any pending SIGCONT signals for that process shall be discarded. Conversely, when SIGCONT is generated for a process, all pending stop signals for that process shall be discarded. When SIGCONT is generated for a process that is stopped, the process shall be continued, even if the SIGCONT signal is blocked or ignored. If SIGCONT is blocked and not ignored, it shall remain pending until it is either unblocked or a stop signal is generated for the process.
- An implementation shall document any conditions not specified by this part of ISO/IEC 9945 under which the implementation generates signals. (See 1.3.1.2.)

476 **3.3.1.3 Signal Actions**

There are three types of actions that can be associated with a signal: SIG_DFL, SIG_IGN, or a *pointer to a function*. Initially, all signals shall be set to SIG_DFL or SIG_IGN prior to entry of the *main()* routine (see 3.1.2). The actions prescribed by these values are as follows:

- 481 (1) SIG_DFL signal-specific default action
- (a) The default actions for the signals defined in this part of
 ISO/IEC 9945 are specified in Table 3-1 and Table 3-2.
- If the default action is to stop the process, the execution of that pro-(b) 484 cess is temporarily suspended. When a process stops, a SIGCHLD 485 signal shall be generated for its parent process, unless the parent 486 process has set the SA_NOCLDSTOP flag (see 3.3.4). While a process 487 is stopped, any additional signals that are sent to the process shall 488 not be delivered until the process is continued except SIGKILL, 489 which always terminates the receiving process. A process that is a 490 member of an orphaned process group shall not be allowed to stop 491 in response to the SIGTSTP, SIGTTIN, or SIGTTOU signals. In cases 492 where delivery of one of these signals would stop such a process, the 493 494 signal shall be discarded.

506

507

508

512

513

514

515

518

519

520

521

522

523

524

525

526

527

528

529

530

- (c) Setting a signal action to SIG_DFL for a signal that is pending, and
 whose default action is to ignore the signal (for example, SIGCHLD),
 shall cause the pending signal to be discarded, whether or not it is
 blocked.
- 499 (2) SIG_IGN ignore signal
- 500(a) Delivery of the signal shall have no effect on the process. The
behavior of a process is undefined after it ignores a SIGFPE, SIGILL,
or SIGSEGV signal that was not generated by the kill() function or
the raise() function defined by the C Standard {2}.
- 504(b) The system shall not allow the action for the signals SIGKILL or505SIGSTOP to be set to SIG_IGN.
 - (c) Setting a signal action to SIG_IGN for a signal that is pending shall cause the pending signal to be discarded, whether or not it is blocked.
- 509(d) If a process sets the action for the SIGCHLD signal to SIG_IGN, the510behavior is unspecified.
- 511 (3) *pointer to a function* catch signal
 - (a) On delivery of the signal, the receiving process is to execute the signal-catching function at the specified address. After returning from the signal-catching function, the receiving process shall resume execution at the point at which it was interrupted.
- 516 (b) The signal-catching function shall be entered as a C language func-517 tion call as follows:

void func (int signo);

where *func* is the specified signal-catching function and *signo* is the signal number of the signal being delivered.

- (c) The behavior of a process is undefined after it returns normally from a signal-catching function for a SIGFPE, SIGILL, or SIGSEGV signal that was not generated by the *kill()* function or the *raise()* function defined by the C Standard {2}.
 - (d) The system shall not allow a process to catch the signals SIGKILL and SIGSTOP.
- (e) If a process establishes a signal-catching function for the SIGCHLD signal while it has a terminated child process for which it has not waited, it is unspecified whether a SIGCHLD signal is generated to indicate that child process.
- 531(f)When signal-catching functions are invoked asynchronously with532process execution, the behavior of some of the functions defined by533this part of ISO/IEC 9945 is unspecified if they are called from a534signal-catching function. The following table defines a set of func-535tions that shall be reentrant with respect to signals (that is, appli-536cations may invoke them, without restriction, from signal-catching537functions).

1

538	_exit()	fstat()	read()	sysconf()
539	access()	getegid()	rename()	tcdrain()
540	alarm()	geteuid()	rmdir()	tcflow()
541	cfgetispeed()	getgid()	setgid()	tcflush()
542	cfgetospeed()	getgroups()	setpgid()	tcgetattr()
543	cfsetispeed()	getpgrp()	setsid()	tcgetpgrp()
544	cfsetospeed()	getpid()	setuid()	<pre>tcsendbreak()</pre>
545	chdir()	getppid()	sigaction()	tcsetattr()
546	chmod()	getuid()	sigaddset()	tcsetpgrp()
547	chown()	kill()	sigdelset()	time()
548	close()	link()	sigemptyset()	times()
549	creat()	lseek()	sigfillset()	umask()
550	<i>dup2</i> ()	mkdir()	sigismember()	uname()
551	dup()	mkfifo()	sigpending()	unlink()
552	execle()	open()	sigprocmask()	utime()
553	execve()	pathconf()	sigsuspend()	wait()
554	fcntl()	pause()	sleep()	waitpid()
555	fork()	pipe()	stat()	write()

All POSIX.1 functions not in the preceding table and all functions 556 defined in the C Standard {2} not stated to be callable from a 557 signal-catching function are considered to be *unsafe* with respect to 558 signals. In the presence of signals, all functions defined by this part 559 of ISO/IEC 9945 or by the C Standard {2} shall behave as defined (by 560 the defining standard) when called from or interrupted by a signal-561 catching function, with a single exception: when a signal interrupts 562 an unsafe function and the signal-catching function calls an unsafe 563 function, the behavior is undefined. 564

3.3.1.4 Signal Effects on Other Functions

Signals affect the behavior of certain functions defined by this part of 566 ISO/IEC 9945 if delivered to a process while it is executing such a function. If the 567 action of the signal is to terminate the process, the process shall be terminated 568 and the function shall not return. If the action of the signal is to stop the process, 569 the process shall stop until continued or terminated. Generation of a SIGCONT 570 signal for the process causes the process to be continued, and the original function 571 shall continue at the point where the process was stopped. If the action of the sig-572 nal is to invoke a signal-catching function, the signal-catching function shall be 573 invoked; in this case, the original function is said to be *interrupted* by the signal. 574 If the signal-catching function executes a return, the behavior of the interrupted 575 function shall be as described individually for that function. Signals that are 576 ignored shall not affect the behavior of any function; signals that are blocked shall 577 not affect the behavior of any function until they are delivered. 578

I

I

I

L

1

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

579 3.3.2 Send a Signal to a Process

- 580 Function: *kill()*
- 581 **3.3.2.1 Synopsis**

582 #include <sys/types.h>

- 583 #include <signal.h>
- 584 int kill(pid_t pid, int sig);

585 **3.3.2.2 Description**

The *kill()* function shall send a signal to a process or a group of processes specified by *pid*. The signal to be sent is specified by *sig* and is either one from the list given in 3.3.1.1 or zero. If *sig* is zero (the null signal), error checking is performed, but no signal is actually sent. The null signal can be used to check the validity of *pid*.

For a process to have permission to send a signal to a process designated by *pid*, the real or effective user ID of the sending process must match the real or effective user ID of the receiving process, unless the sending process has appropriate privileges. If (_POSIX_SAVED_IDS) is defined, the saved set-user-ID of the receiving process shall be checked in place of its effective user ID.

- If *pid* is greater than zero, *sig* shall be sent to the process whose process ID is equal to *pid*.
- If *pid* is zero, *sig* shall be sent to all processes (excluding an unspecified set of system processes) whose process group ID is equal to the process group ID of the sender and for which the process has permission to send a signal.
- 601 If pid is -1, the behavior of the kill() function is unspecified.
- ⁶⁰² If *pid* is negative, but not -1, *sig* shall be sent to all processes (excluding an ⁶⁰³ unspecified set of system processes) whose process group ID is equal to the abso-⁶⁰⁴ lute value of *pid* and for which the process has permission to send a signal.
- If the value of *pid* causes *sig* to be generated for the sending process, and if *sig* is not blocked, either *sig* or at least one pending unblocked signal shall be delivered to the sending process before the kill() function returns.
- If the implementation supports the SIGCONT signal, the user ID tests described
 above shall not be applied when sending SIGCONT to a process that is a member
 of the same session as the sending process.
- An implementation that provides extended security controls may impose further implementation-defined restrictions on the sending of signals, including the null signal. In particular, the system may deny the existence of some or all of the processes specified by *pid*.
- The kill() function is successful if the process has permission to send *sig* to any of the processes specified by *pid*. If the kill() function fails, no signal shall be sent.

617 **3.3.2.3 Returns**

⁶¹⁸ Upon successful completion, the function shall return a value of zero. Otherwise, ⁶¹⁹ a value of -1 shall be returned and *errno* shall be set to indicate the error.

620 **3.3.2.4 Errors**

- If any of the following conditions occur, the kill() function shall return -1 and set *errno* to the corresponding value:
- 623 [EINVAL] The value of the *sig* argument is an invalid or unsupported sig-624 nal number.
- 625 [EPERM] The process does not have permission to send the signal to any 626 receiving process.
- 627[ESRCH]No process or process group can be found corresponding to that628specified by *pid*.

629 3.3.2.5 Cross-References

630 getpid(), 4.1.1; setsid(), 4.3.2; sigaction(), 3.3.4; <signal.h>, 3.3.1.

631 **3.3.3 Manipulate Signal Sets**

632 Functions: sigemptyset(), sigfillset(), sigaddset(), sigdelset(), sigismember()

633 **3.3.3.1 Synopsis**

- 634 #include <signal.h>
- 635 int sigemptyset(sigset_t *set);
- 636 int sigfillset(sigset_t *set);
- 637 int sigaddset(sigset_t *set, int signo);
- 638 int sigdelset(sigset_t *set, int signo);
- 639 int sigismember(const sigset_t *set, int signo);

640 **3.3.3.2 Description**

- The sigsetops primitives manipulate sets of signals. They operate on data objects
 addressable by the application, not on any set of signals known to the system,
 such as the set blocked from delivery to a process or the set pending for a process
 (see 3.3.1).
- The *sigemptyset()* function initializes the signal set pointed to by the argument set, such that all signals defined in this part of ISO/IEC 9945 are excluded.
- The sigfillset() function initializes the signal set pointed to by the argument set,
 such that all signals defined in this part of ISO/IEC 9945 are included.
- Applications shall call either *sigemptyset()* or *sigfillset()* at least once for each object of type *sigset_t* prior to any other use of that object. If such an object is not

ľ

initialized in this way, but is nonetheless supplied as an argument to any of the
 sigaddset(), sigdelset(), sigismember(), sigaction(), sigprocmask(), sigpending(), or
 sigsuspend() functions, the results are undefined.

The *sigaddset()* and *sigdelset()* functions respectively add or delete the individual signal specified by the value of the argument *signo* to or from the signal set pointed to by the argument *set*.

The *sigismember()* function tests whether the signal specified by the value of the argument *signo* is a member of the set pointed to by the argument *set*.

659 **3.3.3.3 Returns**

660 Upon successful completion, the *sigismember()* function returns a value of one if 661 the specified signal is a member of the specified set, or a value of zero if it is not. 662 Upon successful completion, the other functions return a value of zero. For all of 663 the above functions, if an error is detected, a value of -1 is returned, and *errno* is 664 set to indicate the error.

665 **3.3.3.4 Errors**

For each of the following conditions, if the condition is detected, the *sigaddset()*, *sigdelset()*, and *sigismember()* functions shall return -1 and set *errno* to the corresponding value:

669 [EINVAL] The value of the *signo* argument is an invalid or unsupported 670 signal number.

671 3.3.3.5 Cross-References

sigaction(), 3.3.4; <signal.h>, 3.3.1; sigpending(), 3.3.6; sigprocmask(), 3.3.5; sigsuspend(), 3.3.7.

674 3.3.4 Examine and Change Signal Action

- 675 Function: sigaction()
- 676 **3.3.4.1 Synopsis**
- 677 #include <signal.h>

```
678 int sigaction(int sig, const struct sigaction *act,
679 struct sigaction *oact);
```

680 **3.3.4.2 Description**

The *sigaction()* function allows the calling process to examine or specify (or both) the action to be associated with a specific signal. The argument *sig* specifies the signal; acceptable values are defined in 3.3.1.1.

684 The structure *sigaction*, used to describe an action to be taken, is defined in the

l

ł

685 header < signal. h> to include at least the following members:

686 687	Member <u>Type</u>	Member Name	Description
688	<i>void</i> (*)()	sa_handler	SIG_DFL, SIG_IGN, or pointer to a function.
689	sigset_t	sa_mask	Additional set of signals to be blocked during
690			execution of signal-catching function.
691	int	sa_flags	Special flags to affect behavior of signal.

Implementations may add extensions as permitted in 1.3.1.1, point (2). Adding
extensions to this structure, which might change the behavior of the application
with respect to this standard when those fields in the structure are uninitialized,
also requires that the extensions be enabled as required by 1.3.1.1.

If the argument act is not NULL, it points to a structure specifying the action to 696 be associated with the specified signal. If the argument oact is not NULL, the 697 action previously associated with the signal is stored in the location pointed to by 698 the argument oact. If the argument act is NULL, signal handling is unchanged by 699 this function call; thus, the call can be used to enquire about the current handling 700 of a given signal. The sa_handler field of the signation structure identifies the 701 action to be associated with the specified signal. If the sa handler field specifies a 702 signal-catching function, the *sa_mask* field identifies a set of signals that shall be 703 added to the signal mask of the process before the signal-catching function is 704 invoked. The SIGKILL and SIGSTOP signals shall not be added to the signal mask 705 using this mechanism; this restriction shall be enforced by the system without 706 causing an error to be indicated. 707

708 The *sa_flags* field can be used to modify the behavior of the specified signal.

- The following flag bit, defined in the header <signal.h>, can be set in sa_flags:
- 710
 Symbolic
 Description

 711
 Constant
 Do not generate SIGCHLD when children stop.

If sig is SIGCHLD and the SA_NOCLDSTOP flag is not set in sa_flags, and the
implementation supports the SIGCHLD signal, a SIGCHLD signal shall be generated for the calling process whenever any of its child processes stop. If sig is
SIGCHLD and the SA_NOCLDSTOP flag is set in sa_flags, the implementation
shall not generate a SIGCHLD signal in this way.

718 When a signal is caught by a signal-catching function installed by the sigaction()719 function, a new signal mask is calculated and installed for the duration of the 720 signal-catching function [or until a call to either the sigprocmask() or sig-721 suspend() function is made]. This mask is formed by taking the union of the 722 current signal mask and the value of the sa_mask for the signal being delivered, 723 and then including the signal being delivered. If and when the user's signal 724 handler returns normally, the original signal mask is restored.

Once an action is installed for a specific signal, it remains installed until another action is explicitly requested [by another call to the *sigaction()* function] or until one of the *exec* functions is called. If the previous action for sig had been established by the signal() function, defined in the C Standard (2), the values of the fields returned in the structure pointed to by oact are unspecified and, in particular, $oact->sv_handler$ is not necessarily the same value passed to the signal() function. However, if a pointer to the same structure or a copy thereof is passed to a subsequent call to the sigaction() function via the act argument, handling of the signal shall be as if the original call to the signal() function were repeated.

⁷³⁵ If the *sigaction()* function fails, no new signal handler is installed.

It is unspecified whether an attempt to set the action for a signal that cannot be
 caught or ignored to SIG_DFL is ignored or causes an error to be returned with
 errno set to [EINVAL].

739 **3.3.4.3 Returns**

⁷⁴⁰ Upon successful completion, a value of zero is returned. Otherwise, a value of -1⁷⁴¹ is returned and *errno* is set to indicate the error.

- 742 **3.3.4.4 Errors**
- If any of the following conditions occur, the sigaction() function shall return -1and set *errno* to the corresponding value:

745[EINVAL]The value of the sig argument is an invalid or unsupported sig-746nal number, or an attempt was made to catch a signal that can-747not be caught or to ignore a signal that cannot be ignored. See7483.3.1.1.

For each of the following conditions, when the condition is detected and the implementation treats it as an error, the sigaction() function shall return a value of -1and set *errno* to the corresponding value.

[EINVAL] An attempt was made to set the action to SIG_DFL for a signal
 that cannot be caught or ignored (or both).

- 754 3.3.4.5 Cross-References
- 755 kill(), 3.3.2; <signal.h>, 3.3.1; sigprocmask(), 3.3.5; sigsetops, 3.3.3; sig-756 suspend(), 3.3.7.

757 3.3.5 Examine and Change Blocked Signals

- 758 Function: sigprocmask()
- 759 **3.3.5.1 Synopsis**
- 760 #include <signal.h>
- 761 int sigprocmask(int how, const sigset_t *set, sigset_t *oset);

I

I

762 **3.3.5.2 Description**

The *sigprocmask()* function is used to examine or change (or both) the signal mask of the calling process. If the value of the argument *set* is not NULL, it points to a set of signals to be used to change the currently blocked set.

The value of the argument *how* indicates the manner in which the set is changed and shall consist of one of the following values, as defined in the header <signal.h>:

769	Name	Description
770	SIG_BLOCK	The resulting set shall be the union of the current set and
771		the signal set pointed to by the argument set.
772	SIG_UNBLOCK	The resulting set shall be the intersection of the current set
773		and the complement of the signal set pointed to by the argu-
774		ment set.
775	SIG_SETMASK	The resulting set shall be the signal set pointed to by the
776		argument set.

If the argument *oset* is not NULL, the previous mask is stored in the space pointed to by *oset*. If the value of the argument *set* is NULL, the value of the argument *how* is not significant and the signal mask of the process is unchanged by this function call; thus, the call can be used to enquire about currently blocked signals.

- If there are any pending unblocked signals after the call to the sigprocmask()
 function, at least one of those signals shall be delivered before the sigprocmask()
 function returns.
- It is not possible to block the SIGKILL and SIGSTOP signals; this shall be enforced
 by the system without causing an error to be indicated.

If any of the SIGFPE, SIGILL, or SIGSEGV signals are generated while they are
blocked, the result is undefined, unless the signal was generated by a call to the *kill()* function or the *raise()* function defined by the C Standard {2}.

- If the sigprocmask() function fails, the signal mask of the process is not changed
 by this function call.
- 792 **3.3.5.3 Returns**
- ⁷⁹³ Upon successful completion a value of zero is returned. Otherwise, a value of -1⁷⁹⁴ is returned and *errno* is set to indicate the error.
- 795 **3.3.5.4 Errors**
- If any of the following conditions occur, the sigprocmask() function shall return -1
 and set errno to the corresponding value:
- 798 [EINVAL] The value of the *how* argument is not equal to one of the 799 defined values.

800 3.3.5.5 Cross-References

sigaction(), 3.3.4; <signal.h>, 3.3.1; sigpending(), 3.3.6; sigsetops, 3.3.3; sigsuspend(), 3.3.7.

803 3.3.6 Examine Pending Signals

804 Function: *sigpending()*

805 **3.3.6.1 Synopsis**

806 #include <signal.h>

807 int sigpending(sigset_t *set);

808 3.3.6.2 Description

The *sigpending()* function shall store the set of signals that are blocked from delivery and pending for the calling process in the space pointed to by the argument *set*.

812 **3.3.6.3 Returns**

⁸¹³ Upon successful completion, a value of zero is returned. Otherwise, a value of -1⁸¹⁴ is returned and *errno* is set to indicate the error.

815 **3.3.6.4 Errors**

This part of ISO/IEC 9945 does not specify any error conditions that are required to be detected for the *sigpending()* function. Some errors may be detected under conditions that are unspecified by this part of ISO/IEC 9945.

819 3.3.6.5 Cross-References

signal.h>, 3.3.1; sigprocmask(), 3.3.5; sigsetops, 3.3.3.

821 3.3.7 Wait for a Signal

- 822 Function: sigsuspend()
- 823 **3.3.7.1 Synopsis**
- 824 #include <signal.h>
- 825 int sigsuspend(const sigset_t *sigmask);

826 **3.3.7.2 Description**

- The *sigsuspend()* function replaces the signal mask of the process with the set of signals pointed to by the argument *sigmask* and then suspends the process until delivery of a signal whose action is either to execute a signal-catching function or to terminate the process.
- If the action is to terminate the process, the *sigsuspend()* function shall not return. If the action is to execute a signal-catching function, the *sigsuspend()* shall return after the signal-catching function returns, with the signal mask restored to the set that existed prior to the *sigsuspend()* call.
- It is not possible to block those signals that cannot be ignored, as documented in 3.3.1; this shall be enforced by the system without causing an error to be indicated.

838 **3.3.7.3 Returns**

Since the sigsuspend() function suspends process execution indefinitely, there is no successful completion return value. A value of -1 is returned and *errno* is set to indicate the error.

842 3.3.7.4 Errors

- If any of the following conditions occur, the sigsuspend() function shall return -1and set *errno* to the corresponding value:
- 845 [EINTR] A signal is caught by the calling process, and control is 846 returned from the signal-catching function.

847 3.3.7.5 Cross-References

848 pause(), 3.4.2; sigaction(), 3.3.4; <signal.h>, 3.3.1; sigpending(), 3.3.6; sigproc-849 mask(), 3.3.5; sigsetops, 3.3.3.

850 **3.4 Timer Operations**

A process can suspend itself for a specific period of time with the *sleep()* function or suspend itself indefinitely with the *pause()* function until a signal arrives. The *alarm()* function schedules a signal to arrive at a specific time, so a *pause()* suspension need not be indefinite.

855 3.4.1 Schedule Alarm

856 Function: alarm()

857 3.4.1.1 Synopsis

858 unsigned int alarm(unsigned int seconds);

859 **3.4.1.2 Description**

The *alarm()* function shall cause the system to send the calling process a SIGALRM signal after the number of real-time seconds specified by *seconds* have elapsed.

Processor scheduling delays may cause the process actually not to begin handling the signal until after the desired time.

Alarm requests are not stacked; only one SIGALRM generation can be scheduled in this manner. If the SIGALRM has not yet been generated, the call will result in rescheduling the time at which the SIGALRM will be generated.

868 If seconds is zero, any previously made *alarm()* request is canceled.

869 **3.4.1.3 Returns**

If there is a previous *alarm()* request with time remaining, the *alarm()* function shall return a nonzero value that is the number of seconds until the previous request would have generated a SIGALRM signal. Otherwise, the *alarm()* function shall return zero.

- 874 **3.4.1.4 Errors**
- The *alarm()* function is always successful, and no return value is reserved to indicate an error.
- 877 **3.4.1.5 Cross-References**
- exec, 3.1.2; fork(), 3.1.1; pause(), 3.4.2; sigaction(), 3.3.4; <signal.h>, 3.3.1.
- 879 **3.4.2 Suspend Process Execution**
- 880 Function: pause()
- 881 **3.4.2.1 Synopsis**
- 882 int pause(void);

883 **3.4.2.2 Description**

The *pause()* function suspends the calling process until delivery of a signal whose action is either to execute a signal-catching function or to terminate the process.

886 If the action is to terminate the process, the *pause()* function shall not return.

If the action is to execute a signal-catching function, the *pause()* function shall return after the signal-catching function returns. I

889 **3.4.2.3 Returns**

Since the pause() function suspends process execution indefinitely, there is no successful completion return value. A value of -1 is returned and *errno* is set to indicate the error.

893 **3.4.2.4 Errors**

If any of the following conditions occur, the *pause()* function shall return -1 and set *errno* to the corresponding value:

- 896 [EINTR] A signal is caught by the calling process, and control is 897 returned from the signal-catching function.
- 898 3.4.2.5 Cross-References
- *alarm()*, 3.4.1; *kill()*, 3.3.2; *wait*, 3.2.1; 3.3.1.4.
- 900 3.4.3 Delay Process Execution
- 901 Function: sleep()
- 902 3.4.3.1 Synopsis
- 903 unsigned int sleep(unsigned int seconds);

904 **3.4.3.2 Description**

The *sleep()* function shall cause the current process to be suspended from execution until either the number of real-time seconds specified by the argument *seconds* have elapsed or a signal is delivered to the calling process and its action is to invoke a signal-catching function or to terminate the process. The suspension time may be longer than requested due to the scheduling of other activity by the system.

If a SIGALRM signal is generated for the calling process during execution of the *sleep()* function and the SIGALRM signal is being ignored or blocked from delivery,
it is unspecified whether *sleep()* returns when the SIGALRM signal is scheduled.
If the signal is being blocked, it is also unspecified whether it remains pending
after the *sleep()* function returns or is discarded.

916 If a SIGALRM signal is generated for the calling process during execution of the 917 sleep() function, except as a result of a prior call to the alarm() function, and if 918 the SIGALRM signal is not being ignored or blocked from delivery, it is unspecified 919 whether that signal has any effect other than causing the sleep() function to 920 return.

If a signal-catching function interrupts the sleep() function and either examines
or changes the time a SIGALRM is scheduled to be generated, the action associated with the SIGALRM signal, or whether the SIGALRM signal is blocked from
delivery, the results are unspecified.

L

If a signal-catching function interrupts the *sleep()* function and calls the *siglongjmp()* or *longjmp()* function to restore an environment saved prior to the *sleep()* call, the action associated with the SIGALRM signal and the time at which a SIGALRM signal is scheduled to be generated are unspecified. It is also unspecified whether the SIGALRM signal is blocked, unless the process's signal mask is restored as part of the environment (see 8.3.1).

931 3.4.3.3 Returns

If the *sleep()* function returns because the requested time has elapsed, the value returned shall be zero. If the *sleep()* function returns due to delivery of a signal, the value returned shall be the unslept amount (the requested time minus the time actually slept) in seconds.

936 **3.4.3.4 Errors**

The *sleep()* function is always successful, and no return value is reserved to indicate an error.

- 939 3.4.3.5 Cross-References
- 940 alarm(), 3.4.1; pause(), 3.4.2; sigaction(), 3.3.4.

Section 4: Process Environment

1 4	4.1	Process	Identification
-----	-----	---------	----------------

- 2 4.1.1 Get Process and Parent Process IDs
- 3 Functions: getpid(), getppid()

4 **4.1.1.1 Synopsis**

- 5 #include <sys/types.h>
- 6 pid_t getpid(void);
- 7 pid_t getppid(void);

8 4.1.1.2 Description

- 9 The *getpid()* function returns the process ID of the calling process.
- 10 The *getppid()* function returns the parent process ID of the calling process.

11 4.1.1.3 Returns

12 See 4.1.1.2.

13 **4.1.1.4 Errors**

The *getpid()* and *getppid()* functions are always successful, and no return value is reserved to indicate an error.

16 4.1.1.5 Cross-References

17 *exec*, 3.1.2; *fork*(), 3.1.1; *kill*(), 3.3.2.

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

18 4.2 User Identification

19 4.2.1 Get Real User, Effective User, Real Group, and Effective Group IDs

- 20 Functions: getuid(), geteuid(), getgid(), getegid()
- 21 **4.2.1.1 Synopsis**
- 22 #include <sys/types.h>
- 23 uid_t getuid(void);
- 24 uid_t geteuid(void);
- 25 gid_t getgid(void);
- 26 gid_t getegid(void);

27 **4.2.1.2 Description**

- 28 The *getuid()* function returns the real user ID of the calling process.
- 29 The *geteuid()* function returns the effective user ID of the calling process.
- 30 The *getgid()* function returns the real group ID of the calling process.
- 31 The *getegid()* function returns the effective group ID of the calling process.

32 4.2.1.3 Returns

- 33 See 4.2.1.2.
- 34 **4.2.1.4 Errors**
- The getuid(), geteuid(), getgid(), and getegid() functions are always successful, and no return value is reserved to indicate an error.
- 37 4.2.1.5 Cross-References
- *setuid()*, 4.2.2.

39 4.2.2 Set User and Group IDs

- 40 Functions: setuid(), setgid()
- 41 **4.2.2.1 Synopsis**
- 42 #include <sys/types.h>
- 43 int setuid(uid_t uid);
- 44 int setgid(gid_t gid);

45 **4.2.2.2 Description**

- 46 If {_POSIX_SAVED_IDS} is defined:
- 47 (1) If the process has appropriate privileges, the *setuid()* function sets the
 48 real user ID, effective user ID, and the saved set-user-ID to *uid*.
- (2) If the process does not have appropriate privileges, but *uid* is equal to the real user ID or the saved set-user-ID, the *setuid()* function sets the effective user ID to *uid*; the real user ID and saved set-user-ID remain unchanged by this function call.
- (3) If the process has appropriate privileges, the setgid() function sets the
 real group ID, effective group ID, and the saved set-group-ID to gid.
 - (4) If the process does not have appropriate privileges, but gid is equal to the real group ID or the saved set-group-ID, the setgid() function sets the effective group ID to gid; the real group ID and saved set-group-ID remain unchanged by this function call.
- 59 Otherwise:

55

56

57

58

- (1) If the process has appropriate privileges, the *setuid()* function sets the
 real user ID and effective user ID to *uid*.
- (2) If the process does not have appropriate privileges, but *uid* is equal to the
 real user ID, the *setuid()* function sets the effective user ID to *uid*; the
 real user ID remains unchanged by this function call.
- (3) If the process has appropriate privileges, the setgid() function sets the
 real group ID and effective group ID to gid.
- (4) If the process does not have appropriate privileges, but gid is equal to the
 real group ID, the setgid() function sets the effective group ID to gid; the
 real group ID remains unchanged by this function call.
- Any supplementary group IDs of the calling process remain unchanged by thesefunction calls.

72 **4.2.2.3 Returns**

⁷³ Upon successful completion, a value of zero is returned. Otherwise, a value of -1⁷⁴ is returned and *errno* is set to indicate the error.

75 **4.2.2.4 Errors**

- If any of the following conditions occur, the setuid() function shall return -1 and
 set errno to the corresponding value:
- 78 [EINVAL] The value of the *uid* argument is invalid and not supported by 79 the implementation.
- 80[EPERM]The process does not have appropriate privileges and uid does81not match the real user ID or, if {_POSIX_SAVED_IDS} is82defined, the saved set-user-ID.

- If any of the following conditions occur, the *setgid()* function shall return -1 and set *errno* to the corresponding value:
- [EINVAL] The value of the *gid* argument is invalid and not supported by
 the implementation.
- 87[EPERM]The process does not have appropriate privileges and gid does88not match the real group ID or, if {_POSIX_SAVED_IDS} is89defined, the saved set-group-ID.
- 90 4.2.2.5 Cross-References
- 91 exec, 3.1.2; getuid(), 4.2.1.
- 92 4.2.3 Get Supplementary Group IDs
- 93 Function: getgroups()
- 94 **4.2.3.1** Synopsis
- 95 #include <sys/types.h>
- 96 int getgroups(int gidsetsize, gid_t grouplist[]);

97 **4.2.3.2 Description**

The getgroups() function fills in the array grouplist with the supplementary group IDs of the calling process. The gidsetsize argument specifies the number of elements in the supplied array grouplist. The actual number of supplementary group IDs stored in the array is returned. The values of array entries with indices larger than or equal to the returned value are undefined.

As a special case, if the *gidsetsize* argument is zero, *getgroups()* returns the number of supplemental group IDs associated with the calling process without modifying the array pointed to by the *grouplist* argument.

106 **4.2.3.3 Returns**

Upon successful completion, the number of supplementary group IDs is returned.
This value is zero if {NGROUPS_MAX} is zero. A return value of -1 indicates
failure, and *errno* is set to indicate the error.

110 **4.2.3.4 Errors**

III If any of the following conditions occur, the getgroups() function shall return -1and set *errno* to the corresponding value:

113 [EINVAL] The *gidsetsize* argument is not equal to zero and is less than 114 the number of supplementary group IDs.

115 4.2.3.5 Cross-References

- 116 *setgid*(), 4.2.2.
- 117 **4.2.4 Get User Name**
- 118 Functions: getlogin()
- 119 **4.2.4.1 Synopsis**
- 120 char *getlogin(void);

121 **4.2.4.2 Description**

The getlogin() function returns a pointer to a string giving a user name associated with the calling process, which is the login name associated with the calling process.

125 If getlogin() returns a non-NULL pointer, that pointer points to the name under 126 which the user logged in, even if there are several login names with the same 127 user ID.

128

129 **4.2.4.3 Returns**

The *getlogin()* function returns a pointer to a string containing the user's login name, or a **NULL** pointer if the user's login name cannot be found.

The return value from *getlogin()* may point to static data and, therefore, may be overwritten by each call.

134

135 **4.2.4.4 Errors**

This part of ISO/IEC 9945 does not specify any error conditions that are required to be detected for the *getlogin()* function. Some errors may be detected under conditions that are unspecified by this part of ISO/IEC 9945.

139 4.2.4.5 Cross-References

140 getpwnam(), 9.2.2; getpwuid(), 9.2.2.

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

1

141 **4.3 Process Groups**

- 142 4.3.1 Get Process Group ID
- 143 Function: getpgrp()
- 144 **4.3.1.1 Synopsis**
- 145 #include <sys/types.h>
- 146 pid_t getpgrp(void);
- 147 **4.3.1.2 Description**
- 148 The getpgrp() function returns the process group ID of the calling process.
- 149 **4.3.1.3 Returns**
- 150 See 4.3.1.2.

151 **4.3.1.4 Errors**

The *getpgrp()* function is always successful, and no return value is reserved to indicate an error.

154 4.3.1.5 Cross-References

155 setpgid(), 4.3.3; setsid(), 4.3.2; sigaction(), 3.3.4.

156 4.3.2 Create Session and Set Process Group ID

- 157 Function: setsid()
- 158 **4.3.2.1 Synopsis**
- 159 #include <sys/types.h>
- 160 pid_t setsid(void);

161 **4.3.2.2 Description**

If the calling process is not a process group leader, the *setsid()* function shall create a new session. The calling process shall be the session leader of this new session, shall be the process group leader of a new process group, and shall have no controlling terminal. The process group ID of the calling process shall be set equal to the process ID of the calling process. The calling process shall be the only process in the new process group and the only process in the new session.

l

l

168 **4.3.2.3 Returns**

Upon successful completion, the setsid() function returns the value of the process group ID of the calling process. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

172 **4.3.2.4 Errors**

If any of the following conditions occur, the setsid() function shall return -1 and set *errno* to the corresponding value:

175[EPERM]The calling process is already a process group leader, or the
process group ID of a process other than the calling process176matches the process ID of the calling process.

178 **4.3.2.5 Cross-References**

exec, 3.1.2; _exit(), 3.2.2; fork(), 3.1.1; getpid(), 4.1.1; kill(), 3.3.2; setpgid(), 4.3.3;
 sigaction(), 3.3.4.

181 4.3.3 Set Process Group ID for Job Control

- 182 Function: *setpgid()*
- 183 **4.3.3.1 Synopsis**
- 184 #include <sys/types.h>
- 185 int setpgid(pid_t pid, pid_t pgid);

186 **4.3.3.2 Description**

187 If {_POSIX_JOB_CONTROL} is defined:

The setpgid() function is used to either join an existing process group or create a new process group within the session of the calling process. The process group ID of a session leader shall not change. Upon successful completion, the process group ID of the process with a process ID that matches *pid* shall be set to *pgid*. As a special case, if *pid* is zero, the process ID of the calling process shall be used. Also, if *pgid* is zero, the process ID of the indicated process shall be used.

- 195 Otherwise:
- Either the implementation shall support the *setpgid()* function as describedabove or the *setpgid()* function shall fail.

198 **4.3.3.3 Returns**

Upon successful completion, the setpgid() function returns a value of zero. Otherwise, a value of -1 is returned and errno is set to indicate the error.

201 **4.3.3.4 Errors**

If any of the following conditions occur, the setpgid() function shall return -1 and set *errno* to the corresponding value:

204 205 206	[EACCES]	The value of the <i>pid</i> argument matches the process ID of a child process of the calling process, and the child process has successfully executed one of the <i>exec</i> functions.
207 208	[EINVAL]	The value of the <i>pgid</i> argument is less than zero or is not a value supported by the implementation.
209	[ENOSYS]	The <i>setpgid()</i> function is not supported by this implementation.
210	[EPERM]	The process indicated by the <i>pid</i> argument is a session leader.
211 212 213		The value of the <i>pid</i> argument is valid, but matches the process ID of a child process of the calling process, and the child process is not in the same session as the calling process.
214 215 216 217		The value of the <i>pgid</i> argument does not match the process ID of the process indicated by the <i>pid</i> argument, and there is no process with a process group ID that matches the value of the <i>pgid</i> argument in the same session as the calling process.
218 219	[ESRCH]	The value of the <i>pid</i> argument does not match the process ID of the calling process or of a child process of the calling process.

4.3.3.5 Cross-References

221 getpgrp(), 4.3.1; setsid(), 4.3.2; tcsetpgrp(), 7.2.4; exec, 3.1.2.

222 4.4 System Identification

- 4.4.1 Get System Name
- Function: *uname()*

225 **4.4.1.1 Synopsis**

- 226 #include <sys/utsname.h>
- 227 int uname(struct utsname *name);

228 **4.4.1.2 Description**

- The *uname()* function stores information identifying the current operating system in the structure pointed to by the argument *name*.
- The structure *utsname* is defined in the header <sys/utsname.h> and contains at least the members shown in Table 4-1.

1

Member Name	Description
sysname	Name of this implementation of the operating system.
nodename	Name of this node within an implementation-specified communications network.
release	Current release level of this implementation.
version	Current version level of this release.
machine	Name of the hardware type on which the system is running.

Table 4-1 - uname() Structure Members

Each of these data items is a null-terminated array of *char*. 243

The format of each member is implementation defined. The system documenta-244 tion (see 1.3.1.2) shall specify the source and format of each member and may 245 specify the range of values for each member. 246

The inclusion of the *nodename* member in this structure does not imply that it is 247 sufficient information for interfacing to communications networks. 248

4.4.1.3 Returns 249

Upon successful completion, a nonnegative value is returned. Otherwise, a value 250 of -1 is returned and *errno* is set to indicate the error. 251

4.4.1.4 Errors 252

This part of ISO/IEC 9945 does not specify any error conditions that are required 253 to be detected for the *uname()* function. Some errors may be detected under con-254 ditions that are unspecified by this part of ISO/IEC 9945. 255

4.5 Time 256

- 4.5.1 Get System Time 257
- Function: *time()* 258
- 4.5.1.1 Synopsis 259
- #include <time.h> 260
- 261 time t time(time t *tloc);

I

262 **4.5.1.2 Description**

- 263 The *time()* function returns the value of time in seconds since the Epoch.
- The argument *tloc* points to an area where the return value is also stored. If *tloc* is a NULL pointer, no value is stored.

266 **4.5.1.3 Returns**

Upon successful completion, time() returns the value of time. Otherwise, a value of $((time_t) - 1)$ is returned and *errno* is set to indicate the error.

269 **4.5.1.4 Errors**

- 270 This part of ISO/IEC 9945 does not specify any error conditions that are required
- to be detected for the *time()* function. Some errors may be detected under conditions that are unspecified by this part of ISO/IEC 9945.

4.5.2 Get Process Times

- 274 Function: *times()*
- 275 **4.5.2.1 Synopsis**
- 276 #include <sys/times.h>
- 277 clock_t times(struct tms *buffer);

278 **4.5.2.2 Description**

The *times()* function shall fill the structure pointed to by *buffer* with timeaccounting information. The type $clock_t$ and the *tms* structure are defined in <sys/times.h>; the *tms* structure shall contain at least the following members:

282 283	Member <u>Type</u>	Member <u>Name</u>	Description
284	clock_t	tms_utime	User CPU time.
285	clock_t	tms_stime	System CPU time.
286	clock_t	tms_cutime	User CPU time of terminated child processes.
287	$clock_t$	tms_cstime	System CPU time of terminated child processes.

All times are measured in terms of the number of clock ticks used.

The times of a terminated child process are included in the tms_cutime and tms_cstime elements of the parent when a wait() or waitpid() function returns the process ID of this terminated child. See 3.2.1. If a child process has not waited for its terminated children, their times shall not be included in its times.

The value *tms_utime* is the CPU time charged for the execution of user instructions.

I

- The value *tms_stime* is the CPU time charged for execution by the system on behalf of the process.
- The value *tms_cutime* is the sum of the *tms_utimes* and *tms_cutimes* of the child processes.
- The value *tms_cstime* is the sum of the *tms_stimes* and *tms_cstimes* of the child processes.

301 **4.5.2.3 Returns**

Upon successful completion, times() shall return the elapsed real time, in clock ticks, since an arbitrary point in the past (for example, system start-up time). This point does not change from one invocation of times() within the process to another. The return value may overflow the possible range of type $clock_t$. If the times() function fails, a value of $((clock_t) - 1)$ is returned and errno is set to indicate the error.

308 4.5.2.4 Errors

This part of ISO/IEC 9945 does not specify any error conditions that are required to be detected for the *times()* function. Some errors may be detected under conditions that are unspecified by this part of ISO/IEC 9945.

312 4.5.2.5 Cross-References

sis exec, 3.1.2; fork(), 3.1.1; sysconf(), 4.8.1; time(), 4.5.1; wait(), 3.2.1.

314 4.6 Environment Variables

- 315 4.6.1 Environment Access
- 316 Function: getenv()

317 **4.6.1.1 Synopsis**

- 318 #include <stdlib.h>
- 319 char *getenv(const char *name);

320 **4.6.1.2 Description**

The getenv() function searches the environment list (see 2.6) for a string of the form name=value and returns a pointer to value if such a string is present. If the specified name cannot be found, a NULL pointer is returned.

1

ł

1

ł

324 4.6.1.3 Returns

Upon successful completion, the getenv() function returns a pointer to a string containing the value for the specified name, or a NULL pointer if the specified name cannot be found. The return value from getenv() may point to static data and, therefore, may be overwritten by each call. Unsuccessful completion shall result in the return of a NULL pointer.

330 4.6.1.4 Errors

This part of ISO/IEC 9945 does not specify any error conditions that are required to be detected for the *getenv()* function. Some errors may be detected under conditions that are unspecified by this part of ISO/IEC 9945.

334 4.6.1.5 Cross-References

335 3.1.2; 2.6.

336 4.7 Terminal Identification

337 4.7.1 Generate Terminal Pathname

- 338 Function: *ctermid()*
- 339 **4.7.1.1 Synopsis**
- 340 #include <stdio.h>
- 341 char *ctermid(char *s);

342 **4.7.1.2 Description**

- The *ctermid()* function generates a string that, when used as a pathname, refers to the current controlling terminal for the current process.
- 345 If the *ctermid()* function returns a pathname, access to the file is not guaranteed.

346 **4.7.1.3 Returns**

If s is a NULL pointer, the string is generated in an area that may be static (and, therefore, may be overwritten by each call), the address of which is returned. Otherwise, s is assumed to point to an array of *char* of at least L_ctermid bytes; the string is placed in this array and the value of s is returned. The symbolic constant L_ctermid is defined in <stdio.h> and shall have a value greater than zero.

The *ctermid()* function shall return an empty string if the pathname that would refer to the controlling terminal cannot be determined or if the function is unsuccessful.

1

I

I

1

I

I

356 **4.7.1.4 Errors**

This part of ISO/IEC 9945 does not specify any error conditions that are required to be detected for the *ctermid()* function. Some errors may be detected under conditions that are unspecified by this part of ISO/IEC 9945.

360 4.7.1.5 Cross-References

361 *ttyname()*, 4.7.2.

362 4.7.2 Determine Terminal Device Name

363 Functions: *ttyname()*, *isatty()*

364 **4.7.2.1** Synopsis

- 365 char *ttyname(int fildes);
- 366 int isatty(int fildes);

367 **4.7.2.2 Description**

- The *ttyname()* function returns a pointer to a string containing a null-terminated pathname of the terminal associated with file descriptor *fildes*.
- The return value of *ttyname()* may point to static data that is overwritten by each call.
- The *isatty()* function returns 1 if *fildes* is a valid file descriptor associated with a terminal, zero otherwise.

374 **4.7.2.3 Returns**

The *ttyname()* function returns a NULL pointer if *fildes* is not a valid file descriptor associated with a terminal or if the pathname cannot be determined.

377 4.7.2.4 Errors

This part of ISO/IEC 9945 does not specify any error conditions that are required to be detected for the *ttyname()* or *isatty()* functions. Some errors may be detected under conditions that are unspecified by this part of ISO/IEC 9945.

I

381 **4.8 Configurable System Variables**

- 382 **4.8.1 Get Configurable System Variables**
- 383 Function: sysconf()
- 384 **4.8.1.1 Synopsis**
- 385 #include <unistd.h>
- 386 long sysconf(int name);

387 **4.8.1.2 Description**

205

The sysconf() function provides a method for the application to determine the current value of a configurable system limit or option (variable).

The name argument represents the system variable to be queried. The implementation shall support all of the variables listed in Table 4-2 and may support others. The variables in Table 4-2 come from <limits.h> or <unistd.h> and the symbolic constants, defined in <unistd.h>, that are the corresponding values used for name.

396		ole System variables	
397	Variable	name Value	
398	(ARG_MAX)	{_SC_ARG_MAX}	
399	{CHILD_MAX}	[_SC_CHILD_MAX]	
400	clock ticks/second	{_SC_CLK_TCK}	
401	{NGROUPS_MAX}	[_SC_NGROUPS_MAX]	
402	{OPEN_MAX}	[_SC_OPEN_MAX]	
403	{STREAM_MAX}	[_SC_STREAM_MAX]	
404	{TZNAME_MAX}	{_SC_TZNAME_MAX}	
405	{_POSIX_JOB_CONTROL}	[_SC_JOB_CONTROL]	
406	{_POSIX_SAVED_IDS}	{_SC_SAVED_IDS}	
407	{_POSIX_VERSION}	{_SC_VERSION}	
408			

Table 4-2 - Configurable System Variables

409 **4.8.1.3 Returns**

410 If name is an invalid value, sysconf() shall return -1. If the variable correspond-411 ing to name is associated with functionality that is not supported by the system, 412 sysconf() shall return -1 without changing the value of errno.

Otherwise, the *sysconf()* function returns the current variable value on the system. The value returned shall not be more restrictive than the corresponding value described to the application when it was compiled with the implementation's <limits.h> or <unistd.h>. The value shall not change during the lifetime of the calling process.

l

I

418 **4.8.1.4 Errors**

If any of the following conditions occur, the sysconf() function shall return -1 and set *errno* to the corresponding value:

421 [EINVAL] The value of the *name* argument is invalid.

422 4.8.1.5 Special Symbol {CLK_TCK}

The special symbol {CLK_TCK} shall yield the same result 423 as sysconf (SC CLK TCK). It shall be defined in <time.h>. The symbol 424 {CLK_TCK} may be evaluated by the implementation at run time or may be a con-425 stant. This special symbol is obsolescent. 426



Section 5: Files and Directories

The functions in this section perform the operating system services dealing with the creation and removal of files and directories and the detection and modification of their characteristics. They also provide the primary methods a process will use to gain access to files and directories for subsequent I/O operations (see Section 6).

6 5.1 Directories

7 5.1.1 Format of Directory Entries

- 8 The header <dirent.h> defines a structure and a defined type used by the *direc-*9 *tory* routines.
- 10 The internal format of directories is unspecified.
- 11 The *readdir()* function returns a pointer to an object of type *struct dirent* that 12 includes the member:

13	Member	Member	Description
14	Type	Name	
15	char []	d_name	Null-terminated filename

The array of *char d_name* is of unspecified size, but the number of bytes preceding | the terminating null character shall not exceed {NAME_MAX}.

18 5.1.2 Directory Operations

19 Functions: opendir(), readdir(), rewinddir(), closedir()

20 **5.1.2.1 Synopsis**

- 21 #include <sys/types.h>
- 22 #include <dirent.h>
- 23 DIR *opendir(const char *dirname);
- 24 struct dirent *readdir(DIR *dirp);
- 25 void rewinddir(DIR *dirp);
- 26 int closedir(DIR *dirp);

27 **5.1.2.2 Description**

The type DIR, which is defined in the header <dirent.h>, represents a directory 28 stream, which is an ordered sequence of all the directory entries in a particular 29 directory. Directory entries represent files; files may be removed from a directory 30 or added to a directory asynchronously to the operations described in this sub-31 clause (5.1.2). The type *DIR* may be implemented using a file descriptor. In that 32 case, applications will only be able to open up to a total of {OPEN_MAX} files and 33 directories; see 5.3.1. A successful call to any of the *exec* functions shall close any 34 directory streams that are open in the calling process. 35

- The *opendir()* function opens a directory stream corresponding to the directory named by the *dirname* argument. The directory stream is positioned at the first entry.
- The *readdir()* function returns a pointer to a structure representing the directory entry at the current position in the directory stream to which *dirp* refers, and positions the directory stream at the next entry. It returns a **NULL** pointer upon reaching the end of the directory stream.
- The *readdir()* function shall not return directory entries containing empty names.
 It is unspecified whether entries are returned for dot or dot-dot.
- The pointer returned by *readdir()* points to data that may be overwritten by another call to *readdir()* on the same directory stream. This data shall not be overwritten by another call to *readdir()* on a different directory stream.
- The *readdir()* function may buffer several directory entries per actual read operation; the *readdir()* function shall mark for update the *st_atime* field of the directory each time the directory is actually read.
- The *rewinddir()* function resets the position of the directory stream to which *dirp* refers to the beginning of the directory. It also causes the directory stream to refer to the current state of the corresponding directory, as a call to *opendir()* would have done. It does not return a value.
- If a file is removed from or added to the directory after the most recent call to *opendir()* or *rewinddir()*, whether a subsequent call to *readdir()* returns an entry for that file is unspecified.
- The *closedir()* function closes the directory stream referred to by *dirp* and returns a value of zero if successful. Otherwise, it returns -1 indicating an error. Upon return, the value of *dirp* may no longer point to an accessible object of type *DIR*. If a file descriptor is used to implement type *DIR*, that file descriptor shall be closed.
- If the *dirp* argument passed to any of these functions does not refer to a currently open directory stream, the effect is undefined.
- The result of using a directory stream after one of the *exec* family of functions is undefined. After a call to the fork() function, either the parent or the child (but not both) may continue processing the directory stream using readdir() or rewinddir() or both. If both the parent and child processes use these functions, the result is undefined. Either or both processes may use closedir().

L

I

I

70 **5.1.2.3 Returns**

Upon successful completion, opendir() returns a pointer to an object of type DIR.
Otherwise, a value of NULL is returned and errno is set to indicate the error.

Upon successful completion, readdir() returns a pointer to an object of type struct
dirent. When an error is encountered, a value of NULL is returned and errno is
set to indicate the error. When the end of the directory is encountered, a value of
NULL is returned and errno is unchanged by this function call.

Upon successful completion, *closedir()* returns a value of zero. Otherwise, a value
 of -1 is returned and *errno* is set to indicate the error.

79 **5.1.2.4 Errors**

If any of the following conditions occur, the *opendir()* function shall return a value
 of NULL and set *errno* to the corresponding value:

82 [EACCES] Search permission is denied for a component of the path prefix 83 of *dirname*, or read permission is denied for the directory itself.

84 [ENAMETOOLONG]

- 85The length of the dirname argument exceeds {PATH_MAX}, or a86pathname component is longer than {NAME_MAX} while87{_POSIX_NO_TRUNC} is in effect.
- [ENOENT] The named directory does not exist, or *dirname* points to an
 empty string.

90 [ENOTDIR] A component of *dirname* is not a directory.

For each of the following conditions, when the condition is detected, the *opendir()* function shall return a value of NULL and set *errno* to the corresponding value:

- 93 [EMFILE] Too many file descriptors are currently open for the process.
- 94 [ENFILE] Too many file descriptors are currently open in the system.

For each of the following conditions, when the condition is detected, the *readdir()* function shall return a value of NULL and set *errno* to the corresponding value:

- 97 [EBADF] The *dirp* argument does not refer to an open directory stream.
- For each of the following conditions, when the condition is detected, the *closedir()* function shall return -1 and set *errno* to the corresponding value:
- 100 [EBADF] The *dirp* argument does not refer to an open directory stream.
- 101 5.1.2.5 Cross-References
- 102 <dirent.h>, 5.1.1.

103 **5.2 Working Directory**

- 104 5.2.1 Change Current Working Directory
- 105 Function: chdir()
- 106 **5.2.1.1 Synopsis**
- 107 int chdir(const char *path);

108 **5.2.1.2 Description**

The *path* argument points to the pathname of a directory. The *chdir()* function causes the named directory to become the current working directory, that is, the starting point for path searches of pathnames not beginning with slash.

112 If the *chdir()* function fails, the current working directory shall remain 113 unchanged by this function call.

114 **5.2.1.3 Returns**

¹¹⁵ Upon successful completion, a value of zero is returned. Otherwise, a value of -1¹¹⁶ is returned and *errno* is set to indicate the error.

117 **5.2.1.4 Errors**

118 If any of the following conditions occur, the chdir() function shall return -1 and 119 set *errno* to the corresponding value:

- 120 [EACCES] Search permission is denied for any component of the path-121 name.
- 122 [ENAMETOOLONG]
- 123The path argument exceeds {PATH_MAX} in length, or a path-124name component is longer than {NAME_MAX} while125{_POSIX_NO_TRUNC} is in effect.
- 126 [ENOTDIR] A component of the pathname is not a directory.
- 127 [ENOENT] The named directory does not exist or *path* is an empty string.
- 128 **5.2.1.5 Cross-References**
- 129 *getcwd()*, 5.2.2.
l

I

l

I

130 5.2.2 Get Working Directory Pathname

- 131 Function: getcwd()
- 132 **5.2.2.1 Synopsis**
- 133 char *getcwd(char *buf, size_t size);

134 **5.2.2.2 Description**

The getcwd() function copies an absolute pathname of the current working directory to the array of char pointed to by the argument buf and returns a pointer to the result. The size argument is the size in bytes of the array of char pointed to by the buf argument. If buf is a NULL pointer, the behavior of getcwd() is undefined.

140 **5.2.2.3 Returns**

141 If successful, the *buf* argument is returned. A NULL pointer is returned if an 142 error occurs and the variable *errno* is set to indicate the error. The contents of *buf* 143 after an error are undefined.

144 **5.2.2.4 Errors**

If any of the following conditions occur, the getcwd() function shall return a valueof NULL and set errno to the corresponding value:

147 [EINVAL] The *size* argument is zero.

148 [ERANGE] The *size* argument is greater than zero but smaller than the 149 length of the pathname plus 1.

For each of the following conditions, if the condition is detected, the *getcwd()* function shall return a value of NULL and set *errno* to the corresponding value:

- 152 [EACCES] Read or search permission was denied for a component of the 153 pathname.
- 154 5.2.2.5 Cross-References
- 155 *chdir()*, 5.2.1.

1

1

156 **5.3 General File Creation**

- 157 **5.3.1 Open a File**
- 158 Function: open()
- 159 **5.3.1.1 Synopsis**
- 160 #include <sys/types.h>
 161 #include <sys/stat.h>
- 162 #include <fcntl.h>
- 163 int open(const char *path, int oflag, ...);

164 **5.3.1.2 Description**

The *open()* function establishes the connection between a file and a file descriptor. It creates an open file description that refers to a file and a file descriptor that refers to that open file description. The file descriptor is used by other I/O functions to refer to that file. The *path* argument points to a pathname naming a file.

The open() function shall return a file descriptor for the named file that is the 169 lowest file descriptor not currently open for that process. The open file description 170 is new, and therefore the file descriptor does not share it with any other process 171 in the system. The file offset shall be set to the beginning of the file. The 172 FD_CLOEXEC file descriptor flag associated with the new file descriptor shall be 173 cleared. The file status flags and file access modes of the open file description 174 shall be set according to the value of oflag. The value of oflag is the bitwise 175 inclusive OR of values from the following list. See 6.5.1 for the definitions of the 176 symbolic constants. Applications shall specify exactly one of the first three values 177 (file access modes) below in the value of *oflag*: 178

- 179 O_RDONLY Open for reading only.
- 180 O_WRONLY Open for writing only.
- 181 O_RDWR Open for reading and writing. The result is undefined if this
 182 flag is applied to a FIFO.
- 183 Any combination of the remaining flags may be specified in the value of oflag:
- 184 O_APPEND If set, the file offset shall be set to the end of the file prior to 185 each write.
- O CREAT This option requires a third argument, *mode*, which is of type 186 *mode_t*. If the file exists, this flag has no effect, except as noted 187 under O_EXCL, below. Otherwise, the file is created; the file's 188 user ID shall be set to the effective user ID of the process; the 189 file's group ID shall be set to the group ID of the directory in 190 which the file is being created or to the effective group ID of the 191 process. The file permission bits (see 5.6.1) shall be set to the 192 value of *mode* except those set in the file mode creation mask of 193 the process (see 5.3.3). When bits in mode other than the file 194 permission bits are set, the effect is unspecified. The mode 195

I

196 197		argument does not affect whether the file is opened for reading, for writing, or for both.
198 199 200 201 202 203	O_EXCL	If O_EXCL and O_CREAT are set, <i>open()</i> shall fail if the file exists. The check for the existence of the file and the creation of the file if it does not exist shall be atomic with respect to other processes executing <i>open()</i> naming the same filename in the same directory with O_EXCL and O_CREAT set. If O_EXCL is set and O_CREAT is not set, the result is undefined.
204 205 206	O_NOCTTY	If set, and <i>path</i> identifies a terminal device, the $open()$ function shall not cause the terminal device to become the controlling terminal for the process (see 7.1.1.3).
207	O_NONBLOC	K
208		(1) When opening a FIFO with O_RDONLY or O_WRONLY set:
209 210 211 212 213		 (a) If O_NONBLOCK is set: An open() for reading-only shall return without delay. An open() for writing-only shall return an error if no process currently has the file open for reading.
214 215 216 217		 (b) If O_NONBLOCK is clear: An open() for reading-only shall block until a process opens the file for writing. An open() for writing-only shall block until a process opens the file for reading.
218 219		(2) When opening a block special or character special file that supports nonblocking opens:
220 221 222 223		 (a) If O_NONBLOCK is set: The open() shall return without waiting for the dev- ice to be ready or available. Subsequent behavior of the device is device-specific.
224 225 226		(b) If O_NONBLOCK is clear: The open() shall wait until the device is ready or available before returning.
227		(3) Otherwise, the behavior of O_NONBLOCK is unspecified.
228 229 230 231 232 233 234	O_TRUNC	If the file exists and is a regular file, and the file is successfully opened O_RDWR or O_WRONLY, it shall be truncated to zero length and the mode and owner shall be unchanged by this function call. O_TRUNC shall have no effect on FIFO special files or terminal device files. Its effect on other file types is implementation defined. The result of using O_TRUNC with O_RDONLY is undefined.
235	If O_CREAT is se	et and the file did not previously exist, upon successful completion

If O_CREAT is set and the file did not previously exist, upon successful completion
the open() function shall mark for update the st_atime, st_ctime, and st_mtime
fields of the file and the st_ctime and st_mtime fields of the parent directory.

If O_TRUNC is set and the file did previously exist, upon successful completion the open() function shall mark for update the st_ctime and st_mtime fields of the file.

241 **5.3.1.3 Returns**

Upon successful completion, the function shall open the file and return a nonnegative integer representing the lowest numbered unused file descriptor. Otherwise, it shall return -1 and shall set *errno* to indicate the error. No files shall be created or modified if the function returns -1.

246 **5.3.1.4 Errors**

If any of the following conditions occur, the *open()* function shall return -1 and set *errno* to the corresponding value:

249 250	[EACCES]	Search permission is denied on a component of the path prefix, or the file exists and the permissions specified by <i>oflag</i> are
251		denied, or the file does not exist and write permission is denied
252		for the parent directory of the file to be created, or O_TRUNC is
253		specified and write permission is denied.
254	[EEXIST]	O_CREAT and O_EXCL are set and the named file exists.
255	[EINTR]	The open() operation was interrupted by a signal.
256 257	[EISDIR]	The named file is a directory, and the <i>oflag</i> argument specifies write or read/write access.
258	[EMFILE]	Too many file descriptors are currently in use by this process.
259	[ENAMETOO	LONG]
260		The length of the <i>path</i> string exceeds {PATH_MAX}, or a path-
261		name component is longer than {NAME_MAX} while
262		{_POSIX_NO_TRUNC} is in effect.
263	[ENFILE]	Too many files are currently open in the system.
264	[ENOENT]	O_CREAT is not set and the named file does not exist, or
265		O_CREAT is set and either the path prefix does not exist or the
266		<i>path</i> argument points to an empty string.
267	[ENOSPC]	The directory or file system that would contain the new file can-
268		not be extended.
269	[ENOTDIR]	A component of the path prefix is not a directory.
270	[ENXIO]	O_NONBLOCK is set, the named file is a FIFO, O_WRONLY is
271		set, and no process has the file open for reading.
272	[EROFS]	The named file resides on a read-only file system and either
273		O_WRONLY, O_RDWR, O_CREAT (if the file does not exist), or
274		O_TRUNC is set in the <i>oflag</i> argument.

Part 1: SYSTEM API [C LANGUAGE]

275 5.3.1.5 Cross-References

- 276 close(), 6.3.1; creat(), 5.3.2; dup(), 6.2.1; exec, 3.1.2; fcntl(), 6.5.2; <fcntl.h>,
- 277 6.5.1; lseek(), 6.5.3; read(), 6.4.1; <signal.h>, 3.3.1; stat(), 5.6.2;
- 278 <sys/stat.h>, 5.6.1; write(), 6.4.2; umask(), 5.3.3; 3.3.1.4.

279 5.3.2 Create a New File or Rewrite an Existing One

- 280 Function: creat()
- 281 **5.3.2.1 Synopsis**
- 282 #include <sys/types.h>
- 283 #include <sys/stat.h>
- 284 #include <fcntl.h>
- 285 int creat(const char *path, mode_t mode);

286 **5.3.2.2 Description**

- 287 The function call:
- 288 creat(path, mode);
- 289 is equivalent to:
- 290 open (path, O_WRONLY | O_CREAT | O_TRUNC, mode);

291 5.3.2.3 Cross-References

292 open(), 5.3.1; <sys/stat.h>, 5.6.1.

293 5.3.3 Set File Creation Mask

294 Function: *umask()*

295 **5.3.3.1 Synopsis**

- 296 #include <sys/types.h>
- 297 #include <sys/stat.h>
- 298 mode_t umask(mode_t cmask);

299 **5.3.3.2 Description**

The umask() routine sets the file mode creation mask of the process to cmask and returns the previous value of the mask. Only the file permission bits (see 5.6.1) of *cmask* are used; the meaning of the other bits is implementation defined.

The file mode creation mask of the process is used during open(), creat(), mkdir(), and mkfifo() calls to turn off permission bits in the mode argument supplied. Bit positions that are set in cmask are cleared in the mode of the created file.

306 **5.3.3.3 Returns**

The file permission bits in the value returned by umask() shall be the previous value of the file mode creation mask. The state of any other bits in that value is unspecified, except that a subsequent call to umask() with that returned value as cmask shall leave the state of the mask the same as its state before the first call, including any unspecified (by this part of ISO/IEC 9945) use of those bits.

312 **5.3.3.4 Errors**

The umask() function is always successful, and no return value is reserved to indicate an error.

- 315 5.3.3.5 Cross-References
- s16 chmod(), 5.6.4; creat(), 5.3.2; mkdir(), 5.4.1; mkfifo(), 5.4.2; open(), 5.3.1;
- 317 <sys/stat.h>, 5.6.1.
- 318 5.3.4 Link to a File
- 319 Function: link()
- 320 **5.3.4.1 Synopsis**
- 321 int link(const char *existing, const char *new);

322 **5.3.4.2 Description**

The argument *existing* points to a pathname naming an existing file. The argument *new* points to a pathname naming the new directory entry to be created. Implementations may support linking of files across file systems. The *link()* function shall atomically create a new link for the existing file and increment the link count of the file by one.

- If the *link()* function fails, no link shall be created, and the link count of the file shall remain unchanged by this function call.
- The *existing* argument shall not name a directory unless the user has appropriate | privileges and the implementation supports using *link()* on directories.
- The implementation may require that the calling process has permission to access the existing file.
- ³³⁴ Upon successful completion, the *link()* function shall mark for update the *st_ctime*
- field of the file. Also, the *st_ctime* and *st_mtime* fields of the directory that contains the new entry are marked for update.

337 **5.3.4.3 Returns**

338 Upon successful completion, link() shall return a value of zero. Otherwise, a 339 value of -1 is returned and *errno* is set to indicate the error.

340 **5.3.4.4 Errors**

If any of the following conditions occur, the link() function shall return -1 and set
 errno to the corresponding value:

- 343[EACCES]A component of either path prefix denies search permission; or344the requested link requires writing in a directory with a mode345that denies write permission; or the calling process does not346have permission to access the existing file, and this is required347by the implementation.
- 348 [EEXIST] The link named by new exists.
- 349[EMLINK]The number of links to the file named by existing would exceed350{LINK_MAX}.

351 [ENAMETOOLONG]

- 352The length of the existing or new string exceeds {PATH_MAX},353or a pathname component is longer than {NAME_MAX} while354{_POSIX_NO_TRUNC} is in effect.
- 355[ENOENT]A component of either path prefix does not exist, the file named356by existing does not exist, or either existing or new points to an |357empty string.
- 358 [ENOSPC] The directory that would contain the link cannot be extended.
- 359 [ENOTDIR] A component of either path prefix is not a directory.
- 360[EPERM]The file named by existing is a directory, and either the calling361process does not have appropriate privileges or the implemen-362tation prohibits using link() on directories.
- 363 [EROFS] The requested link requires writing in a directory on a read-364 only file system.
- 365[EXDEV]The link named by new and the file named by existing are on366different file systems, and the implementation does not support367links between file systems.

368 5.3.4.5 Cross-References

rename(), 5.5.3; *unlink()*, 5.5.1.

I

370 **5.4 Special File Creation**

- 371 **5.4.1 Make a Directory**
- 372 Function: *mkdir()*
- 373 **5.4.1.1 Synopsis**
- 374 #include <sys/types.h>
- 375 #include <sys/stat.h>
- 376 int mkdir(const char *path, mode_t mode);

377 **5.4.1.2 Description**

The *mkdir()* routine creates a new directory with name *path*. The file permission bits of the new directory are initialized from *mode*. The file permission bits of the *mode* argument are modified by the file creation mask of the process (see 5.3.3). When bits in *mode* other than the file permission bits are set, the meaning of these additional bits is implementation defined.

The owner ID of the directory is set to the effective user ID of the process. The directory's group ID shall be set to the group ID of the directory in which the directory is being created or to the effective group ID of the process.

386 The newly created directory shall be an empty directory.

Upon successful completion, the mkdir() function shall mark for update the st_atime, st_ctime, and st_mtime fields of the directory. Also, the st_ctime and st_mtime fields of the directory that contains the new entry are marked for update.

391 **5.4.1.3 Returns**

A return value of zero indicates success. A return value of -1 indicates that an error has occurred, and an error code is stored in *errno*. No directory shall be created if the return value is -1.

395 **5.4.1.4 Errors**

If any of the following conditions occur, the mkdir() function shall return -1 and set *errno* to the corresponding value:

- 398 [EACCES] Search permission is denied on a component of the path prefix,
 399 or write permission is denied on the parent directory of the
 400 directory to be created.
- 401 [EEXIST] The named file exists.
- 402 [EMLINK] The link count of the parent directory would exceed 403 {LINK_MAX}.

4 [ENA	ETOOLONG]	
)5	The length of the path argument e	exceeds {PATH_MAX}, or a
)6	pathname component is longer	than {NAME_MAX} while
)7	{_POSIX_NO_TRUNC} is in effect.	
)8 [ENC	NT] A component of the path prefix does	not exist, or the <i>path</i> argu-
)9	ment points to an empty string.	
lo [ENC	PC] The file system does not contain end	ough space to hold the con-
11	tents of the new directory or to exte	end the parent directory of
12	the new directory.	
I3 [ENC	DIR] A component of the path prefix is not	t a directory.
14 [ERO	3] The parent directory of the directory read-only file system.	v being created resides on a
10 [ENC 11 12 13 [ENC 14 [ERC	 PC] The file system does not contain end tents of the new directory or to extend the new directory. DIR] A component of the path prefix is not B) The parent directory of the directory read-only file system. 	ough space t end the pare t a directory. y being creat

- 416 **5.4.1.5 Cross-References**
- 417 *chmod()*, 5.6.4; *stat()*, 5.6.2; <*sys/stat.h>*, 5.6.1; *umask()*, 5.3.3.

418 5.4.2 Make a FIFO Special File

- 419 Function: *mkfifo()*
- 420 **5.4.2.1** Synopsis
- 421 #include <sys/types.h>
- 422 #include <sys/stat.h>
- 423 int mkfifo(const char *path, mode_t mode);

424 **5.4.2.2 Description**

The *mkfifo()* routine creates a new FIFO special file named by the pathname pointed to by *path*. The file permission bits of the new FIFO are initialized from *mode*. The file permission bits of the *mode* argument are modified by the file creation mask of the process (see 5.3.3). When bits in *mode* other than the file permission bits are set, the effect is implementation defined.

The owner ID of the FIFO shall be set to the effective user ID of the process. The group ID of the FIFO shall be set to the group ID of the directory in which the FIFO is being created or to the effective group ID of the process.

Upon successful completion, the *mkfifo()* function shall mark for update the *st_atime*, *st_ctime*, and *st_mtime* fields of the file. Also, the *st_ctime* and *st_mtime*fields of the directory that contains the new entry are marked for update.

436 **5.4.2.3 Returns**

Upon successful completion, a value of zero is returned. Otherwise, a value of -1
is returned, no FIFO is created, and *errno* is set to indicate the error.

439 **5.4.2.4 Errors**

If any of the following conditions occur, the mkfifo() function shall return -1 and set *errno* to the corresponding value:

- 442 [EACCES] Search permission is denied on a component of the path prefix,
 443 or write permission is denied on the parent directory of the file
 444 to be created.
- 445 [EEXIST] The named file already exists.

446 [ENAMETOOLONG]

- 447The length of the path string exceeds {PATH_MAX}, or a path-448name component is longer than {NAME_MAX} while449{_POSIX_NO_TRUNC} is in effect.
- 450 [ENOENT] A component of the path prefix does not exist, or the *path* argu-451 ment points to an empty string.
- 452 [ENOSPC] The directory that would contain the new file cannot be 453 extended, or the file system is out of file allocation resources.
- 454 [ENOTDIR] A component of the path prefix is not a directory.
- 455 [EROFS] The named file resides on a read-only file system.
- 456 5.4.2.5 Cross-References
- 457 chmod(), 5.6.4; exec, 3.1.2; pipe(), 6.1.1; stat(), 5.6.2; <sys/stat.h>, 5.6.1; 458 umask(), 5.3.3.

459 **5.5 File Removal**

460 **5.5.1 Remove Directory Entries**

- 461 Function: *unlink()*
- 462 **5.5.1.1 Synopsis**
- 463 int unlink(const char *path);

464 **5.5.1.2 Description**

The *unlink()* function shall remove the link named by the pathname pointed to by *path* and decrement the link count of the file referenced by the link.

467 When the link count of the file becomes zero and no process has the file open, the 468 space occupied by the file shall be freed and the file shall no longer be accessible. 469 If one or more processes have the file open when the last link is removed, the link 470 shall be removed before unlink() returns, but the removal of the file contents shall 471 be postponed until all references to the file have been closed

471 be postponed until all references to the file have been closed.

- The *path* argument shall not name a directory unless the process has appropriate privileges and the implementation supports using *unlink()* on directories. Appli-
- 474 cations should use *rmdir()* to remove a directory.

Upon successful completion, the unlink() function shall mark for update the st_ctime and st_mtime fields of the parent directory. Also, if the link count of the file is not zero, the st_ctime field of the file shall be marked for update.

478 **5.5.1.3 Returns**

Upon successful completion, a value of zero shall be returned. Otherwise, a value of -1 shall be returned and *errno* shall be set to indicate the error. If -1 is returned, the named file shall not be changed by this function call.

482 **5.5.1.4 Errors**

If any of the following conditions occur, the *unlink()* function shall return -1 and set *errno* to the corresponding value:

- 485 [EACCES] Search permission is denied for a component of the path prefix, 486 or write permission is denied on the directory containing the 487 link to be removed.
- 488[EBUSY]The directory named by the path argument cannot be unlinked489because it is being used by the system or another process and490the implementation considers this to be an error.

491 [ENAMETOOLONG]

- 492The length of the path argument exceeds {PATH_MAX}, or a493pathname component is longer than {NAME_MAX} while494{_POSIX_NO_TRUNC} is in effect.
- 495 [ENOENT] The named file does not exist, or the *path* argument points to 496 an empty string.
- 497 [ENOTDIR] A component of the path prefix is not a directory.
- 498[EPERM]The file named by path is a directory, and either the calling499process does not have appropriate privileges or the implemen-500tation prohibits using unlink() on directories.
- 501 [EROFS] The directory entry to be unlinked resides on a read-only file 502 system.

503 5.5.1.5 Cross-References

⁵⁰⁴ *close()*, 6.3.1; *link()*, 5.3.4; *open()*, 5.3.1; *rename()*, 5.5.3; *rmdir()*, 5.5.2.

l

- 505 5.5.2 Remove a Directory
- 506 Function: *rmdir()*
- 507 **5.5.2.1 Synopsis**
- 508 int rmdir(const char *path);

509 **5.5.2.2 Description**

The *rmdir()* function removes a directory whose name is given by *path*. The directory shall be removed only if it is an empty directory.

512 If the named directory is the root directory or the current working directory of any 513 process, it is unspecified whether the function succeeds or whether it fails and 514 sets *errno* to [EBUSY].

If the link count of the directory becomes zero and no process has the directory open, the space occupied by the directory shall be freed and the directory shall no longer be accessible. If one or more processes have the directory open when the last link is removed, the dot and dot-dot entries, if present, are removed before rmdir() returns and no new entries may be created in the directory, but the directory is not removed until all references to the directory have been closed.

521 Upon successful completion, the rmdir() function shall mark for update the 522 st_ctime and st_mtime fields of the parent directory.

523 **5.5.2.3 Returns**

⁵²⁴ Upon successful completion, a value of zero shall be returned. Otherwise, a value ⁵²⁵ of -1 shall be returned and *errno* shall be set to indicate the error. If -1 is ⁵²⁶ returned, the named directory shall not be changed by this function call.

527 **5.5.2.4 Errors**

If any of the following conditions occur, the rmdir() function shall return -1 and set *errno* to the corresponding value:

530	[EACCES]	Search permission is denied on a component of the path prefix,
531		or write permission is denied on the parent directory of the
532		directory to be removed.
533	[EBUSY]	The directory named by the <i>path</i> argument cannot be removed
534		because it is being used by another process and the implemen-
535		tation considers this to be an error.
536	[EEXIST] or	[ENOTEMPTY]
537		The <i>path</i> argument names a directory that is not an empty
538		directory.
539	[ENAMETOC	DLONG]
540		The length of the <i>path</i> argument exceeds {PATH_MAX}, or a
541		pathname component is longer than {NAME_MAX} while

542		{_POSIX_NO_TRUNC} is in effect.
543 544	[ENOENT]	The <i>path</i> argument names a nonexistent directory or points to an empty string.
545	[ENOTDIR]	A component of the path is not a directory.
546 547	[EROFS]	The directory entry to be removed resides on a read-only file system.

- 548 5.5.2.5 Cross-References
- 549 *mkdir()*, 5.4.1; *unlink()*, 5.5.1.
- 550 **5.5.3 Rename a File**
- 551 Function: rename()
- 552 **5.5.3.1 Synopsis**
- 553 int rename(const char *old, const char *new);

554 **5.5.3.2 Description**

The *rename()* function changes the name of a file. The *old* argument points to the pathname of the file to be renamed. The *new* argument points to the new pathname of the file.

If the *old* argument and the *new* argument both refer to links to the same existing file, the *rename()* function shall return successfully and perform no other action.

If the *old* argument points to the pathname of a file that is not a directory, the *new* argument shall not point to the pathname of a directory. If the link named by the *new* argument exists, it shall be removed and *old* renamed to *new*. In this case, a link named *new* shall exist throughout the renaming operation and shall refer either to the file referred to by *new* or *old* before the operation began. Write access permission is required for both the directory containing *old* and the directory containing *new*.

If the *old* argument points to the pathname of a directory, the *new* argument shall not point to the pathname of a file that is not a directory. If the directory named by the *new* argument exists, it shall be removed and *old* renamed to *new*. In this case, a link named *new* shall exist throughout the renaming operation and shall refer either to the file referred to by *new* or *old* before the operation began. Thus, if *new* names an existing directory, it shall be required to be an empty directory.

The *new* pathname shall not contain a path prefix that names *old*. Write access permission is required for the directory containing *old* and the directory containing *new*. If the *old* argument points to the pathname of a directory, write access permission may be required for the directory named by *old*, and, if it exists, the directory named by *new*.

If the link named by the *new* argument exists and the link count of the file becomes zero when it is removed and no process has the file open, the space occupied by the file shall be freed and the file shall no longer be accessible. If one or more processes have the file open when the last link is removed, the link shall be removed before *rename()* returns, but the removal of the file contents shall be postponed until all references to the file have been closed.

⁵⁸⁴ Upon successful completion, the *rename()* function shall mark for update the ⁵⁸⁵ st_ctime and st_mtime fields of the parent directory of each file.

586 **5.5.3.3 Returns**

⁵⁸⁷ Upon successful completion, a value of zero shall be returned. Otherwise, a value ⁵⁸⁸ of -1 shall be returned and *errno* shall be set to indicate the error. If -1 is ⁵⁸⁹ returned, neither the file named by *old* nor the file named by *new*, if either exists, ⁵⁹⁰ shall be changed by this function call.

591 **5.5.3.4 Errors**

⁵⁹² If any of the following conditions occur, the *rename()* function shall return -1 and ⁵⁹³ set *errno* to the corresponding value:

594 595 59 6 597	[EACCES]	A component of either path prefix denies search permission, or one of the directories containing <i>old</i> or <i>new</i> denies write per- missions, or write permission is required and is denied for a directory pointed to by the <i>old</i> or <i>new</i> arguments.
598 599 600	[EBUSY]	The directory named by <i>old</i> or <i>new</i> cannot be renamed because it is being used by the system or another process and the imple- mentation considers this to be an error.
601 602 603	[EEXIST] or [E	NOTEMPTY] The link named by <i>new</i> is a directory containing entries other than dot and dot-dot.
604 605	[EINVAL]	The <i>new</i> directory pathname contains a path prefix that names the <i>old</i> directory.
606 607	[EISDIR]	The <i>new</i> argument points to a directory, and the <i>old</i> argument points to a file that is not a directory.
608 609 610 611	[ENAMETOOL(DNG] The length of the <i>old</i> or <i>new</i> argument exceeds (PATH_MAX), or a pathname component is longer than {NAME_MAX} while {_POSIX_NO_TRUNC} is in effect.
612 613	[EMLINK]	The file named by <i>old</i> is a directory, and the link count of the parent directory of <i>new</i> would exceed {LINK_MAX}.
614 615	[ENOENT]	The link named by the <i>old</i> argument does not exist, or either <i>old</i> or <i>new</i> points to an empty string.
616	[ENOSPC]	The directory that would contain <i>new</i> cannot be extended.

- [ENOTDIR] A component of either path prefix is not a directory, or the old 617 argument names a directory and the new argument names a 618 nondirectory file. 619 The requested operation requires writing in a directory on a 620 [EROFS] read-only file system. 621 [EXDEV] The links named by *new* and *old* are on different file systems, 622 and the implementation does not support links between file 623 systems. 624
- 625 5.5.3.5 Cross-References
- 626 link(), 5.3.4; rmdir(), 5.5.2; unlink(), 5.5.1.
- 627 **5.6 File Characteristics**

628 5.6.1 File Characteristics: Header and Data Structure

The header <sys/stat.h> defines the structure *stat*, which includes the members shown in Table 5-1, returned by the functions *stat()* and *fstat()*.

Member Mem'zer Type Name		Description		
mode_t	st_mode	File mode (see 5.6.1.2).		
ino_t	st_ino	File serial number.		
dev_t	st_dev	ID of device containing this file.		
nlink_t	st_nlink	Number of links.		
uid_t	st_uid	User ID of the owner of the file.		
gid_t	st_gid	Group ID of the group of the file.		
off_t	st_size	For regular files, the file size in bytes. For other file types, the use of		
		this field is unspecified.		
time_t	st_atime	Time of last access.		
time_t	st_mtime	Time of last data modification.		
time_t	st_ctime	Time of last file status change.		

Table 5-1 - stat Structure

647 NOTE: File serial number and device ID taken together uniquely identify the file within the system.

All of the described members shall appear in the stat structure. The structure
members st_mode, st_ino, st_dev, st_uid, st_gid, st_atime, st_ctime, and st_mtime
shall have meaningful values for all file types defined in this part of ISO/IEC 9945.
The value of the member st_nlink shall be set to the number of links to the file.

652 **5.6.1.1 <sys/stat.h> File Types**

The following macros shall test whether a file is of the specified type. The value m supplied to the macros is the value of st_mode from a *stat* structure. The macro evaluates to a nonzero value if the test is true, zero if the test is false.

656S_ISDIR(m)Test macro for a directory file.657S_ISCHR(m)Test macro for a character special file.658S_ISBLK(m)Test macro for a block special file.659S_ISREG(m)Test macro for a regular file.

661 **5.6.1.2 <sys/stat.h> File Modes**

 $S_{ISFIFO}(m)$

660

The file modes portion of values of type $mode_t$, such as the st_mode value, are bit-encoded with the following masks and bits:

Test macro for a pipe or a FIFO special file.

664 665	S_IRWXU	Read, write, se sions mask for	earch (if a directory), or execute (otherwise) permis- the file owner class.
666		S_IRUSR	Read permission bit for the file owner class.
667		S_IWUSR	Write permission bit for the file owner class.
668 669		S_IXUSR	Search (if a directory) or execute (otherwise) per- missions bit for the file owner class.
670 671	S_IRWXG	Read, write, se sions mask for	arch (if a directory), or execute (otherwise) permis- the file group class.
672		S_IRGRP	Read permission bit for the file group class.
673		S_IWGRP	Write permission bit for the file group class.
674 675		S_IXGRP	Search (if a directory) or execute (otherwise) per- missions bit for the file group class.
676 677	S_IRWXO	Read, write, se sions mask for	arch (if a directory), or execute (otherwise) permis- the file other class.
678		S_IROTH	Read permission bit for the file other class.
679		S_IWOTH	Write permission bit for the file other class.
680 681		S_IXOTH	Search (if a directory) or execute (otherwise) per- missions bit for the file other class.
682 683 684 685	S_ISUID	Set user ID or shall be set to a program (see on any write.	n execution. The effective user ID of the process that of the owner of the file when the file is run as $exec$). On a regular file, this bit should be cleared
686 687 688	S_ISGID	Set group ID or the group of th On a regular fi	n execution. Set effective group ID on the process to ne file when the file is run as a program (see <i>exec</i>). le, this bit should be cleared on any write.

Part 1: SYSTEM API [C LANGUAGE]

The bits defined by S_IRUSR, S_IWUSR, S_IXUSR, S_IRGRP, S_IWGRP, S_IXGRP, 689 S_IROTH, S_IWOTH, S_IXOTH, S_ISUID, and S_ISGID shall be unique. S_IRWXU 690 shall be the bitwise inclusive OR of S_IRUSR, S_IWUSR, and S_IXUSR. S_IRWXG 691 shall be the bitwise inclusive OR of S_IRGRP, S_IWGRP, and S_IXGRP. S_IRWXO 692 shall be the bitwise inclusive OR of S_IROTH, S_IWOTH, and S_IXOTH. Imple-693 mentations may OR other implementation-defined bits into S_IRWXU, S_IRWXG, 694 and S_IRWXO, but they shall not overlap any of the other bits defined in this part 695 of ISO/IEC 9945. The file permission bits are defined to be those corresponding to 696 the bitwise inclusive OR of S_IRWXU, S_IRWXG, and S_IRWXO. 697

698 5.6.1.3 <sys/stat.h> Time Entries

- 699 The time-related fields of *struct stat* are as follows:
- 700 *st_atime* Accessed file data, for example, *read()*.
- 701 st_mtime Modified file data, for example, write().
- 702 st_ctime Changed file status, for example, chmod().
- These times are updated as described in 2.3.5.
- 704
- 705 Times are given in seconds since the Epoch.

706 5.6.1.4 Cross-References

chmod(), 5.6.4; chown(), 5.6.5; creat(), 5.3.2; exec, 3.1.2; link(), 5.3.4; mkdir(),
5.4.1; mkfifo(), 5.4.2; pipe(), 6.1.1; read(), 6.4.1; unlink(), 5.5.1; utime(), 5.6.6;

- 709 write(), 6.4.2; remove() [C Standard {2}].
- 710 5.6.2 Get File Status
- 711 Functions: stat(), fstat()
- 712 **5.6.2.1 Synopsis**
- 713 #include <sys/types.h>
- 714 #include <sys/stat.h>
- 715 int stat(const char *path, struct stat *buf);

716 int fstat(int fildes, struct stat *buf);

717 **5.6.2.2 Description**

The *path* argument points to a pathname naming a file. Read, write, or execute permission for the named file is not required, but all directories listed in the pathname leading to the file must be searchable. The *stat()* function obtains information about the named file and writes it to the area pointed to by the *buf* argument.

Similarly, the *fstat()* function obtains information about an open file known by thefile descriptor *fildes*.

I

Į

ł

An implementation that provides additional or alternate file access control mechanisms may, under implementation-defined conditions, cause the *stat()* and *fstat()* functions to fail. In particular, the system may deny the existence of the file specified by *path*.

- Both functions update any time-related fields, as described in 2.3.5, before writing
 into the *stat* structure.
- The *buf* is taken to be a pointer to a *stat* structure, as defined in the header (sys/stat.h>, into which information is placed concerning the file.
- 732 **5.6.2.3 Returns**
- Upon successful completion, a value of zero shall be returned. Otherwise, a value of -1 shall be returned and *errno* shall be set to indicate the error.
- 735 **5.6.2.4 Errors**
- If any of the following conditions occur, the stat() function shall return -1 and set *errno* to the corresponding value:
- 738 [EACCES] Search permission is denied for a component of the path prefix.
- 739 [ENAMETOOLONG]
- 740

741

- The length of the *path* argument exceeds {PATH_MAX}, or a pathname component is longer than {NAME_MAX} while {_POSIX_NO_TRUNC} is in effect.
- 743 [ENOENT] The named file does not exist, or the *path* argument points to 744 an empty string.
- 745 [ENOTDIR] A component of the path prefix is not a directory.

If any of the following conditions occur, the fstat() function shall return -1 and set *errno* to the corresponding value:

- 748 [EBADF] The *fildes* argument is not a valid file descriptor.
- 749 **5.6.2.5 Cross-References**
- ⁷⁵⁰ creat(), 5.3.2; dup(), 6.2.1; fcntl(), 6.5.2; open(), 5.3.1; pipe(), 6.1.1;
- 751 <sys/stat.h>, 5.6.1.
- 752 5.6.3 Check File Accessibility
- 753 Function: access()
- 754 **5.6.3.1 Synopsis**
- 755 #include <unistd.h>
- 756 int access(const char *path, int amode);

757 **5.6.3.2 Description**

The *access()* function checks the accessibility of the file named by the pathname pointed to by the *path* argument for the file access permissions indicated by *amode*, using the real user ID in place of the effective user ID and the real group ID in place of the effective group ID.

The value of *amode* is either the bitwise inclusive OR of the access permissions to be checked (R_OK, W_OK, and X_OK) or the existence test (F_OK). See 2.9.1 for the description of these symbolic constants.

If any access permission is to be checked, each shall be checked individually, as
 described in 2.3.2. If the process has appropriate privileges, an implementation
 may indicate success for X_OK even if none of the execute file permission bits are
 set.

769 **5.6.3.3 Returns**

If the requested access is permitted, a value of zero shall be returned. Otherwise,
a value of -1 shall be returned and *errno* shall be set to indicate the error.

772 **5.6.3.4 Errors**

If any of the following conditions occur, the access() function shall return -1 and set *errno* to the corresponding value:

775[EACCES]The permissions specified by amode are denied, or search per-776mission is denied on a component of the path prefix.

777 [ENAMETOOLONG]

778The length of the path argument exceeds {PATH_MAX}, or a779pathname component is longer than {NAME_MAX} while780{_POSIX_NO_TRUNC} is in effect.

- [ENOENT] The *path* argument points to an empty string or to the name of
 a file that does not exist.
- 783 [ENOTDIR] A component of the path prefix is not a directory.
- [EROFS] Write access was requested for a file residing on a read-only file
 system.
- For each of the following conditions, if the condition is detected, the *access()* function shall return -1 and set *errno* to the corresponding value:
- 788 [EINVAL] An invalid value was specified for amode.
- 789 **5.6.3.5 Cross-References**
- 790 chmod(), 5.6.4; stat(), 5.6.2; <unistd.h>, 2.9.

I

791 **5.6.4 Change File Modes**

- 792 Function: chmod()
- 793 **5.6.4.1 Synopsis**
- 794 #include <sys/types.h>
- 795 #include <sys/stat.h>
- 796 int chmod(const char *path, mode_t mode);

797 **5.6.4.2 Description**

The *path* argument shall point to a pathname naming a file. If the effective user 798 ID of the calling process matches the file owner or the calling process has 799 appropriate privileges, the *chmod()* function shall set the S_ISUID, S_ISGID, and 800 the file permission bits, as described in 5.6.1, of the named file from the 801 corresponding bits in the *mode* argument. These bits define access permissions 802 for the user associated with the file, the group associated with the file, and all oth-803 ers, as described in 2.3.2. Additional implementation-defined restrictions may 804 cause the S_ISUID and S_ISGID bits in mode to be ignored. 805

If the calling process does not have appropriate privileges, if the group ID of the file does not match the effective group ID or one of the supplementary group IDs, and if the file is a regular file, bit S_ISGID (set group ID on execution) in the mode of the file shall be cleared upon successful return from *chmod*().

- The effect on file descriptors for files open at the time of the *chmod()* function is implementation defined.
- ⁸¹² Upon successful completion, the chmod() function shall mark for update the ⁸¹³ st_ctime field of the file.

814 **5.6.4.3 Returns**

⁸¹⁵ Upon successful completion, the function shall return a value of zero. Otherwise, ⁸¹⁶ a value of -1 shall be returned and *errno* shall be set to indicate the error. If -1 is

returned, no change to the file mode shall have occurred.

818 **5.6.4.4 Errors**

- If any of the following conditions occur, the chmod() function shall return -1 and set *errno* to the corresponding value:
- 821 [EACCES] Search permission is denied on a component of the path prefix.
- 822 [ENAMETOOLONG]
- 823The length of the path argument exceeds {PATH_MAX}, or a824pathname component is longer than {NAME_MAX} while825{_POSIX_NO_TRUNC} is in effect.
- 826 [ENOTDIR] A component of the path prefix is not a directory.

- 827[ENOENT]The named file does not exist or the *path* argument points to an
empty string.
- 829[EPERM]The effective user ID does not match the owner of the file, and
the calling process does not have the appropriate privileges.
- 831 [EROFS] The named file resides on a read-only file system.

832 5.6.4.5 Cross-References

- s33 chown(), 5.6.5; mkdir(), 5.4.1; mkfifo(), 5.4.2; stat(), 5.6.2; <sys/stat.h>, 5.6.1.
- 834 **5.6.5 Change Owner and Group of a File**
- 835 Function: chown()
- 836 **5.6.5.1 Synopsis**
- 837 #include <sys/types.h>
- 838 int chown(const char *path, uid_t owner, gid_t group);

839 **5.6.5.2 Description**

- The path argument points to a pathname naming a file. The user ID and group ID of the named file are set to the numeric values contained in *owner* and *group* respectively.
- 843 Only processes with an effective user ID equal to the user ID of the file or 844 with appropriate privileges may change the ownership of a file. If 845 {_POSIX_CHOWN_RESTRICTED} is in effect for *path*:
- 846 (1) Changing the owner is restricted to processes with appropriate847 privileges.
- (2) Changing the group is permitted to a process without appropriate
 privileges, but with an effective user ID equal to the user ID of the file, if
 and only if *owner* is equal to the user ID of the file and *group* is equal
 either to the effective group ID of the calling process or to one of its supplementary group IDs.
- If the *path* argument refers to a regular file, the set-user-ID (S_ISUID) and setgroup-ID (S_ISGID) bits of the file mode shall be cleared upon successful return from *chown*(), unless the call is made by a process with appropriate privileges, in which case it is implementation defined whether those bits are altered. If the *chown*() function is successfully invoked on a file that is not a regular file, these bits may be cleared. These bits are defined in 5.6.1.
- Upon successful completion, the chown() function shall mark for update the st_ctime field of the file.

861 **5.6.5.3 Returns**

⁸⁶² Upon successful completion, a value of zero shall be returned. Otherwise, a value ⁸⁶³ of -1 shall be returned and *errno* shall be set to indicate the error. If -1 is ⁸⁶⁴ returned, no change shall be made in the owner and group of the file.

865 **5.6.5.4 Errors**

If any of the following conditions occur, the chown() function shall return -1 and set *errno* to the corresponding value:

868 [EACCES] Search permission is denied on a component of the path prefix.

869 [ENAMETOOLONG]

- 870The length of the path argument exceeds {PATH_MAX}, or a871pathname component is longer than {NAME_MAX} while872{_POSIX_NO_TRUNC} is in effect.
- 873 [ENOTDIR] A component of the path prefix is not a directory.
- [ENOENT] The named file does not exist, or the *path* argument points to an empty string.
- 876[EPERM]The effective user ID does not match the owner of the file, or the
calling process does not have appropriate privileges and
{_POSIX_CHOWN_RESTRICTED} indicates that such privilege is
required.
- 880 [EROFS] The named file resides on a read-only file system.

For each of the following conditions, if the condition is detected, the chown() function shall return -1 and set errno to the corresponding value:

- 883 [EINVAL] The owner or group ID supplied is invalid and not supported by 884 the implementation.
- 885 5.6.5.5 Cross-References
- 886 chmod(), 5.6.4; <sys/stat.h>, 5.6.1.

5.6.6 Set File Access and Modification Times

888 Function: *utime()*

889 **5.6.6.1 Synopsis**

890 #include <sys/types.h>
891 #include <utime.h>

892 int utime(const char *path, const struct utimbuf *times);

ł

893 **5.6.6.2 Description**

The argument *path* points to a pathname naming a file. The *utime()* function sets the access and modification times of the named file.

If the *times* argument is NULL, the access and modification times of the file are set to the current time. The effective user ID of the process must match the owner of the file, or the process must have write permission to the file or appropriate privileges, to use the *utime()* function in this manner.

900 If the *times* argument is not NULL, it is interpreted as a pointer to a *utimbuf* 901 structure, and the access and modification times are set to the values contained in 902 the designated structure. Only the owner of the file and processes with appropri-903 ate privileges shall be permitted to use the *utime()* function in this way.

The *utimbuf* structure is defined by the header <utime.h> and includes the following members:

906 907	Member <u>Type</u>	Member Name	Description
908	time_t	actime	Access time
909	time_t	modtime	Modification time

910 The times in the *utimbuf* structure are measured in seconds since the Epoch.

Implementations may add extensions as permitted in 1.3.1.1, point (2). Adding
extensions to this structure, which might change the behavior of the application
with respect to this standard when those fields in the structure are uninitialized,
also requires that the extensions be enabled as required by 1.3.1.1.

915 Upon successful completion, the utime() function shall mark for update the 916 st_ctime field of the file.

917 **5.6.6.3 Returns**

Upon successful completion, the function shall return a value of zero. Otherwise,
a value of -1 shall be returned, *errno* is set to indicate the error, and the file times
shall not be affected.

921 5.6.6.4 Errors

If any of the following conditions occur, the utime() function shall return -1 and set *errno* to the corresponding value:

924[EACCES]Search permission is denied by a component of the path prefix,925or the *times* argument is NULL and the effective user ID of the926process does not match the owner of the file and write access is927denied.

928 [ENAMETOOLONG]

929	The length	of	the	path	ar	gument	exceeds	s {PATH_I	MAX},	or	а
930	pathname	cor	npor	nent	is	longer	than	NAME_M	AX}	whil	le
931	{_POSIX_N	D_T	RUN	C} is	in e	ffect.					

932 933	[ENOENT]	The named file does not exist or the <i>path</i> argument points to an empty string.
934	[ENOTDIR]	A component of the path prefix is not a directory.
935 936 937 938	[EPERM]	The <i>times</i> argument is not NULL, the effective user ID of the calling process has write access to the file, but does not match the owner of the file, and the calling process does not have the appropriate privileges.
939	[EROFS]	The <i>named</i> file resides on a read-only file system.

- 940 **5.6.6.5 Cross-References**
- 941 <sys/stat.h>, 5.6.1.

942 **5.7 Configurable Pathname Variables**

- 943 5.7.1 Get Configurable Pathname Variables
- 944 Functions: pathconf(), fpathconf()
- 945 **5.7.1.1 Synopsis**
- 946 #include <unistd.h>
- 947 long pathconf(const char *path, int name);
- 948 long fpathconf(int fildes, int name);

949 **5.7.1.2 Description**

- The *pathconf()* and *fpathconf()* functions provide a method for the application to determine the current value of a configurable limit or option (*variable*) that is associated with a file or directory.
- For *pathconf()*, the *path* argument points to the pathname of a file or directory. For *fpathconf()*, the *fildes* argument is an open file descriptor.
- The name argument represents the variable to be queried relative to that file or directory. The implementation shall support all of the variables listed in Table 5-2 and may support others. The variables in Table 5-2 come from limits.h> or <unistd.h> and the symbolic constants, defined in <unistd.h>, that are the corresponding values used for name.
- 960 **5.7.1.3 Returns**
- If name is an invalid value, the pathconf() and fpathconf() functions shall return -1.
- If the variable corresponding to name has no limit for the path or file descriptor,
 the pathconf() and fpathconf() functions shall return -1 without changing errno.

965

	Variable	name Value	Notes
	(LINK_MAX)	{_PC_LINK_MAX}	(1)
	{MAX_CANON}	{_PC_MAX_CANON}	(2)
	{MAX_INPUT}	{_PC_MAX_INPUT}	(2)
	{NAME_MAX}	{_PC_NAME_MAX}	(3), (4)
	{PATH_MAX}	{_PC_PATH_MAX}	(4), (5)
	(PIPE_BUF)	{_PC_PIPE_BUF}	(6)
	{_POSIX_CHOWN_RESTRICTED}	{_PC_CHOWN_RESTRICTED}	(7)
	{_POSIX_NO_TRUNC}	{_PC_NO_TRUNC}	(3, 4)
	{_POSIX_VDISABLE}	{_PC_VDISABLE}	(2)
OTES	5:		
(1)	If <i>path</i> or <i>fildes</i> refers to a directory, th	e value returned applies to the	e directory itself.
(2)	If <i>path</i> or <i>fildes</i> does not refer to a term supports an association of the variable	ninal file, it is unspecified whet name with the specified file.	her an implementatior
(3)	If <i>path</i> or <i>fildes</i> refers to a directory, the directory.	he value returned applies to th	e filenames within the
(4)	If <i>path</i> or <i>fildes</i> does not refer to a dir supports an association of the variable	rectory, it is unspecified whet name with the specified file.	her an implementatior
(5)	If <i>path</i> or <i>fildes</i> refers to a directory, th pathname when the specified directory	e value returned is the maxim is the working directory.	um length of a relative
(6)	If <i>path</i> refers to a FIFO, or <i>fildes</i> refers referenced object itself. If <i>path</i> or <i>fild</i> any FIFOs that exist or can be created other type of file, it is unspecified whet variable name with the specified file.	to a pipe or a FIFO, the value es refers to a directory, the va d within the directory. If path ther an implementation suppor	returned applies to the lue returned applies to a or <i>fildes</i> refers to any ts an association of the
	If path or fildes refers to a directory, t	he value returned applies to a	ny files defined in this

Table 5-2 Configurable Pathname Variables

⁹⁹⁶ If the implementation needs to use *path* to determine the value of *name* and the ⁹⁹⁷ implementation does not support the association of *name* with the file specified by ⁹⁹⁸ *path*, or if the process did not have the appropriate privileges to query the file ⁹⁹⁹ specified by *path*, or *path* does not exist, the *pathconf()* function shall return -1.

1000 If the implementation needs to use *fildes* to determine the value of *name* and the 1001 implementation does not support the association of *name* with the file specified by 1002 *fildes*, or if *fildes* is an invalid file descriptor, the *fpathconf()* function shall 1003 return -1.

Otherwise, the *pathconf()* and *fpathconf()* functions return the current variable value for the file or directory without changing *errno*. The value returned shall not be more restrictive than the corresponding value described to the application when it was compiled with the implementation's <limits.h> or <unistd.h>.

- 1008 **5.7.1.4 Errors**
- If any of the following conditions occur, the *pathconf()* and *fpathconf()* functions
 shall return -1 and set *errno* to the corresponding value:
- 1011 [EINVAL] The value of *name* is invalid.

For each of the following conditions, if the condition is detected, the *pathconf()* function shall return -1 and set *errno* to the corresponding value:

1014[EACCES]Search permission is denied for a component of the path prefix.1015[EINVAL]The implementation does not support an association of the vari-1016able name with the specified file.

1017 [ENAMETOOLONG]

- 1018The length of the path argument exceeds {PATH_MAX}, or a1019pathname component is longer than {NAME_MAX} while1020{_POSIX_NO_TRUNC} is in effect.
- 1021[ENOENT]The named file does not exist, or the path argument points to
an empty string.
- 1023 [ENOTDIR] A component of the path prefix is not a directory.

For each of the following conditions, if the condition is detected, the *fpathconf()* function shall return -1 and set *errno* to the corresponding value:

- 1026 [EBADF] The *fildes* argument is not a valid file descriptor.
- 1027[EINVAL]The implementation does not support an association of the vari-1028able name with the specified file.

L

L

Section 6: Input and Output Primitives

1 The functions in this section deal with input and output from files and pipes.

2 Functions are also specified that deal with the coordination and management of

3 file descriptors and I/O activity.

4 **6.1 Pipes**

5 6.1.1 Create an Inter-Process Channel

- 6 Function: *pipe()*
- 7 **6.1.1.1 Synopsis**
- 8 int pipe(int fildes[2]);

9 **6.1.1.2 Description**

The *pipe()* function shall create a pipe and place two file descriptors, one each into the arguments *fildes[0]* and *fildes[1]*, that refer to the open file descriptions for the read and write ends of the pipe. Their integer values shall be the two lowest available at the time of the *pipe()* function call. The O_NONBLOCK and FD_CLOEXEC flags shall be clear on both file descriptors. [The *fcntl()* function can be used to set these flags.]

Data can be written to file descriptor *fildes*[1] and read from file descriptor
 fildes[0]. A read on file descriptor *fildes*[0] shall access the data written to file
 descriptor *fildes*[1] on a first-in-first-out basis.

A process has the pipe open for reading if it has a file descriptor open that refers to the read end, *fildes*[0]. A process has the pipe open for writing if it has a file descriptor open that refers to the write end, *fildes*[1].

Upon successful completion, the *pipe()* function shall mark for update the st_atime, st_ctime, and st_mtime fields of the pipe.

24 6.1.1.3 Returns

Upon successful completion, the function shall return a value of zero. Otherwise,
a value of -1 shall be returned and *errno* shall be set to indicate the error.

27 **6.1.1.4 Errors**

- If any of the following conditions occur, the *pipe()* function shall return -1 and set *errno* to the corresponding value:
- 30[EMFILE]More than {OPEN_MAX}-2 file descriptors are already in use by31this process.
- [ENFILE] The number of simultaneously open files in the system would
 exceed a system-imposed limit.
- 34 6.1.1.5 Cross-References
- *fcntl()*, 6.5.2; *open()*, 5.3.1; *read()*, 6.4.1; *write()*, 6.4.2.

36 6.2 File Descriptor Manipulation

37 6.2.1 Duplicate an Open File Descriptor

- 38 Functions: dup(), dup2()
- 39 **6.2.1.1 Synopsis**
- 40 int dup(int fildes);
- 41 int dup2(int fildes, int fildes2);

42 **6.2.1.2 Description**

The dup() and dup2() functions provide an alternate interface to the service provided by the *fcntl()* function using the F_DUPFD command. The call:

- 45 fid = dup (fildes);
- 46 shall be equivalent to:

fid = fcntl (fildes, F_DUPFD, 0);

48 The call:

47

49

fid = dup2 (fildes, fildes2);

50 shall be equivalent to:

```
51 close (fildes2);
52 fid = fcntl (fildes, F DUPFD, fildes2);
```

- 53 except for the following:
- (1) If *fildes2* is negative or greater than or equal to {OPEN_MAX}, the *dup2()* function shall return -1 and *errno* shall be set to [EBADF].
- If *fildes* is a valid file descriptor and is equal to *fildes2*, the *dup2()* function shall return *fildes2* without closing it.

- (3) If *fildes* is not a valid file descriptor, *dup2()* shall fail and not close
 fildes2.
- (4) The value returned shall be equal to the value of *fildes2* upon successful
 completion or shall be -1 upon failure.

62 6.2.1.3 Returns

⁶³ Upon successful completion, the function shall return a file descriptor. Other-⁶⁴ wise, a value of -1 shall be returned and *errno* shall be set to indicate the error.

65 **6.2.1.4 Errors**

- If any of the following conditions occur, the dup() function shall return -1 and set errno to the corresponding value:
- 68 [EBADF] The argument *fildes* is not a valid open file descriptor.

69 [EMFILE] The number of file descriptors would exceed {OPEN_MAX}.

If any of the following conditions occur, the *dup2()* function shall return -1 and
set *errno* to the corresponding value:

- 72[EBADF]The argument fildes is not a valid open file descriptor, or the
argument fildes2 is negative or greater than or equal to
{OPEN_MAX}.
- 75 [EINTR] The *dup2()* function was interrupted by a signal.

76 6.2.1.5 Cross-References

close(), 6.3.1; creat(), 5.3.2; exec, 3.1.2; fcntl(), 6.5.2; open(), 5.3.1; pipe(), 6.1.1.

78 6.3 File Descriptor Deassignment

- 79 **6.3.1 Close a File**
- 80 Function: close()

81 **6.3.1.1 Synopsis**

82 int close(int fildes);

83 **6.3.1.2 Description**

The *close()* function shall deallocate (i.e., make available for return by subsequent *open()*s, etc., executed by the process) the file descriptor indicated by *fildes*. All outstanding record locks owned by the process on the file associated with the file descriptor shall be removed (that is, unlocked). ł

ł

- If the *close()* function is interrupted by a signal that is to be caught, it shall return -1 with *errno* set to [EINTR], and the state of *fildes* is unspecified.
- When all file descriptors associated with a pipe or FIFO special file have been closed, any data remaining in the pipe or FIFO shall be discarded.
- When all file descriptors associated with an open file description have been closed,
 the open file description shall be freed.

If the link count of the file is zero, when all file descriptors associated with the file
have been closed, the space occupied by the file shall be freed and the file shall no
longer be accessible.

97 **6.3.1.3 Returns**

⁹⁸ Upon successful completion, a value of zero shall be returned. Otherwise, a value
 ⁹⁹ of -1 shall be returned and *errno* shall be set to indicate the error.

100 **6.3.1.4 Errors**

If any of the following conditions occur, the *close()* function shall return -1 and
 set *errno* to the corresponding value:

- 103 [EBADF] The *fildes* argument is not a valid file descriptor.
- 104 [EINTR] The *close* function was interrupted by a signal.
- 105 6.3.1.5 Cross-References

creat(), 5.3.2; dup(), 6.2.1; exec, 3.1.2; fcntl(), 6.5.2; fork(), 3.1.1; open(), 5.3.1;
 pipe(), 6.1.1; unlink(), 5.5.1; 3.3.1.4.

- 108 6.4 Input and Output
- 109 6.4.1 Read from a File
- 110 Function: read()
- 111 **6.4.1.1 Synopsis**
- 112 ssize_t read(int fildes, void *buf, size_t nbyte);
- 113 **6.4.1.2 Description**

The *read()* function shall attempt to read *nbyte* bytes from the file associated with the open file descriptor, *fildes*, into the buffer pointed to by *buf*.

116 If *nbyte* is zero, the *read()* function shall return zero and have no other results.

On a regular file or other file capable of seeking, *read()* shall start at a position in the file given by the file offset associated with *fildes*. Before successful return

I

- from *read()*, the file offset shall be incremented by the number of bytes actually read.
- 121 On a file not capable of seeking, the *read()* shall start from the current position. 122 The value of a file offset associated with such a file is undefined.

Upon successful completion, the *read()* function shall return the number of bytes actually read and placed in the buffer. This number shall never be greater than *nbyte*. The value returned may be less than *nbyte* if the number of bytes left in the file is less than *nbyte*, if the *read()* request was interrupted by a signal, or if the file is a pipe (or FIFO) or special file and has fewer than *nbyte* bytes immediately available for reading. For example, a *read()* from a file associated with a terminal may return one typed line of data.

- 130 If a read() is interrupted by a signal before it reads any data, it shall return -1131 with *errno* set to [EINTR].
- If a read() is interrupted by a signal after it has successfully read some data,
 either it shall return -1 with errno set to [EINTR], or it shall return the number of
 bytes read. A read() from a pipe or FIFO shall never return with errno set to
 [EINTR] if it has transferred any data.
- No data transfer shall occur past the current end-of-file. If the starting position is
 at or after the end-of-file, zero shall be returned. If the file refers to a device special file, the result of subsequent read() requests is implementation defined.
- 139 If the value of *nbyte* is greater than (SSIZE_MAX), the result is implementation 140 defined.
- 141 When attempting to read from an empty pipe (or FIFO):
- (1) If no process has the pipe open for writing, read() shall return zero to
 indicate end-of-file.
- 144 (2) If some process has the pipe open for writing and O_NONBLOCK is set,
 145 read() shall return -1 and set errno to [EAGAIN].
- If some process has the pipe open for writing and O_NONBLOCK is clear,
 read() shall block until some data is written or the pipe is closed by all
 processes that had the pipe open for writing.
- When attempting to read a file (other than a pipe or FIFO) that supports nonblocking reads and has no data currently available:
- 151 (1) If O_NONBLOCK is set, read() shall return -1 and set errno to [EAGAIN].
- 152 (2) If O_NONBLOCK is clear, read() shall block until some data becomes
 153 available.
- 154 The use of the O_NONBLOCK flag has no effect if there is some data available.
- For any portion of a regular file, prior to the end-of-file, that has not been written, *read()* shall return bytes with value zero.
- Upon successful completion where *nbyte* is greater than zero, the *read()* function
 shall mark for update the *st_atime* field of the file.

I

159 **6.4.1.3 Returns**

160 Upon successful completion, read() shall return an integer indicating the number 161 of bytes actually read. Otherwise, read() shall return a value of -1 and set *errno* 162 to indicate the error, and the content of the buffer pointed to by *buf* is 163 indeterminate.

164 **6.4.1.4 Errors**

If any of the following conditions occur, the *read()* function shall return -1 and set
 errno to the corresponding value:

- 167[EAGAIN]The O_NONBLOCK flag is set for the file descriptor and the pro-168cess would be delayed in the read operation.
- 169 [EBADF] The *fildes* argument is not a valid file descriptor open for 170 reading.
- 171[EINTR]The read operation was interrupted by a signal, and either no172data was transferred or the implementation does not report173partial transfer for this file.
- 174[EIO]The implementation supports job control, the process is in a175background process group and is attempting to read from its176controlling terminal, and either the process is ignoring or block-177ing the SIGTTIN signal or the process group of the process is178orphaned. This error may also be generated when conditions179unspecified by this part of ISO/IEC 9945 occur.

180 6.4.1.5 Cross-References

- creat(), 5.3.2; dup(), 6.2.1; fcntl(), 6.5.2; lseek(), 6.5.3; open(), 5.3.1; pipe(), 6.1.1;
 3.3.1.4; 7.1.1.
- 183 **6.4.2 Write to a File**
- 184 Function: write()
- 185 **6.4.2.1 Synopsis**
- 186 ssize_t write(int fildes, const void *buf, size_t nbyte);
- 187 **6.4.2.2 Description**

The *write()* function shall attempt to write *nbyte* bytes from the buffer pointed to by *buf* to the file associated with the open file descriptor, *fildes*.

If *nbyte* is zero and the file is a regular file, the *write()* function shall return zero and have no other results. If *nbyte* is zero and the file is not a regular file, the results are unspecified.

On a regular file or other file capable of seeking, the actual writing of data shall proceed from the position in the file indicated by the file offset associated with

- *fildes.* Before successful return from *write()*, the file offset shall be incremented
 by the number of bytes actually written. On a regular file, if this incremented file
 offset is greater than the length of the file, the length of the file shall be set to this
 file offset.
- On a file not capable of seeking, the *write()* shall start from the current position.The value of a file offset associated with such a file is undefined.
- If the O_APPEND flag of the file status flags is set, the file offset shall be set to the end of the file prior to each write, and no intervening file modification operation shall be allowed between changing the file offset and the write operation.
- If a *write()* requests that more bytes be written than there is room for (for example, the physical end of a medium), only as many bytes as there is room for shall be written. For example, suppose there is space for 20 bytes more in a file before reaching a limit. A write of 512 bytes would return 20. The next write of a nonzero number of bytes would give a failure return (except as noted below).
- Upon successful completion, the *write()* function shall return the number of bytes actually written to the file associated with *fildes*. This number shall never be greater than *nbyte*.
- If a write() is interrupted by a signal before it writes any data, it shall return -1 with *errno* set to [EINTR].
- If write() is interrupted by a signal after it successfully writes some data, either it
 shall return -1 with errno set to [EINTR], or it shall return the number of bytes
 written. A write() to a pipe or FIFO shall never return with errno set to [EINTR] if
 it has transferred any data and nbyte is less than or equal to (PIPE_BUF).
- If the value of *nbyte* is greater than (SSIZE_MAX), the result is implementation defined.
- After a *write()* to a regular file has successfully returned:
 - (1) Any successful read() from each byte position in the file that was modified by that write() shall return the data specified by the write() for that position, until such byte positions are again modified.
- (2) Any subsequent successful *write()* to the same byte position in the file
 shall overwrite that file data. The phrase "subsequent successful *write()*"
 in the previous sentence is intended to be viewed from a system perspective [i.e., *read()* followed by a systemwide subsequent *write()*].
- Write requests to a pipe (or FIFO) shall be handled in the same manner as write requests to a regular file, with the following exceptions:
 - (1) There is no file offset associated with a pipe, hence each write request shall append to the end of the pipe.
- (2) Write requests of {PIPE_BUF} bytes or less shall not be interleaved with data from other processes doing writes on the same pipe. Writes of greater than {PIPE_BUF} bytes may have data interleaved, on arbitrary boundaries, with writes by other processes, whether or not the O_NONBLOCK flag of the file status flags is set.

223

230

243

244

245

246

247

252

253

- (3) If the O_NONBLOCK flag is clear, a write request may cause the process
 to block, but on normal completion it shall return *nbyte*.
- (4) If the O_NONBLOCK flag is set, write() requests shall be handled dif ferently, in the following ways:
 - (a) The *write*() function shall not block the process.
- 242 (b) A write request for {PIPE_BUF} or fewer bytes shall either:
 - [1] If there is sufficient space available in the pipe, transfer all the data and return the number of bytes requested.
 - [2] If there is not sufficient space available in the pipe, transfer no data and return -1 with *errno* set to [EAGAIN].
 - (c) A write request for more than {PIPE_BUF} bytes shall either:
- 248[1]When at least one byte can be written, transfer what it can
and return the number of bytes written. When all data previ-
ously written to the pipe has been read, it shall transfer at
least {PIPE_BUF} bytes.
 - [2] When no data can be written, transfer no data and return -1 with *errno* set to [EAGAIN].

When attempting to write to a file descriptor (other than a pipe or FIFO) that supports nonblocking writes and cannot accept the data immediately:

- (1) If the O_NONBLOCK flag is clear, write() shall block until the data can be
 accepted.
- (2) If the O_NONBLOCK flag is set, write() shall not block the process. If
 some data can be written without blocking the process, write() shall
 write what it can and return the number of bytes written. Otherwise, it
 shall return -1 and errno shall be set to [EAGAIN].

Upon successful completion where *nbyte* is greater than zero, the *write()* function shall mark for update the *st_ctime* and *st_mtime* fields of the file.

264 **6.4.2.3 Returns**

Upon successful completion, write() shall return an integer indicating the number of bytes actually written. Otherwise, it shall return a value of -1 and set *errno* to indicate the error.

268 **6.4.2.4 Errors**

If any of the following conditions occur, the write() function shall return -1 and set *errno* to the corresponding value:

271 272	[EAGAIN]	The O_NONBLOCK flag is set for the file descriptor and the process would be delayed in the write operation.
273 274	[EBADF]	The <i>fildes</i> argument is not a valid file descriptor open for writing.

- 275 [EFBIG] An attempt was made to write a file that exceeds an 276 implementation-defined maximum file size.
- 277[EINTR]The write operation was interrupted by a signal, and either no
data was transferred or the implementation does not report
partial transfers for this file.
- 280[EIO]The implementation supports job control, the process is in a
background process group and is attempting to write to its con-
trolling terminal, TOSTOP is set, the process is neither ignoring
nor blocking SIGTTOU signals, and the process group of the pro-
cess is orphaned. This error may also be generated when condi-
tions unspecified by this part of ISO/IEC 9945 occur.
- 286 [ENOSPC] There is no free space remaining on the device containing the 287 file.
- 288[EPIPE]An attempt is made to write to a pipe (or FIFO) that is not open289for reading by any process. A SIGPIPE signal shall also be sent290to the process.
- 291 6.4.2.5 Cross-References

creat(), 5.3.2; *dup*(), 6.2.1; *fcntl*(), 6.5.2; *lseek*(), 6.5.3; *open*(), 5.3.1; *pipe*(), 6.1.1; 3.3.1.4.

294 **6.5 Control Operations on Files**

295 6.5.1 Data Definitions for File Control Operations

The header <fcntl.h> defines the following requests and arguments for the fcntl() and open() functions. The values within each of the tables within this clause (Table 6-1 through Table 6-7) shall be unique numbers. In addition, the values of the entries for oflag values, file status flags, and file access modes shall be unique.

- 301 6.5.2 File Control
- 302 Function: fcntl()
- 303 6.5.2.1 Synopsis
- 304 #include <sys/types.h>
- 305 #include <unistd.h>
- 306 #include <fcntl.h>
- 307 int fcntl(int fildes, int cmd, ...);

326

333

341 342

308		Table 6-1 - cmd Values for fcntl()	- cmd Values for $fcntl()$	
310	Constant	Description		
311	F_DUPFD	Duplicate file descriptor.		
312	F_GETFD	Get file descriptor flags.		
313	F_GETLK	Get record locking information.		
314	F_SETFD	Set file descriptor flags.		
315	F_GETFL	Get file status flags.		
316	F SETFL	Set file status flags.		
317	FSETLK	Set record locking information.		
318	F_setlkw	Set record locking information; wait if blocked.		
319		- /		

Table 6-1 - cmd Values for fcntl()

Table 6-2 – File Descriptor Flags Used for fcntl()

Description
e the file descriptor upon execution of an <i>exec</i> -family tion.
5

Table 6-3 – *l_type* Values for Record Locking With *fcntl()*

328	Constant	Description	
329	F_RDLCK	Shared or read lock.	
330	F_UNLCK	Unlock.	
331	F_WRLCK	Exclusive or write lock.	
332			

Table 6-4 - oflag Values for open()

335	Constant	Description	
336	O_CREAT	Create file if it does not exist.	
337	O_EXCL	Exclusive use flag.	
338	O_NOCTTY	Do not assign a controlling terminal.	
339	O_TRUNC	Truncate flag.	
340			

Table 6-5 – File Status Flags Used for open() and fcntl()

343	Constant	Description
344 345	O_APPEND O_NONBLOCK	Set append mode. No delay.
346		are delay.
347 348	Table 0-0 -	The Access modes Used for open() and jent()
------------	-------------	---
349	Constant	Description
350	O_RDONLY	Open for reading only.
351	O_RDWR	Open for reading and writing.
352	O_WRONLY	Open for writing only.
353		

Table 6-6 - File Access Modes Used for open() and fcntl()

Table 6-7 – Mask for Use With File Access Modes

000		
356	Constant	Description
357	O_ACCMODE	Mask for file access modes.
358		

359 **6.5.2.2 Description**

354

- The function *fcntl()* provides for control over open files. The argument *fildes* is a file descriptor.
- The available values for *cmd* are defined in the header <fcntl.h> (see 6.5.1), which shall include:
- 364F_DUPFDReturn a new file descriptor that is the lowest numbered avail-
able (i.e., not already open) file descriptor greater than or equal
to the third argument, arg, taken as an integer of type *int*. The
new file descriptor refers to the same open file description as
the original file descriptor and shares any locks.
- The FD_CLOEXEC flag associated with the new file descriptor is cleared to keep the file open across calls to the *exec* family of functions.
- 372F_GETFDGet the file descriptor flags, as defined in Table 6-2, that are373associated with the file descriptor fildes. File descriptor flags374are associated with a single file descriptor and do not affect375other file descriptors that refer to the same file.
- 376F_SETFDSet the file descriptor flags, as defined in Table 6-2, that are
associated with *fildes* to the third argument, *arg*, taken as type378int. If the FD_CLOEXEC flag is zero, the file shall remain open
across *exec* functions; otherwise, the file shall be closed upon
successful execution of an *exec* function.
- Get the file status flags, as defined in Table 6-5, and file access F_GETFL 381 modes for the open file description associated with *fildes*. The 382 file access modes defined in Table 6-6 can be extracted from the 383 return value using the mask O_ACCMODE, which is defined in 384 <fcntl.h>. File status flags and file access modes are associ-385 ated with the open file description and do not affect other file 386 descriptors that refer to the same file with different open file 387 descriptions. 388

389F_SETFLSet the file status flags, as defined in Table 6-5, for the open file
description associated with *fildes* from the corresponding bits in
the third argument, arg, taken as type *int*. Bits corresponding
392
393
393
394Table 6-6) and the oflag
are ignored.394If any bits in arg other than those mentioned here are changed
by the application, the result is unspecified.

The following commands are available for advisory record locking. Advisory record locking shall be supported for regular files, and may be supported for other files.

- 399 F_GETLK Get the first lock that blocks the lock description pointed to by
the third argument, arg, taken as a pointer to type struct flock400(see below). The information retrieved overwrites the informa-
tion passed to fcntl() in the flock structure. If no lock is found
that would prevent this lock from being created, the structure
shall be left unchanged by this function call except for the lock
type, which shall be set to F_UNLCK.
- F SETLK Set or clear a file segment lock according to the lock description 406 pointed to by the third argument, arg, taken as a pointer to 407 type struct flock (see below). F_SETLK is used to establish 408 409 shared (or read) locks (F_RDLCK) or exclusive (or write) locks, (F_WRLCK), as well as to remove either type of lock (F_UNLCK). 410 F RDLCK, F WRLCK, and F UNLCK are defined by the 411 <fcntl.h> header. If a shared or exclusive lock cannot be set, 412 *fcntl()* shall return immediately. 413
- 414F_SETLKWThis command is the same as F_SETLK except that if a shared415or exclusive lock is blocked by other locks, the process shall416wait until the request can be satisfied. If a signal that is to be417caught is received while fcntl() is waiting for a region, the418fcntl() shall be interrupted. Upon return from the signal419handler of the process, fcntl() shall return -1 with errno set to420[EINTR], and the lock operation shall not be done.
- The *flock* structure, defined by the <fcntl.h> header, describes an advisory lock.
 It includes the members shown in Table 6-8.

When a shared lock has been set on a segment of a file, other processes shall be able to set shared locks on that segment or a portion of it. A shared lock prevents any other process from setting an exclusive lock on any portion of the protected area. A request for a shared lock shall fail if the file descriptor was not opened with read access.

An exclusive lock shall prevent any other process from setting a shared lock or an exclusive lock on any portion of the protected area. A request for an exclusive lock shall fail if the file descriptor was not opened with write access.

The value of l_whence is SEEK_SET, SEEK_CUR, or SEEK_END to indicate that the relative offset, l_start bytes, will be measured from the start of the file, current position, or end of the file, respectively. The value of l_len is the number of consecutive bytes to be locked. If l_len is negative, the result is undefined. The

Т

Member Type	Member Name	Description
short	l_type	F_RDLCK, F_WRLCK, or F_UNLCK.
short	l_whence	Flag for starting offset.
off_t	l_start	Relative offset in bytes.
off_t	l_len	Size; if 0, then until EOF.
pid_t	l_pid	Process ID of the process holding the lock, returned with F_GETLI

Table 6-8 - flock Structure

l_pid field is only used with F_GETLK to return the process ID of the process holding a blocking lock. After a successful F_GETLK request, the value of *l_whence*shall be SEEK_SET.

448 Locks may start and extend beyond the current end of a file, but shall not start or 449 extend before the beginning of the file. A lock shall be set to extend to the largest 450 possible value of the file offset for that file if l_len is set to zero. If the *flock struct* 451 has l_whence and l_start that point to the beginning of the file, and l_len of zero, 452 the entire file shall be locked.

There shall be at most one type of lock set for each byte in the file. Before a suc-453 cessful return from an F_SETLK or an F_SETLKW request when the calling pro-454 cess has previously existing locks on bytes in the region specified by the request, 455 the previous lock type for each byte in the specified region shall be replaced by the 456 new lock type. As specified above under the descriptions of shared locks and 457 exclusive locks, an F_SETLK or an F_SETLKW request shall (respectively) fail or 458 block when another process has existing locks on bytes in the specified region and 459 the type of any of those locks conflicts with the type specified in the request. 460

461 All locks associated with a file for a given process shall be removed when a file 462 descriptor for that file is closed by that process or the process holding that file 463 descriptor terminates. Locks are not inherited by a child process created using 464 the *fork()* function.

A potential for deadlock occurs if a process controlling a locked region is put to sleep by attempting to lock the locked region of another process. If the system detects that sleeping until a locked region is unlocked would cause a deadlock, the fcntl() function shall fail with an [EDEADLK] error.

469 **6.5.2.3 Returns**

470 Upon successful completion, the value returned shall depend on *cmd*. The vari-471 ous return values are shown in Table 6-9.

Otherwise, a value of -1 shall be returned and *errno* shall be set to indicate the error.

'4 '5		Table 6-9 $- fcntl()$ Return Values
6	Request	Return Value
7 8	F_DUPFD F GETFD	A new file descriptor. Value of the flags defined in Table 6-2, but the return value shall not be negative.
))	F_SETFD F_GETFL	Value of file status flags and access modes, but the return value shall not be negative.
	F_SETFL F_GETLK F_SETLK	Value other than -1. Value other than -1. Value other than -1.
	f_setlkw	Value other than -1.

6.5.2.4 Errors 486

If any of the following conditions occur, the fcntl() function shall return -1 and set 487 *errno* to the corresponding value: 488

489	[EACCES] or	[EAGAIN]
490		The argument cmd is F_SETLK, the type of lock (l_type) is a
491		shared lock (F_RDLCK) or exclusive lock (F_WRLCK), and the
492		segment of a file to be locked is already exclusive-locked by
493		another process; or the type is an exclusive lock and some por-
494		tion of the segment of a file to be locked is already shared-
495		locked or exclusive-locked by another process.
496	[EBADF]	The <i>fildes</i> argument is not a valid file descriptor.
497		The argument <i>cmd</i> is F_SETLK or F_SETLKW, the type of lock
498		(l_type) is a shared lock (F_RDLCK), and fildes is not a valid file
499		descriptor open for reading.
500		The argument <i>cmd</i> is F_SETLK or F_SETLKW, the type of lock
501		(l_type) is an exclusive lock (F_WRLCK), and fildes is not a valid
502		file descriptor open for writing.
503	[EINTR]	The argument <i>cmd</i> is F_SETLKW, and the function was inter-
504		rupted by a signal.
505	[EINVAL]	The argument <i>cmd</i> is F_DUPFD, and the third argument is
506		negative or greater than or equal to {OPEN_MAX}.
507		The argument cmd is F_GETLK, F_SETLK, or F_SETLKW and
508		the data to which arg points is not valid, or <i>fildes</i> refers to a file
509		that does not support locking.
510	[EMFILE]	The argument <i>cmd</i> is F_DUPFD and {OPEN_MAX} file descrip-
511		tors are currently in use by this process, or no file descriptors
512		greater than or equal to arg are available.
513	[ENOLCK]	The argument <i>cmd</i> is F_SETLK or F_SETLKW, and satisfying
514		the lock or unlock request would result in the number of locked
515		regions in the system exceeding a system-imposed limit.

- For each of the following conditions, if the condition is detected, the fcntl() function shall return -1 and set *errno* to the corresponding value:
- 518 [EDEADLK] The argument *cmd* is F_SETLKW, and a deadlock condition was 519 detected.

520 6.5.2.5 Cross-References

521 close(), 6.3.1; exec, 3.1.2; open(), 5.3.1; <fcntl.h>, 6.5.1; 3.3.1.4.

522 6.5.3 Reposition Read/Write File Offset

523 Function: lseek()

524 **6.5.3.1 Synopsis**

- 525 #include <sys/types.h>
- 526 #include <unistd.h>
- 527 off_t lseek(int fildes, off_t offset, int whence);

528 **6.5.3.2 Description**

- The *fildes* argument is an open file descriptor. The *lseek()* function shall set the file offset for the open file description associated with *fildes* as follows:
- 531 (1) If whence is SEEK_SET, the offset is set to offset bytes.
- 532 (2) If whence is SEEK_CUR, the offset is set to its current value plus offset
 533 bytes.
- (3) If whence is SEEK_END, the offset is set to the size of the file plus offset
 bytes.
- The symbolic constants SEEK_SET, SEEK_CUR, and SEEK_END are defined in the header <unistd.h>.
- 538 Some devices are incapable of seeking. The value of the file offset associated with 539 such a device is undefined. The behavior of the *lseek()* function on such devices is 540 implementation defined.
- The *lseek()* function shall allow the file offset to be set beyond the end of existing data in the file. If data is later written at this point, subsequent reads of data in the gap shall return bytes with the value zero until data is actually written into the gap.
- 545 The *lseek()* function shall not, by itself, extend the size of a file.

546 **6.5.3.3 Returns**

547 Upon successful completion, the function shall return the resulting offset location 548 as measured in bytes from the beginning of the file. Otherwise, it shall return a 549 value of $((off_t) - 1)$, shall set *errno* to indicate the error, and the file offset shall 550 remain unchanged by this function call. ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

551 **6.5.3.4 Errors**

If any of the following conditions occur, the lseek() function shall return -1 and set *errno* to the corresponding value:

- 554 [EBADF] The *fildes* argument is not a valid file descriptor.
- 555 [EINVAL] The *whence* argument is not a proper value, or the resulting file 556 offset would be invalid.
- 557 [ESPIPE] The *fildes* argument is associated with a pipe or FIFO.

558 6.5.3.5 Cross-References

creat(), 5.3.2; dup(), 6.2.1; fcntl(), 6.5.2; open(), 5.3.1; read(), 6.4.1; sigaction(), 3.3.4; write(), 6.4.2; <unistd.h>, 2.9.

Section 7: Device- and Class-Specific Functions

1 7.1 General Terminal Interface

This section describes a general terminal interface that shall be provided. It shall be supported on any asynchronous communication ports if the implementation provides them. It is implementation defined whether this interface supports network connections or synchronous ports or both. The conformance document shall describe which device types are supported by these interfaces. Certain functions in this section apply only to the controlling terminal of a process; where this is the case, it is so noted.

9 7.1.1 Interface Characteristics

10 7.1.1.1 Opening a Terminal Device File

When a terminal file is opened, it normally causes the process to wait until a connection is established. In practice, application programs seldom open these files; they are opened by special programs and become the standard input, output, and error files of an application.

As described in 5.3.1, opening a terminal device file with the O_NONBLOCK flag
clear shall cause the process to block until the terminal device is ready and available. The CLOCAL flag can also affect open(). See 7.1.2.4.

18 7.1.1.2 Process Groups

A terminal may have a foreground process group associated with it. This foreground process group plays a special role in handling signal-generating input characters, as discussed below in 7.1.1.9.

If the implementation supports job control (if {_POSIX_JOB_CONTROL} is defined; 22 see 2.9), command interpreter processes supporting job control can allocate the 23 terminal to different jobs, or process groups, by placing related processes in a sin-24 25 gle process group and associating this process group with the terminal. The foreground process group of a terminal may be set or examined by a process, assum-26 ing the permission requirements in this section are met; see 7.2.3 and 7.2.4. The 27 terminal interface aids in this allocation by restricting access to the terminal by 28 processes that are not in the foreground process group; see 7.1.1.4. 29

When there is no longer any process whose process ID or process group ID matches the process group ID of the foreground process group, the terminal shall have no foreground process group. It is unspecified whether the terminal has a foreground process group when there is no longer any process whose process group ID matches the process group ID of the foreground process group, but there is a process whose process ID matches. No actions defined by this part of ISO/IEC 9945, other than allocation of a controlling terminal as described in 7.1.1.3 or a successful call to *tcsetpgrp()*, shall cause a process group to become the foreground process group of a terminal.

39 **7.1.1.3 The Controlling Terminal**

A terminal may belong to a process as its controlling terminal. Each process of a 40 session that has a controlling terminal has the same controlling terminal. A ter-41 minal may be the controlling terminal for at most one session. The controlling 42 terminal for a session is allocated by the session leader in an implementation-43 defined manner. If a session leader has no controlling terminal and opens a ter-44 minal device file that is not already associated with a session without using the 45 O NOCTTY option (see 5.3.1), it is implementation defined whether the terminal 46 becomes the controlling terminal of the session leader. If a process that is not a 47 session leader opens a terminal file, or the O_NOCTTY option is used on open(), 48 that terminal shall not become the controlling terminal of the calling process. 49 When a controlling terminal becomes associated with a session, its foreground 50 process group shall be set to the process group of the session leader. 51

The controlling terminal is inherited by a child process during a fork() function 52 call. A process relinquishes its controlling terminal when it creates a new session 53 with the setsid() function; other processes remaining in the old session that had 54 this terminal as their controlling terminal continue to have it. Upon the close of 55 the last file descriptor in the system (whether or not it is in the current session) 56 associated with the controlling terminal, it is unspecified whether all processes 57 that had that terminal as their controlling terminal cease to have any controlling 58 terminal. Whether and how a session leader can reacquire a controlling terminal 59 after the controlling terminal has been relinquished in this fashion is unspecified. 60 A process does not relinquish its controlling terminal simply by closing all of its 61 file descriptors associated with the controlling terminal if other processes con-62 tinue to have it open. 63

64 When a controlling process terminates, the controlling terminal is disassociated 65 from the current session, allowing it to be acquired by a new session leader. Sub-66 sequent access to the terminal by other processes in the earlier session may be 67 denied, with attempts to access the terminal treated as if modem disconnect had 68 been sensed.

69 **7.1.1.4 Terminal Access Control**

If a process is in the foreground process group of its controlling terminal, read 70 operations shall be allowed as described in 7.1.1.5. For those implementations 71 that support job control, any attempts by a process in a background process group 72 to read from its controlling terminal shall cause its process group to be sent a 73 SIGTTIN signal unless one of the following special cases apply: If the reading pro-74 cess is ignoring or blocking the SIGTTIN signal, or if the process group of the read-75 ing process is orphaned, the read() returns -1 with errno set to [EIO], and no sig-76 nal is sent. The default action of the SIGTTIN signal is to stop the process to 77 which it is sent. See 3.3.1.1. 78

If a process is in the foreground process group of its controlling terminal, write 79 operations shall be allowed as described in 7.1.1.8. Attempts by a process in a 80 background process group to write to its controlling terminal shall cause the pro-81 cess group to be sent a SIGTTOU signal unless one of the following special cases 82 apply: If TOSTOP is not set, or if TOSTOP is set and the process is ignoring or 83 blocking the SIGTTOU signal, the process is allowed to write to the terminal and 84 the SIGTTOU signal is not sent. If TOSTOP is set, and the process group of the 85 writing process is orphaned, and the writing process is not ignoring or blocking 86 SIGTTOU, the write() returns -1 with errno set to [EIO], and no signal is sent. 87

Certain calls that set terminal parameters are treated in the same fashion as
 write, except that TOSTOP is ignored; that is, the effect is identical to that of ter minal writes when TOSTOP is set. See 7.2.

91 7.1.1.5 Input Processing and Reading Data

A terminal device associated with a terminal device file may operate in fullduplex mode, so that data may arrive even while output is occurring. Each terminal device file has associated with it an *input queue*, into which incoming data is stored by the system before being read by a process. The system may impose a limit, {MAX_INPUT}, on the number of bytes that may be stored in the input queue. The behavior of the system when this limit is exceeded is implementation defined.

Two general kinds of input processing are available, determined by whether the terminal device file is in canonical mode or noncanonical mode. These modes are described in 7.1.1.6 and 7.1.1.7. Additionally, input characters are processed according to the c_iflag (see 7.1.2.2) and c_lflag (see 7.1.2.5) fields. Such processing can include *echoing*, which in general means transmitting input characters immediately back to the terminal when they are received from the terminal. This is useful for terminals that can operate in full-duplex mode.

The manner in which data is provided to a process reading from a terminal device file is dependent on whether the terminal device file is in canonical or noncanonical mode.

Another dependency is whether the O_NONBLOCK flag is set by open() or fcntl(). If the O_NONBLOCK flag is clear, then the read request shall be blocked until data is available or a signal has been received. If the O_NONBLOCK flag is set, then the read request shall be completed, without blocking, in one of three ways:

- (1) If there is enough data available to satisfy the entire request, the read()
 shall complete successfully and return the number of bytes read.
- (2) If there is not enough data available to satisfy the entire request, the
 read() shall complete successfully, having read as much data as possible,
 and return the number of bytes it was able to read.
- (3) If there is no data available, the read() shall return -1 with errno set to
 [EAGAIN].

When data is available depends on whether the input processing mode is canonical or noncanonical. The following subclauses, 7.1.1.6 and 7.1.1.7, describe each of these input processing modes.

123 7.1.1.6 Canonical Mode Input Processing

In canonical mode input processing, terminal input is processed in units of lines. 124 A line is delimited by a newline (' n') character, an end-of-file (EOF) character, or 125 an end-of-line (EOL) character. See 7.1.1.9 for more information on EOF and EOL. 126 This means that a read request shall not return until an entire line has been 127 typed or a signal has been received. Also, no matter how many bytes are 128 requested in the read call, at most one line shall be returned. It is not, however, 129 necessary to read a whole line at once; any number of bytes, even one, may be 130 requested in a read without losing information. 131

132 If {MAX_CANON} is defined for this terminal device, it is a limit on the number of 133 bytes in a line. The behavior of the system when this limit is exceeded is imple-134 mentation defined. If {MAX_CANON} is not defined, there is no such limit; 135 see 2.8.5.

Erase and kill processing occur when either of two special characters, the ERASE 136 and KILL characters (see 7.1.1.9), is received. This processing affects data in the 137 input queue that has not yet been delimited by a newline (NL), EOF, or EOL char-138 acter. This undelimited data makes up the current line. The ERASE character 139 deletes the last character in the current line, if there is any. The KILL character 140 deletes all data in the current line, if there is any. The ERASE and KILL charac-141 ters have no effect if there is no data in the current line. The ERASE and KILL 142 143 characters themselves are not placed in the input queue.

144 7.1.1.7 Noncanonical Mode Input Processing

In noncanonical mode input processing, input bytes are not assembled into lines, and erase and kill processing does not occur. The values of the MIN and TIME members of the c_cc array are used to determine how to process the bytes received.

MIN represents the minimum number of bytes that should be received when the *read()* function successfully returns. TIME is a timer of 0,1 second granularity that is used to time out short-term or bursty data transmissions. If MIN is greater than {MAX_INPUT}, the response to the request is undefined. The four possible values for MIN and TIME and their interactions are described below.

154 **7.1.1.7.1** Case A: MIN > 0, TIME > 0

In this case TIME serves as an interbyte timer and is activated after the first byte 155 is received. Since it is an interbyte timer, it is reset after a byte is received. The 156 interaction between MIN and TIME is as follows: as soon as one byte is received, 157 the interbyte timer is started. If MIN bytes are received before the interbyte timer 158 expires (remember that the timer is reset upon receipt of each byte), the read is 159 satisfied. If the timer expires before MIN bytes are received, the characters 160 received to that point are returned to the user. Note that if TIME expires, at least 161 one byte shall be returned because the timer would not have been enabled unless 162 a byte was received. In this case (MIN > 0, TIME > 0), the read shall block until 163 the MIN and TIME mechanisms are activated by the receipt of the first byte or 164 until a signal is received. If data is in the buffer at the time of the read(), the 165 result shall be as if data had been received immediately after the read(). 166

Part 1: SYSTEM API [C LANGUAGE]

167 **7.1.1.7.2** Case B: MIN > 0, TIME = 0

In this case, since the value of TIME is zero, the timer plays no role and only MIN
is significant. A pending read is not satisfied until MIN bytes are received (i.e.,
the pending read shall block until MIN bytes are received) or a signal is received.
A program that uses this case to read record-based terminal I/O may block
indefinitely in the read operation.

173 **7.1.1.7.3** Case C: MIN = 0, TIME > 0

In this case, since MIN = 0, TIME no longer represents an interbyte timer. It now 174 serves as a read timer that is activated as soon as the read() function is pro-175 cessed. A read is satisfied as soon as a single byte is received or the read timer 176 expires. Note that in this case if the timer expires, no bytes shall be returned. If 177 the timer does not expire, the only way the read can be satisfied is if a byte is 178 received. In this case, the read shall not block indefinitely waiting for a byte; if no 179 byte is received within TIME*0,1 seconds after the read is initiated, the read()180 shall return a value of zero, having read no data. If data is in the buffer at the 181 time of the read(), the timer shall be started as if data had been received immedi-182 ately after the read(). 183

184 **7.1.1.7.4** Case D: MIN = 0, TIME = 0

The minimum of either the number of bytes requested or the number of bytes currently available shall be returned without waiting for more bytes to be input. If no characters are available, *read()* shall return a value of zero, having read no data.

189 7.1.1.8 Writing Data and Output Processing

When a process writes one or more bytes to a terminal device file, they are processed according to the c_oflag field (see 7.1.2.3). The implementation may provide a buffering mechanism; as such, when a call to write() completes, all of the bytes written have been scheduled for transmission to the device, but the transmission will not necessarily have completed. See also 6.4.2 for the effects of O_NONBLOCK on write().

196 **7.1.1.9 Special Characters**

- Certain characters have special functions on input or output or both. These func tions are summarized as follows:
- 199INTRSpecial character on input and recognized if the ISIG flag (see2007.1.2.5) is enabled. It generates a SIGINT signal that is sent to201all processes in the foreground process group for which the ter-202minal is the controlling terminal. If ISIG is set, the INTR char-203acter is discarded when processed.
- 204QUITSpecial character on input and recognized if the ISIG flag is205enabled. It generates a SIGQUIT signal that is sent to all206processes in the foreground process group for which the termi-207nal is the controlling terminal. If ISIG is set, the QUIT charac-208ter is discarded when processed.

Special character on input and recognized if the ICANON flag is ERASE 209 set. It erases the last character in the current line; see 7.1.1.6. 210 The ERASE character shall not erase beyond the start of a line, 211 as delimited by an NL, EOF, or EOL character. If ICANON is 212 set, the ERASE character is discarded when processed. 213 **KILL** Special character on input and recognized if the ICANON flag is 214 set. It deletes the entire line, as delimited by a NL, EOF, or 215 EOL character. If ICANON is set, the KILL character is dis-216 carded when processed. 217 EOF Special character on input and recognized if the ICANON flag is 218 set. When received, all the bytes waiting to be read are 219 immediately passed to the process, without waiting for a new-220 line, and the EOF is discarded. Thus, if there are no bytes wait-221 ing (that is, the EOF occurred at the beginning of a line), a byte 222 count of zero shall be returned from the *read()*, representing an 223 end-of-file indication. If ICANON is set, the EOF character is 224 discarded when processed. 225 NL Special character on input and recognized if the ICANON flag is 226set. It is the line delimiter (' n'). 227 EOL Special character on input and recognized if the ICANON flag is 228set. It is an additional line delimiter, like NL. 229 SUSP Recognized on input if job control is supported (see 7.1.2.6). If 230 the ISIG flag is enabled, receipt of the SUSP character causes a 231 SIGTSTP signal to be sent to all processes in the foreground pro-232 cess group for which the terminal is the controlling terminal, 233 and the SUSP character is discarded when processed. 234 Special character on both input and output and recognized if 235 STOP the IXON (output control) or IXOFF (input control) flag is set. It 236 237 can be used to temporarily suspend output. It is useful with CRT terminals to prevent output from disappearing before it 238 239 can be read. If IXON is set, the STOP character is discarded when processed. 240 START Special character on both input and output and recognized if 241 the IXON (output control) or IXOFF (input control) flag is set. 242 Can be used to resume output that has been suspended by a 243 STOP character. If IXON is set, the START character is dis-244 carded when processed. 245CR Special character on input and recognized if the ICANON flag is 246 set; it is the 'r', as denoted in the C Standard {2}. When 247 ICANON and ICRNL are set and IGNCR is not set, this character 248 is translated into a NL and has the same effect as a NL 249 character. 250 The NL and CR characters cannot be changed. It is implementation defined 251

whether the START and STOP characters can be changed. The values for INTR, QUIT, ERASE, KILL, EOF, EOL, and SUSP (job control only), shall be changeable to suit individual tastes.

I

1

- If {_POSIX_VDISABLE} is in effect for the terminal file, special character functions
 associated with changeable special control characters can be disabled individu ally; see 7.1.2.6.
- If two or more special characters have the same value, the function performed when that character is received is undefined.

A special character is recognized not only by its value, but also by its context; for 260 example, an implementation may define multibyte sequences that have a mean-261 ing different from the meaning of the bytes when considered individually. Imple-262 mentations may also define additional single-byte functions. These 263 implementation-defined multibyte or single-byte functions are recognized only if 264 the IEXTEN flag is set; otherwise, data is received without interpretation, except 265as required to recognize the special characters defined in this subclause (7.1.1.9). 266

267 **7.1.1.10 Modem Disconnect**

If a modem disconnect is detected by the terminal interface for a controlling ter-268 minal, and if CLOCAL is not set in the c cflag field for the terminal (see 7.1.2.4), 269 the SIGHUP signal is sent to the controlling process associated with the terminal. 270 Unless other arrangements have been made, this causes the controlling process to 271 terminate; see 3.2.2. Any subsequent call to the read() function shall return the 272 value zero, indicating end of file. See 6.4.1. Thus, processes that read a terminal 273 file and test for end-of-file can terminate appropriately after a disconnect. If the 974 [EIO] condition specified in 6.4.1.4 that applies when the implementation supports 275 job control also exists, it is unspecified whether the EOF condition or the [EIO] is 276 returned. Any subsequent write() to the terminal device returns -1, with *errno* 277 set to [EIO], until the device is closed. 278

279 **7.1.1.11 Closing a Terminal Device File**

The last process to close a terminal device file shall cause any output to be sent to the device and any input to be discarded. Then, if HUPCL is set in the control modes and the communications port supports a disconnect function, the terminal device shall perform a disconnect.

284 7.1.2 Parameters That Can Be Set

285 7.1.2.1 *termios* Structure

Routines that need to control certain terminal I/O characteristics shall do so by
using the *termios* structure as defined in the header <termios.h>. The
members of this structure include (but are not limited to) those shown in Table 71.

The types $tcflag_t$ and cc_t shall be defined in the header <termios.h>. They shall be unsigned integral types.

292

293	181	016 1-1	- termios	s Structure
294295 296	Member Type	Array Size	Member Name	Description
297	tcflag_t		c_iflag	Input modes.
298	tcflag_t		c_oflag	Output modes.
299	tcflag_t		c_cflag	Control modes.
300	tcflag_t		c_lflag	Local modes.
301	cc_t	NCCS	c_cc	Control characters.
302				

$m_{\rm e}$ h l = $m_{\rm e}$ 1 . . .

7.1.2.2 Input Modes 303

Values of the c_{iflag} field, shown in Table 7-2, describe the basic terminal input 304 control and are composed of the bitwise inclusive OR of the masks shown, which 305 shall be bitwise distinct. The mask name symbols in this table are defined in 306 <termios.h>. 307

308 309	Table 7	-2 – termios c_iflag Field	
310	Mask Name	Description	
311	BRKINT	Signal interrupt on break.	
312	ICRNL	Map CR to NL on input.	
313	IGNBRK	Ignore break condition.	
314	IGNCR	Ignore CR.	
315	IGNPAR	Ignore characters with parity errors.	
316	INLCR	Map NL to CR on input.	
317	INPCK	Enable input parity check.	
318	ISTRIP	Strip character.	
319	IXOFF	Enable start/stop input control.	
320	IXON	Enable start/stop output control.	
321	PARMRK	Mark parity errors.	
322			

In the context of asynchronous serial data transmission, a break condition is 323 defined as a sequence of zero-valued bits that continues for more than the time to 324 send one byte. The entire sequence of zero-valued bits is interpreted as a single 325 break condition, even if it continues for a time equivalent to more than one byte. 326 In contexts other than asynchronous serial data transmission, the definition of a 327 break condition is implementation defined. 328

If IGNBRK is set, a break condition detected on input is ignored, that is, not put 329 on the input queue and therefore not read by any process. If IGNBRK is not set 330 and BRKINT is set, the break condition shall flush the input and output queues. 331 If the terminal is the controlling terminal of a foreground process group, the 332 break condition shall generate a single SIGINT signal to that foreground process 333 group. If neither IGNBRK nor BRKINT is set, a break condition is read as a single 334 '\0', or if PARMRK is set, as '377', '0', '0'. 335

336 If IGNPAR is set, a byte with a framing or parity error (other than break) is ignored. 337

If PARMRK is set and IGNPAR is not set, a byte with a framing or parity error (other than break) is given to the application as the three-character sequence 340 (377', '0', X, where '377', '0' is a two-character flag preceding each sequence and X is the data of the character received in error. To avoid ambiguity in this case, if ISTRIP is not set, a valid character of '377' is given to the application as '377', '377'. If neither PARMRK nor IGNPAR is set, a framing or parity error (other than break) is given to the application as a single character '0'.

If INPCK is set, input parity checking is enabled. If INPCK is not set, input parity checking is disabled, allowing output parity generation without input parity errors. Note that whether input parity checking is enabled or disabled is independent of whether parity detection is enabled or disabled (see 7.1.2.4). If parity detection is enabled, but input parity checking is disabled, the hardware to which the terminal is connected shall recognize the parity bit, but the terminal special file shall not check whether this bit is set correctly or not.

- If ISTRIP is set, valid input bytes are first stripped to seven bits; otherwise, all eight bits are processed.
- If INLCR is set, a received NL character is translated into a CR character. If
 IGNCR is set, a received CR character is ignored (not read). If IGNCR is not set
 and ICRNL is set, a received CR character is translated into a NL character.
- If IXON is set, start/stop output control is enabled. A received STOP character shall suspend output, and a received START character shall restart output. When IXON is set, START and STOP characters are not read, but merely perform flow control functions. When IXON is not set, the START and STOP characters are read.
- If IXOFF is set, start/stop input control is enabled. The system shall transmit one 362 or more STOP characters, which are intended to cause the terminal device to stop 363 transmitting data, as needed to prevent the input queue from overflowing and 364 causing the undefined behavior described in 7.1.1.5 and shall transmit one or 365 more START characters, which are intended to cause the terminal device to 366 resume transmitting data, as soon as the device can continue transmitting data 367 without risk of overflowing the input queue. The precise conditions under which 368 STOP and START characters are transmitted are implementation defined. 369
- 370 The initial input control value after *open()* is implementation defined.

371 **7.1.2.3 Output Modes**

Values of the c_oflag field describe the basic terminal output control and are composed of the bitwise inclusive OR of the following masks, which shall be bitwise distinct:

375 376

5	<u>Mask Name</u>	Description
6	OPOST	Perform output processing.

377 The mask name symbols for the *c_oflag* field are defined in <termios.h>.

378 If OPOST is set, output data is processed in an implementation-defined fashion so 379 that lines of text are modified to appear appropriately on the terminal device;

- otherwise, characters are transmitted without change.
- ³⁸¹ The initial output control value after *open()* is implementation defined.

382 **7.1.2.4 Control Modes**

388

Values of the c_cflag field, shown in Table 7-3, describe the basic terminal hardware control and are composed of the bitwise inclusive OR of the masks shown, which shall be bitwise distinct; not all values specified are required to be supported by the underlying hardware. The mask name symbols in this table are defined in <termios.h>.

389			
390	Mask Name	Description	
391	CLOCAL	Ignore modem status lines.	
392	CREAD	Enable receiver.	
393	CSIZE	Number of bits per byte:	
394	CS5	5 bits	
395	CS6	6 bits	
396	CS7	7 bits	
397	CS8	8 bits	
398	CSTOPB	Send two stop bits, else one.	
399	HUPCL	Hang up on last close.	
400	PARENB	Parity enable.	
401	PARODD	Odd parity, else even.	
402			

Table 7	7-3		termios	С_	cflag	Field
---------	-----	--	---------	----	-------	-------

The CSIZE bits specify the byte size in bits for both transmission and reception. This size does not include the parity bit, if any. If CSTOPB is set, two stop bits are used; otherwise, one stop bit is used. For example, at 110 baud, two stop bits are normally used.

- ⁴⁰⁷ If CREAD is set, the receiver is enabled; otherwise, no characters shall be ⁴⁰⁸ received.
- If PARENB is set, parity generation and detection is enabled and a parity bit is
 added to each character. If parity is enabled, PARODD specifies odd parity if set;
 otherwise, even parity is used.

If HUPCL is set, the modem control lines for the port shall be lowered when the
last process with the port open closes the port or the process terminates. The
modem connection shall be broken.

- If CLOCAL is set, a connection does not depend on the state of the modem status lines. If CLOCAL is clear, the modem status lines shall be monitored.
- Under normal circumstances, a call to the open() function shall wait for the
 modem connection to complete. However, if the O_NONBLOCK flag is set (see
 5.3.1) or if CLOCAL has been set, the open() function shall return immediately
 without waiting for the connection.
- If the object for which the control modes are set is not an asynchronous serial connection, some of the modes may be ignored; for example, if an attempt is made to

set the baud rate on a network connection to a terminal on another host, the baud
rate may or may not be set on the connection between that terminal and the
machine to which it is directly connected.

426 The initial hardware control value after *open()* is implementation defined.

427 **7.1.2.5 Local Modes**

432

Values of the *c_lflag* field, shown in Table 7-4, describe the control of various functions and are composed of the bitwise inclusive OR of the masks shown, which
shall be bitwise distinct. The mask name symbols in this table are defined in
<termios.h>.

434	Mask Name	Description
435	ECHO	Enable echo.
436	ECHOE	Echo ERASE as an error-correcting backspace.
437	ECHOK	Echo KILL.
438	ECHONL	Echo '\n'.
439	ICANON	Canonical input (erase and kill processing).
440	IEXTEN	Enable extended (implementation-defined) functions.
441	ISIG	Enable signals.
442	NOFLSH	Disable flush after interrupt, quit, or suspend.
143	TOSTOP	Send SIGTTOU for background output.

Table 7-4 – termios c_lflag Field

If ECHO is set, input characters are echoed back to the terminal. If ECHO is not set, input characters are not echoed.

If ECHOE and ICANON are set, the ERASE character shall cause the terminal to
erase the last character in the current line from the display, if possible. If there is
no character to erase, an implementation may echo an indication that this was
the case or do nothing.

- If ECHOK and ICANON are set, the KILL character shall either cause the terminal
 to erase the line from the display or shall echo the '\n' character after the KILL
 character.
- If ECHONL and ICANON are set, the '\n' character shall be echoed even if ECHO
 is not set.
- If ICANON is set, canonical processing is enabled. This enables the erase and kill
 edit functions and the assembly of input characters into lines delimited by NL,
 EOF, and EOL, as described in 7.1.1.6.
- If ICANON is not set, read requests are satisfied directly from the input queue. A
 read shall not be satisfied until at least MIN bytes have been received or the
 timeout value TIME has expired between bytes. The time value represents tenths
 of seconds. See 7.1.1.7 for more details.

If ISIG is set, each input character is checked against the special control characters
 ters INTR, QUIT, and SUSP (job control only). If an input character matches one of
 these control characters, the function associated with that character is performed.

If ISIG is not set, no checking is done. Thus, these special input functions are pos-466 sible only if ISIG is set. 467

If IEXTEN is set, implementation-defined functions shall be recognized from the 468 input data. It is implementation defined how IEXTEN being set interacts with 469 ICANON, ISIG, IXON, or IXOFF. If IEXTEN is not set, then implementation-defined 470 functions shall not be recognized, and the corresponding input characters shall be 471 processed as described for ICANON, ISIG, IXON, and IXOFF. 472

- If NOFLSH is set, the normal flush of the input and output queues associated with 473 the INTR, QUIT, and SUSP (job control only) characters shall not be done. 474
- If TOSTOP is set and the implementation supports job control, the signal SIGTTOU 475 is sent to the process group of a process that tries to write to its controlling termi-476 nal if it is not in the foreground process group for that terminal. This signal, by 477 default, stops the members of the process group. Otherwise, the output generated 478 by that process is output to the current output stream. Processes that are block-479 ing or ignoring SIGTTOU signals are excepted and allowed to produce output, and 480 the SIGTTOU signal is not sent. 481
- The initial local control value after open() is implementation defined. 482

7.1.2.6 Special Control Characters 483

The special control characters values are defined by the array c cc. The subscript 484 name and description for each element in both canonical and noncanonical modes 485 are shown in Table 7-5. The subscript name symbols in this table are defined in 486 <termios.h>. 487

Subsc	ript Usage	
Canonical Mode	Noncanonical Mode	Description
VEOF		EOF character
VEOL		EOL character
VERASE		ERASE character
VINTR	VINTR	INTR character
VKILL		KILL character
	VMIN	MIN value
VQUIT	VQUIT	QUIT character
VSUSP	VSUSP	SUSP character
	VTIME	TIME value
VSTART	VSTART	START character
VSTOP	VSTOP	STOP character

Table	7.5	_	termine c	cc	Special	Control	Characters
Table	1-0	-	<i>lernuos</i> c	CC	Special	CONTROL	Unaracters

The subscript values shall be unique, except that the VMIN and VTIME subscripts 505 may have the same values as the VEOF and VEOL subscripts, respectively. 506

Implementations that do not support job control may ignore the SUSP character 507 value in the *c_cc* array indexed by the VSUSP subscript. 508

- The value of NCCS (the number of elements in the c_{cc} array) is unspecified by this part of ISO/IEC 9945.
- Implementations that do not support changing the START and STOP characters may ignore the character values in the c_{cc} array indexed by the VSTART and VSTOP subscripts when tcsetattr() is called, but shall return the value in use when tcgetattr() is called.

If {_POSIX_VDISABLE} is defined for the terminal device file, and the value of one of the changeable special control characters (see 7.1.1.9) is {_POSIX_VDISABLE}, that function shall be disabled, that is, no input data shall be recognized as the disabled special character. If ICANON is not set, the value of {_POSIX_VDISABLE} has no special meaning for the VMIN and VTIME entries of the c_cc array.

520 The initial values of all control characters are implementation defined.

521 7.1.2.7 Baud Rate Values

524

The baud rate values specified in Table 7-6 can be set into the *termios* structure by the baud rate functions in 7.1.3.

6	Name	Description	Name	Description
7	B0	Hang up	B600	600 baud
	B50	50 baud	B1200	1200 baud
	B75	75 baud	B1800	1800 baud
	B110	110 baud	B2400	2400 baud
	B134	134.5 baud	B4800	4800 baud
	B150	150 baud	B9600	9600 baud
	B200	200 baud	B19200	19200 baud
	B300	300 baud	B38400	38 400 baud

Table 7-6 - termios Baud Rate Values

536 7.1.3 Baud Rate Functions

537 Functions: cfgetispeed(), cfgetospeed(), cfsetispeed(), cfsetospeed()

538 **7.1.3.1 Synopsis**

- 539 #include <termios.h>
- 540 speed_t cfgetospeed(const struct termios *termios_p);
- 541 int cfsetospeed(struct termios *termios_p, speed_t speed);
- 542 speed_t cfgetispeed(const struct termios *termios_p);
- 543 int cfsetispeed(struct termios *termios_p, speed_t speed);

544 **7.1.3.2 Description**

The following interfaces are provided for getting and setting the values of the input and output baud rates in the *termios* structure. The effects on the terminal device described below do not become effective until the *tcsetattr()* function is successfully called, and not all errors are detected until *tcsetattr()* is called as well.

The input and output baud rates are represented in the *termios* structure. The values shown in Table 7-6 are defined. The name symbols in this table are defined in <termios.h>.

- The type $speed_t$ shall be defined in <termios.h> and shall be an unsigned integral type.
- 554 The *termios_p* argument is a pointer to a *termios* structure.
- The *cfgetospeed()* function shall return the output baud rate stored in the *termios* structure to which *termios_p* points.
- The *cfgetispeed()* function shall return the input baud rate stored in the *termios* structure to which *termios_p* points.
- The *cfsetospeed()* function shall set the output baud rate stored in the *termios* structure to which *termios_p* points.
- The *cfsetispeed()* function shall set the input baud rate stored in the *termios* structure to which *termios_p* points.
- 563 Certain values for speeds that are set in the *termios* structure and passed to 564 *tcsetattr()* have special meanings. These are discussed under *tcsetattr()*.
- The cfgetispeed() and cfgetospeed() functions return exactly the value found in the termios data structure, without interpretation.
- 567 Both *cfsetispeed()* and *cfsetospeed()* return a value of zero if successful and -1 to 568 indicate an error. It is unspecified whether these return an error if an unsup-569 ported baud rate is set.
- 570 **7.1.3.3 Returns**
- 571 See 7.1.3.2.

572 **7.1.3.4 Errors**

This part of ISO/IEC 9945 does not specify any error conditions that are required to be detected for the *cfgetispeed()*, *cfgetospeed()*, *cfsetispeed()*, or *cfsetospeed()* functions. Some errors may be detected under conditions that are unspecified by this part of ISO/IEC 9945.

- 577 **7.1.3.5 Cross-References**
- 578 *tcsetattr()*, 7.2.1.

ſ

Part 1: SYSTEM API [C LANGUAGE]

1

579 7.2 General Terminal Interface Control Functions

The functions that are used to control the general terminal function are described in this clause. If the implementation supports job control, unless otherwise noted for a specific command, these functions are restricted from use by background processes. Attempts to perform these operations shall cause the process group to be sent a SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process is allowed to perform the operation and the SIGTTOU signal is not sent.

In all the functions, *fildes* is an open file descriptor. However, the functions affect the underlying terminal file, not just the open file description associated with the file descriptor.

- 590 7.2.1 Get and Set State
- 591 Functions: tcgetattr(), tcsetattr()
- 592 **7.2.1.1 Synopsis**
- 593 **#include** <termios.h>

594 int tcgetattr(int fildes, struct termios *termios_p);

595 int tcsetattr(int fildes, int optional_actions, 596 const struct termios * termios_p);

597 **7.2.1.2 Description**

The tcgetattr() function shall get the parameters associated with the object 598 referred to by fildes and store them in the termios structure referenced by 599 *termios_p*. This function is allowed from a background process; however, the ter-600 minal attributes may be subsequently changed by a foreground process. If the 601 terminal device supports different input and output baud rates, the baud rates 602 stored in the *termios* structure returned by *tcgetattr()* shall reflect the actual baud 603 rates, even if they are equal. If differing baud rates are not supported, the rate 604 returned as the output baud rate shall be the actual baud rate. The rate returned 605 as the input baud rate shall be either the number zero or the output rate (as one 606 of the symbolic values). Permitting either behavior is obsolescent.⁴⁾ 607

The *tcsetattr*() function shall set the parameters associated with the terminal (unless support is required from the underlying hardware that is not available) from the *termios* structure referenced by *termios_p* as follows:

611

(1) If optional_actions is TCSANOW, the change shall occur immediately.

 ⁴⁾ In a future revision of this part of ISO/IEC 9945, a returned value of zero as the input baud rate
 when differing baud rates are not supported may no longer be permitted.

- (2) If optional_actions is TCSADRAIN, the change shall occur after all output
 written to *fildes* has been transmitted. This function should be used
 when changing parameters that affect output.
- 617
- 618 619

(3) If *optional_actions* is TCSAFLUSH, the change shall occur after all output written to the object referred to by *fildes* has been transmitted, and all input that has been received, but not read, shall be discarded before the change is made.

The symbolic constants for the values of *optional_actions* are defined in <termios.h>.

The zero baud rate, B0, is used to terminate the connection. If B0 is specified as the output baud rate when *tcsetattr()* is called, the modem control lines shall no longer be asserted. Normally, this will disconnect the line.

If the input baud rate is equal to the numeral zero in the *termios* structure when *tcsetattr()* is called, the input baud rate will be changed by *tcsetattr()* to the same value as that specified by the value of the output baud rate, exactly as if the input rate had been set to the output rate by *cfsetispeed()*. This usage of zero is obsolescent.

631 The *tcsetattr*() function shall return success if it was able to perform any of the requested actions, even if some of the requested actions could not be performed. 632 633 It shall set all the attributes that the implementation does support as requested and leave all the attributes not supported by the hardware unchanged. If no part 634 of the request can be honored, it shall return -1 and set *errno* to [EINVAL]. If the 635 input and output baud rates differ and are a combination that is not supported, 636 neither baud rate is changed. A subsequent call to tcgetattr() shall return the 637 actual state of the terminal device [reflecting both the changes made and not 638 made in the previous tcsetattr() call]. The tcsetattr() function shall not change the 639 values in the *termios* structure whether or not it actually accepts them. 640

The termios structure may have additional fields not defined by this part of ISO/IEC 9945. The effect of the tcsetattr() function is undefined if the value of the termios structure pointed to by termios_p was not derived from the result of a call to tcgetattr() on fildes; a Strictly Conforming POSIX.1 Application shall modify only fields and flags defined by this part of ISO/IEC 9945 between the call to tcgetattr() and tcsetattr(), leaving all other fields and flags unmodified.

No actions defined by this part of ISO/IEC 9945, other than a call to *tcsetattr*() or a close of the last file descriptor in the system associated with this terminal device, shall cause any of the terminal attributes defined by this part of ISO/IEC 9945 to change.

651 **7.2.1.3 Returns**

⁶⁵² Upon successful completion, a value of zero is returned. Otherwise, a value of -1⁶⁵³ is returned and *errno* is set to indicate the error.

ł

654 **7.2.1.4 Errors**

- If any of the following conditions occur, the tcgetattr() function shall return -1and set *errno* to the corresponding value:
- 657 [EBADF] The *fildes* argument is not a valid file descriptor.
- 658 [ENOTTY] The file associated with *fildes* is not a terminal.
- If any of the following conditions occur, the tcsetattr() function shall return -1 and set *errno* to the corresponding value:
- 661 [EBADF] The *fildes* argument is not a valid file descriptor.
- 662 [EINTR] A signal interrupted the *tcsetattr()* function.
- 663[EINVAL]The optional_actions argument is not a proper value, or an
attempt was made to change an attribute represented in the
termios structure to an unsupported value.
- 666 [ENOTTY] The file associated with *fildes* is not a terminal.

667 7.2.1.5 Cross-References

- 668 <termios.h>, 7.1.2.
- 669 7.2.2 Line Control Functions
- **Functions:** *tcsendbreak()*, *tcdrain()*, *tcflush()*, *tcflow()*

671 **7.2.2.1 Synopsis**

- 672 #include <termios.h>
- 673 int tcsendbreak(int fildes, int duration);
- 674 int tcdrain(int fildes);
- 675 int tcflush(int fildes, int queue_selector);
- 676 int tcflow(int fildes, int action);

677 **7.2.2.2 Description**

If the terminal is using asynchronous serial data transmission, the *tcsendbreak()* function shall cause transmission of a continuous stream of zero-valued bits for a specific duration. If *duration* is zero, it shall cause transmission of zero-valued bits for at least 0,25 seconds and not more that 0,5 seconds. If *duration* is not zero, it shall send zero-valued bits for an implementation-defined period of time.

If the terminal is not using asynchronous serial data transmission, it is imple mentation defined whether the *tcsendbreak()* function sends data to generate a
 break condition (as defined by the implementation) or returns without taking any
 action.

The tcdrain() function shall wait until all output written to the object referred to
by fildes has been transmitted.

Upon successful completion, the *tcflush()* function shall have discarded any data
 written to the object referred to by *fildes* but not transmitted, or data received, but
 not read, depending on the value of *queue_selector*:

- (1) If *queue_selector* is TCIFLUSH, it shall flush data received, but not read.
- (2) If queue_selector is TCOFLUSH, it shall flush data written, but not
 transmitted.
- (3) If queue_selector is TCIOFLUSH, it shall flush both data received but not read and data written but not transmitted.

The *tcflow()* function shall suspend transmission or reception of data on the object referred to by *fildes*, depending on the value of *action*:

- (1) If action is TCOOFF, it shall suspend output.
- 700 (2) If action is TCOON, it shall restart suspended output.
- (3) If action is TCIOFF, the system shall transmit a STOP character, which is intended to cause the terminal device to stop transmitting data to the system. (See the description of IXOFF in 7.1.2.2.)
- (4) If action is TCION, the system shall transmit a START character, which is intended to cause the terminal device to start transmitting data to the system. (See the description of IXOFF in 7.1.2.2.)
- The symbolic constants for the values of *queue_selector* and *action* are defined in <termios.h>.
- The default on the opening of a terminal file is that neither its input nor its output is suspended.

711 **7.2.2.3 Returns**

⁷¹² Upon successful completion, a value of zero is returned. Otherwise, a value of -1⁷¹³ is returned and *errno* is set to indicate the error.

714 7.2.2.4 Errors

If any of the following conditions occur, the tcsendbreak() function shall return -1and set *errno* to the corresponding value:

- 717 [EBADF] The *fildes* argument is not a valid file descriptor.
- 718 [ENOTTY] The file associated with *fildes* is not a terminal.

If any of the following conditions occur, the tcdrain() function shall return -1 and set *errno* to the corresponding value:

- 721 [EBADF] The *fildes* argument is not a valid file descriptor.
- 722 [EINTR] A signal interrupted the *tcdrain()* function.
- 723 [ENOTTY] The file associated with *fildes* is not a terminal.

If any of the following conditions occur, the tcflush() function shall return -1 and set *errno* to the corresponding value:

- 726 [EBADF] The *fildes* argument is not a valid file descriptor.
- 727 [EINVAL] The *queue_selector* argument is not a proper value.
- 728 [ENOTTY] The file associated with *fildes* is not a terminal.
- If any of the following conditions occur, the tcflow() function shall return -1 and set *errno* to the corresponding value:
- 731 [EBADF] The *fildes* argument is not a valid file descriptor.
- 732 [EINVAL] The action argument is not a proper value.
- 733 [ENOTTY] The file associated with *fildes* is not a terminal.
- 734 **7.2.2.5 Cross-References**
- 735 <termios.h>, 7.1.2.

736 7.2.3 Get Foreground Process Group ID

- **Function:** *tcgetpgrp()*
- 738 **7.2.3.1 Synopsis**
- 739 #include <sys/types.h>
- 740 pid_t tcgetpgrp(int fildes);

741 **7.2.3.2 Description**

742 If {_POSIX_JOB_CONTROL} is defined:

- (1) The tcgetpgrp() function shall return the value of the process group ID of
 the foreground process group associated with the terminal.
- (2) The tcgetpgrp() function is allowed from a process that is a member of a background process group; however, the information may be subsequently changed by a process that is a member of a foreground process group.
- 749 Otherwise:
- The implementation shall either support the tcgetpgrp() function as described above or the tcgetpgrp() call shall fail.

752 **7.2.3.3 Returns**

Upon successful completion, tcgetpgrp() returns the process group ID of the foreground process group associated with the terminal. If there is no foreground process group, tcgetpgrp() shall return a value greater than 1 that does not match the process group ID of any existing process group. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

758 **7.2.3.4 Errors**

If any of the following conditions occur, the tcgetpgrp() function shall return -1and set *errno* to the corresponding value:

- 761 [EBADF] The *fildes* argument is not a valid file descriptor.
- [ENOSYS] The tcgetpgrp() function is not supported in this implementa tion.
- [ENOTTY] The calling process does not have a controlling terminal, or the
 file is not the controlling terminal.
- 766 **7.2.3.5 Cross-References**
- 767 setsid(), 4.3.2; setpgid(), 4.3.3; tcsetpgrp(), 7.2.4.

768 7.2.4 Set Foreground Process Group ID

- 769 Function: tcsetpgrp()
- 770 **7.2.4.1 Synopsis**
- 771 #include <sys/types.h>
- int tcsetpgrp(int fildes, pid_t pgrp_id);

773 **7.2.4.2 Description**

174 If {_POSIX_JOB_CONTROL} is defined:

If the process has a controlling terminal, the *tcsetpgrp()* function shall set the foreground process group ID associated with the terminal to *pgrp_id*. The file associated with *fildes* must be the controlling terminal of the calling process, and the controlling terminal must be currently associated with the session of the calling process. The value of *pgrp_id* must match a process group ID of a process in the same session as the calling process.

- 781 Otherwise:
- The implementation shall either support the *tcsetpgrp()* function as described above, or the *tcsetpgrp()* call shall fail.
- 784 **7.2.4.3 Returns**
- Upon successful completion, tcsetpgrp() returns a value of zero. Otherwise, a value of -1 is returned and errno is set to indicate the error.

787 **7.2.4.4 Errors**

If any of the following conditions occur, the tcsetpgrp() function shall return -1and set *errno* to the corresponding value:

790	[EBADF]	The <i>fildes</i> argument is not a valid file descriptor.
791 792	[EINVAL]	The value of the <i>pgrp_id</i> argument is not supported by the implementation.
793 794	[ENOSYS]	The <i>tcsetpgrp()</i> function is not supported in this implementation.
795 796 797	[ENOTTY]	The calling process does not have a controlling terminal, or the file is not the controlling terminal, or the controlling terminal is no longer associated with the session of the calling process.
798 799 800	[EPERM]	The value of <i>pgrp_id</i> is a value supported by the implementation, but does not match the process group ID of a process in the same session as the calling process.



L

Section 8: Language-Specific Services for the C Programming Language

1 8.1 Referenced C Language Routines

The functions listed below are described in the indicated sections of the 2 C Standard {2}. POSIX.1 with the C Language Binding comprises these functions, 3 the extensions to them described in this clause, and the rest of the requirements 4 stipulated in this part of ISO/IEC 9945. The functions appended with plus signs 5 (+) have requirements beyond those set forth in the C Standard {2}. Any imple-6 mentation claiming conformance to POSIX.1 with the C Language Binding shall 7 comply with the requirements outlined in this clause, the requirements stipulated 8 in the rest of this part of ISO/IEC 9945, and the requirements in the indicated sec-9 tions of the C Standard {2}. 10

For requirements concerning conformance to this clause, see 1.3.3 and its subclauses.

13 14	4.2	Diagnostics Functions: assert.
15 16 17	4.3	Character Handling Functions: isalnum, isalpha, iscntrl, isdigit, isgraph, islower, isprint, ispunct, isspace, isupper, isxdigit, tolower, toupper.
18 19	4.4	Localization Functions: setlocale+.
20 21 22	4.5	Mathematics Functions: acos, asin, atan, atan2, cos, sin, tan, cosh, sinh, tanh, exp, frexp, ldexp, log, log10, modf, pow, sqrt, ceil, fabs, floor, fmod.
23 24	4.6	Non-Local Jumps Functions: setjmp, longjmp.
25 26 27 28 29	4.9	Input/Output Functions: clearerr, fclose, feof, ferror, fflush, fgetc, fgets, fopen, fputc, fputs, fread, freopen, fseek, ftell, fwrite, getc, getchar, gets, perror, printf, fprintf, sprintf, putc, putchar, puts, remove, rename+, rewind, scanf, fscanf, sscanf, setbuf, tmpfile, tmpnam, ungetc.
30 31 32	4.10	General Utilities Functions: abs, atof, atoi, atol, rand, srand, calloc, free, malloc, realloc, abort+, exit, getenv+, bsearch, qsort.
33 34 35	4.11	String Handling Functions: strcpy, strncpy, strcat, strncat, strcmp, strncmp, strchr, strcspn, strpbrk, strrchr, strspn, strstr, strtok, strlen.

ł

ł

I

- 36 4.12 Date and Time
- Functions: time, asctime, ctime+, gmtime+, localtime+, mktime+, strftime+.
- Systems conforming to this part of ISO/IEC 9945 shall make no distinction between the "text streams" and the "binary streams" described in the C Standard {2}.
- For the fseek() function, if the specified position is beyond end-of-file, the consequences described in lseek() (see 6.5.3) shall occur.
- The EXIT_SUCCESS macro, as used by the *exit()* function, shall evaluate to a value of zero. Similarly, the EXIT_FAILURE macro shall evaluate to a nonzero value.
- The relationship between a time in seconds since the Epoch used as an argument to gmtime() and the tm structure (defined in <time.h>) is that the result shall be as specified in the expression given in the definition of seconds since the Epoch in 2.2.2.77, where the names in the structure and in the expression correspond. If the time zone UCTO is in effect, this shall also be true for localtime() and mktime().
- 53 8.1.1 Extensions to Time Functions

The contents of the environment variable named **TZ** (see 2.6) shall be used by the functions *ctime()*, *localtime()*, *strftime()*, and *mktime()* to override the default time zone. The value of **TZ** has one of the two forms (spaces inserted for clarity):

- 57 :characters
- 58 or:
- 59 std offset dst offset, rule
- If **TZ** is of the first format (i.e., if the first character is a colon), the characters following the colon are handled in an implementation-defined manner.
- The expanded format (for all TZs whose value does not have a colon as the first character) is as follows:
- 64 stdoffset[dst[offset][,start[/time],end[/time]]]
- 65 Where:
- std and dst Indicates no less than three, nor more than {TZNAME_MAX}, 66 bytes that are the designation for the standard (*std*) or summer 67 (dst) time zone. Only std is required; if dst is missing, then sum-68 mer time does not apply in this locale. Upper- and lowercase 69 70 letters are explicitly allowed. Any characters except a leading colon (:) or digits, the comma (,), the minus (-), the plus (+), and 71 the null character are permitted to appear in these fields, but 72their meaning is unspecified. 73

offset 74

75

76

77

78

79

80

81

82

83

84

85

86

87

90

91

92

93

94

95

96

97

98

99

100

101

102

103

Indicates the value one must add to the local time to arrive at Coordinated Universal Time. The offset has the form:

hh[:mm[:ss]]

The minutes (mm) and seconds (ss) are optional. The hour (hh)shall be required and may be a single digit. The offset following std shall be required. If no offset follows dst, summer time is assumed to be one hour ahead of standard time. One or more digits may be used; the value is always interpreted as a decimal number. The hour shall be between zero and 24, and the minutes (and seconds)-if present-between zero and 59. Use of values outside these ranges causes undefined behavior. If preceded by a "-", the time zone shall be east of the Prime Meridian; otherwise it shall be west (which may be indicated by an optional preceding "+").

rule Indicates when to change to and back from summer time. The 88 *rule* has the form: 89

date/time, date/time

where the first *date* describes when the change from standard to summer time occurs and the second date describes when the change back happens. Each time field describes when, in current local time, the change to the other time is made.

- The format of *date* shall be one of the following:
- The Julian day $n \ (1 \le n \le 365)$. Leap days shall not be Jn counted. That is, in all years-including leap years-February 28 is day 59 and March 1 is day 60. It is impossible to explicitly refer to the occasional February 29.
 - The zero-based Julian day ($0 \le n \le 365$). Leap days n shall be counted, and it is possible to refer to February 29.

Mm.n.d

- 104 The d^{th} day $(0 \le d \le 6)$ of week *n* of month *m* of the 105 year $(1 \le n \le 5, 1 \le m \le 12)$, where week 5 means "the 106 last d day in month m" which may occur in either the 107 fourth or the fifth week). Week 1 is the first week in 108 which the d'th day occurs. Day zero is Sunday. 109
- 110 The *time* has the same format as *offset* except that no leading sign ("-" or "+") shall be allowed. The default, if time is not 111 given, shall be 02:00:00. 112

Whenever ctime(), strftime(), mktime(), or localtime() is called, the time zone 113 names contained in the external variable *tzname* shall be set as if the *tzset()* func-114 tion had been called. 115

Applications are explicitly allowed to change TZ and have the changed TZ apply 116 to themselves. 117

I

118 8.1.2 Extensions to setlocale() Function

- 119 Function: setlocale()
- 120 **8.1.2.1 Synopsis**
- 121 #include <locale.h>
- 122 char *setlocale(int category, const char *locale);
- 123 **8.1.2.2 Description**

The *setlocale()* function sets, changes, or queries the locale of the process according to the values of the *category* and the *locale* arguments. The possible values for *category* include:

- 127 LC_CTYPE 128 LC_COLLATE
- 129 LC_TIME
- 130 LC NUMERIC
- 131 LC_MONETARY
- 132 Implementation-defined additional categories
- For POSIX.1 systems, environment variables are defined that correspond to the named categories above and that have the same spelling.
- The value LC_ALL for *category* names all of the categories of the locale of the process; LC_ALL is a special constant, not a category. There is an environment variable LC_ALL with the semantics noted below.
- The *locale* argument is a pointer to a character string that can be an explicit string, a NULL pointer, or a null string.
- When *locale* is an explicit string, the contents of the string are implementation 140 defined except for the value "C". The value "C" for locale specifies the minimal 141 environment for C-language translation. If setlocale() is not invoked, the "C" 142 locale shall be the locale of the process. The locale name "POSIX" shall be recog-143 144 nized. It shall provide the same semantics as the C locale for those functions defined within this part of ISO/IEC 9945 or by the C Standard (2). Extensions or 145 146 refinements to the POSIX locale beyond those provided by the C locale may be included in future revisions, and other parts of ISO/IEC 9945 are expected to add 147 148 to the requirements of the POSIX locale.
- When *locale* is a **NULL** pointer the locale of the process is queried according to the value of *category*. The content of the string returned is unspecified.
- When *locale* is a null string, the *setlocale()* function takes the name of the new locale for the specified category from the environment as determined by the first condition met below:
- (1) If LC_ALL is defined in the environment and is not null, the value of
 LC_ALL is used.
- (2) If there is a variable defined in the environment with the same name as
 the category and that is not null, the value specified by that environment
 variable is used.

Part 1: SYSTEM API [C LANGUAGE]

- 159 160
- (3) If LANG is defined in the environment and is not null, the value of LANG is used.

If the resulting value is a supported locale, *setlocale()* sets the specified category 161 of the locale of the process to that value and returns the value specified below. If 162 the value does not name a supported locale (and is not null), setlocale() returns a 163 **NULL** pointer, and the locale of the process is not changed by this function call. If 164 no nonnull environment variable is present to supply a value, it is implementa-165 tion defined whether *setlocale()* sets the specified category of the locale of the pro-166 cess to a systemwide default value or to "C" or to "POSIX". The possible actual 167 values of the environment variables are implementation defined and should 168 appear in the system documentation. 169

Setting all of the categories of the locale of the process is similar to successively setting each individual category of the locale of the process, except that all error checking is done before any actions are performed. To set all the categories of the locale of the process, *setlocale()* is invoked as:

In this case, *setlocale()* first verifies that the values of all the environment vari-175 ables it needs according to the precedence above indicate supported locales. If the 176 value of any of these environment-variable searches yields a locale that is not sup-177 ported (and nonnull), the setlocale() function returns a NULL pointer and the 178 locale of the process is not changed. If all environment variables name supported 179 locales, *setlocale()* then proceeds as if it had been called for each category, using 180 the appropriate value from the associated environment variable or from the 181 implementation-defined default if there is no such value. 182

183 8.1.2.3 Returns

A successful call to *setlocale()* returns a string that corresponds to the locale set. The string returned is such that "a subsequent call with that string and its associated category will restore that part of the process's locale" (C Standard {2}). The string returned shall not be modified by the process, but may be overwritten by a subsequent call to the *setlocale()* function. This string is not required to be the value of the environment variable used, if one was used.

190 8.2 C Language Input/Output Functions

191 This clause describes input/output functions of the C Standard {2} and their 192 interactions with other functions defined by this part of ISO/IEC 9945.

All functions specified in the C Standard {2} as operating on a *file name* shall
operate on a *pathname*. All functions specified in the C Standard {2} as creating a
file shall do so as if they called the *creat()* function with a value appropriate to the
C language function for the *path* argument and a value of

- 197 S_IRUSR|S_IWUSR|S_IRGRP|S_IWGRP|S_IROTH|S_IWOTH
- 198 for the *mode* argument.

I

I

I

I

The type *FILE* and the terms *file position indicator* and *stream* are those defined by the C Standard {2}.

A stream is considered local to a single process. After a fork() call, each of the parent and child have distinct streams that share an open file description.

203 8.2.1 Map a Stream Pointer to a File Descriptor

204 Function: fileno()

205 8.2.1.1 Synopsis

- 206 #include <stdio.h>
- 207 int fileno(FILE *stream);
- 208 8.2.1.2 Description

The *fileno()* function returns the integer file descriptor associated with the *stream* (see 5.3.1).

The following symbolic values in the <unistd.h> header (see 2.9) define the file descriptors that shall be associated with the C language *stdin*, *stdout*, and *stderr* when the application is started:

214	Name	Description	<u>Value</u>
215	STDIN_FILENO	Standard input value, stdin.	0
216	STDOUT_FILENO	Standard output value, stdout.	1
217	STDERR_FILENO	Standard error value, stderr.	2

At entry to main(), these streams shall be in the same state as if they had just been opened with *fdopen()* called with a mode consistent with that required by the C Standard {2} and the file descriptor described above.

221 8.2.1.3 Returns

See 8.2.1.2. If an error occurs, a value of -1 is returned and *errno* is set to indicate the error.

224 8.2.1.4 Errors

This part of ISO/IEC 9945 does not specify any error conditions that are required to be detected for the *fileno()* function. Some errors may be detected under conditions that are unspecified by this part of ISO/IEC 9945.

- 228 8.2.1.5 Cross-References
- 229 open(), 5.3.1.

230 8.2.2 Open a Stream on a File Descriptor

- **Function:** *fdopen()*
- 232 8.2.2.1 Synopsis
- 233 #include <stdio.h>
- 234 FILE *fdopen(int fildes, const char *type);
- 235 8.2.2.2 Description
- 236 The *fdopen()* routine associates a stream with a file descriptor.
- 237 The *type* argument is a character string having one of the following values:

238	"r"	Open for reading.
239	"W	Open for writing.
240	"a"	Open for writing at end-of-file.
241	"r+"	Open for update (reading and writing).
242	"w+"	Open for update (reading and writing).
243	"a+"	Open for update (reading and writing) at end-of-file.

The meaning of these flags is exactly as specified by the C Standard {2} for fopen(), except that "w" and "w+" do not cause truncation of the file. Additional values for the type argument may be defined by an implementation.

- The application shall ensure that the mode of the stream is allowed by the mode of the open file.
- The file position indicator associated with the new stream is set to the position indicated by the file offset associated with the file descriptor. The error indicator and end-of-file indicator for the stream shall be cleared.
- 252

253 8.2.2.3 Returns

If successful, the *fdopen()* function returns a pointer to a stream. Otherwise, a
NULL pointer is returned and *errno* is set to indicate the error.

256 8.2.2.4 Errors

This part of ISO/IEC 9945 does not specify any error conditions that are required to be detected for the *fdopen()* function. Some errors may be detected under conditions that are unspecified by this part of ISO/IEC 9945.

- 260 8.2.2.5 Cross-References
- 261 open(), 5.3.1; fopen() [C Standard {2}].

262 8.2.3 Interactions of Other FILE-Type C Functions

A single open file description can be accessed both through streams and through file descriptors. Either a file descriptor or a stream will be called a *handle* on the open file description to which it refers; an open file description may have several handles.

Handles can be created or destroyed by user action without affecting the underlying open file description. Some of the ways to create them include fcntl(), dup(),fdopen(), fileno(), and fork() (which duplicates existing ones into new processes). They can be destroyed by at least fclose(), close(), and the *exec* functions (which close some file descriptors and destroy streams).

- A file descriptor that is never used in an operation that could affect the file offset [for example read(), write(), or lseek()] is not considered a handle in this discussion, but could give rise to one [as a consequence of fdopen(), dup(), or fork(), for example]. This exception does include the file descriptor underlying a stream, whether created with fopen() or fdopen(), as long as it is not used directly by the application to affect the file offset. [The read() and write() functions implicitly affect the file offset; lseek() explicitly affects it.]
- The result of function calls involving any one handle (the *active handle*) are defined elsewhere in this part of ISO/IEC 9945, but if two or more handles are used, and any one of them is a stream, their actions shall be coordinated as described below. If this is not done, the result is undefined.
- A handle that is a stream is considered to be closed when either an *fclose()* or *freopen()* is executed on it [the result of *freopen()* is a new stream for this discussion, which cannot be a handle on the same open file description as its previous value] or when the process owning that stream terminates with *exit()* or *abort()*. A file descriptor is closed by *close()*, *_exit()*, or by one of the *exec* functions when FD_CLOEXEC is set on that file descriptor.
- For a handle to become the active handle, the actions below must be performed 289 between the last other use of the first handle (the current active handle) and the 290 first other use of the second handle (the future active handle). The second handle 291 then becomes the active handle. All activity by the application affecting the file 292 offset on the first handle shall be suspended until it again becomes the active han-293 dle. (If a stream function has as an underlying function that affects the file offset, 294 the stream function will be considered to affect the file offset. The underlying 295 functions are described below.) 296
- The handles need not be in the same process for these rules to apply. Note that after a fork(), two handles exist where one existed before. The application shall assure that, if both handles will ever be accessed, that they will both be in a state where the other could become the active handle first. The application shall prepare for a fork() exactly as if it were a change of active handle. [If the only action performed by one of the processes is one of the *exec* functions or *_exit()* {not *exit()*}, the handle is never accessed in that process.]
- 304 (1) For the first handle, the first applicable condition below shall apply.
 305 After the actions required below are taken, the handle may be closed if it
 306 is still open.
I

(a) If it is a file descriptor, no action is required. 307 (b) If the only further action to be performed on any handle to this open 308 file description is to close it, no action need be taken. 309 If it is a stream that is unbuffered, no action need be taken. (c) 310 (d) If it is a stream that is line-buffered and the last character written 311 to the stream was a newline [that is, as if a $putc(' \n')$ was the 312 most recent operation on that stream], no action need be taken. 313 (e) If it is a stream that is open for writing or append (but not also open 314 for reading), either an fflush() shall occur or the stream shall be 315 316 closed. If the stream is open for reading and it is at the end of the file (f) 317 [feof() is true], no action need be taken. 318 If the stream is open with a mode that allows reading and the 319 (g) underlying open file description refers to a device that is capable of 320 seeking, either an *fflush()* shall occur or the stream shall be closed. 321 (h) Otherwise, the result is undefined. 322 For the second handle: if any previous active handle has called a func-323 (2) tion that explicitly changed the file offset, except as required above for 324 the first handle, the application shall perform an *lseek()* or an *fseek()* (as 325 appropriate to the type of the handle) to an appropriate location. 326 (3) If the active handle ceases to be accessible before the requirements on the 327 first handle above have been met, the state of the open file description 328 becomes undefined. This might occur, for example, during a fork() or an 329 $_{exit}().$ 330 (4) The exec functions shall be considered to make inaccessible all streams 331 that are open at the time they are called, independent of what streams or 332 file descriptors may be available to the new process image. 333 (5) Implementations shall assure that an application, even one consisting of 334 several processes, shall yield correct results (no data is lost or duplicated 335 when writing, all data is written in order, except as requested by seeks) 336 when the rules above are followed, regardless of the sequence of handles 337 used. If the rules above are not followed, the result is unspecified. When 338 these rules are followed, it is implementation defined whether, and under 339 what conditions, all input is seen exactly once. 340 (6) Each function that operates on a stream is said to have zero or more 341 342 underlying functions. This means that the stream function shares certain traits with the underlying functions, but does not require that there 343 be any relation between the implementations of the stream function and 344 its underlying functions. 345 (7) Also, in the subclauses below, additional requirements on the standard 346 347 I/O routines, beyond those in the C Standard {2}, are given.

1

L

1

- 348 **8.2.3.1** fopen()
- The *fopen()* function shall allocate a file descriptor as *open()* does.
- 350 The underlying function is open().
- 351 8.2.3.2 fclose()

The fclose() function shall perform a close() on the file descriptor that is associated with the *FILE* stream. It shall also mark for update the *st_ctime* and *st_mtime* fields of the underlying file, if the stream was writable, and if buffered data had not been written to the file yet.

356 The underlying functions are *write()* and *close()*.

357

- 358 **8.2.3.3** freopen()
- 359 The *freopen()* function has the properties of both *fclose()* and *fopen()*.
- 360 8.2.3.4 fflush()

The fflush() function shall mark for update the st_ctime and st_mtime fields of the underlying file if the stream was writable and if buffered data had not been written to the file yet.

364 The underlying functions are *write()* and *lseek()*.

365

366 8.2.3.5 fgetc(), fgets(), fread(), getc(), getchar(), gets(), scanf(), fscanf()

These functions may mark the st_atime field for update. The st_atime field shall be marked for update by the first successful execution of one of these functions that returns data not supplied by a prior call to ungetc().

The underlying functions are *read()* and *lseek()*.

8.2.3.6 fputc(), fputs(), fwrite(), putc(), putchar(), puts(), printf(), fprintf()

The *st_ctime* and *st_mtime* fields of the file shall be marked for update between the successful execution of one of these functions and the next successful completion of a call to either *fflush()* or *fclose()* on the same stream or a call to *exit()* or *abort()*.

377 The underlying functions are *write()* and *lseek()*.

If *fwrite()* writes greater than zero bytes, but fewer than requested, the error indicator for the stream shall be set. If the underlying *write()* reports an error, *errno*shall not be modified by *fwrite()*, and the error indicator for the stream shall be
set.

If the implementation provides the *vprintf()* and *vfprintf()* functions from the C
Standard {2}, they also shall meet the constraints specified in this part of
ISO/IEC 9945 for (respectively) *printf()* and *fprintf()*.

385 **8.2.3.7** *fseek(), rewind()*

These functions shall mark the *st_ctime* and *st_mtime* fields of the file for update if the stream was writable and if buffered data had not yet been written to the file.

389 The underlying functions are *lseek()* and *write()*.

If the most recent operation, other than ftell(), on a given stream is fflush(), the file offset in the underlying open file description shall be adjusted to reflect the location specified by the fseek().

393 8.2.3.8 perror()

The perror() function shall mark the file associated with the standard error | stream as having been written (st_ctime, st_mtime marked for update) at some time between its successful completion and exit(), abort(), or the completion of fflush() or fclose() on stderr.

- 398 8.2.3.9 *tmpfile()*
- 399 The *tmpfile()* function shall allocate a file descriptor as *fopen()* does.

400 **8.2.3.10** *ftell(*)

The underlying function is *lseek()*. The result of *ftell()* after an *fflush()* shall be the same as the result before the *fflush()*. If the stream is opened in append mode or if the O_APPEND flag is set as a consequence of dealing with other handles on the file, the result of *ftell()* on that stream is unspecified.

405 8.2.3.11 Error Reporting

If any of the functions above return an error indication, the value of *errno* shall be set to indicate the error condition. If that error condition is one that this part of ISO/IEC 9945 specifies to be detected by one of the corresponding underlying functions, the value of *errno* shall be the same as the value specified for the underlying function.

411 **8.2.3.12** *exit(), abort()*

The exit() function shall have the effect of fclose() on every open stream, with the properties of fclose() as described above. The abort() function shall also have these effects if the call to abort() causes process termination, but shall have no effect on streams otherwise. The C Standard {2} specifies the conditions where abort() does or does not cause process termination. For the purposes of that specification, a signal that is blocked shall not be considered caught.

I

1

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

- 418 8.2.4 Operations on Files the *remove*() Function
- 419 The *remove()* function shall have the same effect on file times as *unlink()*.
- 420 8.3 Other C Language Functions
- 421 8.3.1 Nonlocal Jumps
- 422 Functions: *sigsetjmp()*, *siglongjmp()*

423 **8.3.1.1 Synopsis**

- 424 #include <setjmp.h>
- 425 int sigsetjmp(sigjmp_buf env, int savemask);
- 426 void siglongjmp(sigjmp_buf env, int val);

427 **8.3.1.2 Description**

The sigsetjmp() macro shall comply with the definition of the setjmp() macro in the C Standard {2}. If the value of the savemask argument is not zero, the sigsetjmp() function shall also save the current signal mask of the process (see 3.3.1) as part of the calling environment.

The siglongjmp() function shall comply with the definition of the longjmp() function in the C Standard {2}. If and only if the *env* argument was initialized by a call to the sigsetjmp() function with a nonzero savemask argument, the siglongjmp() function shall restore the saved signal mask.

436 8.3.1.3 Cross-References

- 437 sigaction(), 3.3.4; <signal.h>, 3.3.1; sigprocmask(), 3.3.5; sigsuspend(), 3.3.7.
- 438 **8.3.2 Set Time Zone**
- 439 Function: *tzset()*
- 440 8.3.2.1 Synopsis
- 441 #include <time.h>
- 442 void tzset(void);

443 8.3.2.2 Description

The *tzset()* function uses the value of the environment variable TZ to set time
conversion information used by *localtime()*, *ctime()*, *strftime()*, and *mktime()*. If
TZ is absent from the environment, implementation-defined default time-zone
information shall be used.

- 448 The *tzset()* function shall set the external variable *tzname*:
- 449 extern char *tzname[2] = {"std", "dst"};
- where *std* and *dst* are as described in 8.1.1.



Section 9: System Databases

1 9.1 System Databases

The routines described in this section allow an application to access the two system databases that are described below.

- 4 The *group* database contains the following information for each group:
 - (1) Group name

5

6

7

9

- (2) Numerical group ID
 - (3) List of all users allowed in the group
- 8 The *user* database contains the following information for each user:
 - (1) User name
- 10 (2) Numerical user ID
- 11 (3) Numerical group ID
- 12 (4) Initial working directory
 - (5) Initial user program
- 14 If the initial user program field is null, the system default is used.
- 15 If the initial working directory field is null, the interpretation of that field is 16 implementation defined.
- These databases may contain other fields that are unspecified by this part of ISO/IEC 9945.

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

1

19 9.2 Database Access

- 20 9.2.1 Group Database Access
- Functions: getgrgid(), getgrnam()

22 **9.2.1.1 Synopsis**

- 23 #include <sys/types.h>
- 24 #include <grp.h>
- 25 struct group *getgrgid(gid_t gid);
- 26 struct group *getgrnam(const char *name);

27 9.2.1.2 Description

The getgrgid() and getgrnam() routines both return pointers to an object of type struct group containing an entry from the group database with a matching gid or name. This structure, which is defined in <grp.h>, includes the members shown in Table 9-1.

Member Type	Member Name	Description	
char *	gr_name	The name of the group.	
gid_t	gr_gid	The numerical group ID.	
char **	gr_mem	A null-terminated vector of pointers to the individual member n	

40 **9.2.1.3 Returns**

- 41 A NULL pointer is returned on error or if the requested entry is not found.
- The return values may point to static data that is overwritten by each call.

43 **9.2.1.4 Errors**

This part of ISO/IEC 9945 does not specify any error conditions that are required to be detected for the *getgrgid()* or *getgrnam()* functions. Some errors may be detected under conditions that are unspecified by this part of ISO/IEC 9945.

47 9.2.1.5 Cross-References

48 getlogin(), 4.2.4.

49 9.2.2 User Database Access

50 Functions: *getpwuid()*, *getpwnam()*

51 **9.2.2.1 Synopsis**

52 53	<pre>#include <sys types.h=""> #include <pwd.h></pwd.h></sys></pre>	
54	<pre>struct passwd *getpwuid(uid_t uid);</pre>	I
55	<pre>struct passwd *getpwnam(const char *name);</pre>	1

56 **9.2.2.2 Description**

The getpwuid() and getpwnam() functions both return a pointer to an object of type struct passwd containing an entry from the user database with a matching uid or name. This structure, which is defined in <pwd.h>, includes the members shown in Table 9-2.

Member Type	Member Name	Description
char *	pw_name	User name.
uid_t	pw_uid	User ID number.
gid_t	pw_gid	Group ID number.
char *	pw_dir	Initial Working Directory
char *	pw_shell	Initial User Program.

Table 9-2 - passwd Structure

71

61

72 9.2.2.3 Returns

- A NULL pointer is returned on error or if the requested entry is not found.
- 74 The return values may point to static data that is overwritten on each call.

75 9.2.2.4 Errors

This part of ISO/IEC 9945 does not specify any error conditions that are required to be detected for the *getpwuid()* or *getpwnam()* functions. Some errors may be

detected under conditions that are unspecified by this part of ISO/IEC 9945.

- 79 9.2.2.5 Cross-References
- 80 *getlogin()*, 4.2.4.

1



Section 10: Data Interchange Format

1 10.1 Archive/Interchange File Format

A conforming system shall provide a mechanism to copy files from a medium to the file hierarchy and copy files from the file hierarchy to a medium using the interchange formats described here. This part of ISO/IEC 9945 does not define this mechanism.

When this mechanism is used to copy files from the medium by a process without 6 appropriate privileges, the protection information (ownership and access permis-7 sions) shall be set in the same fashion that creat() would when given the mode 8 argument matching the file permissions supplied by the mode field of the 9 extended tar format or the c_{mode} field of the extended cpio format. A process 10 with appropriate privileges shall restore the ownership and the permissions 11 exactly as recorded on the medium, except that the symbolic user and group IDs 12 are used for the tar format, as described in 10.1.1. 13

The *format-creating utility* is used to translate from the file system to the formats defined in this clause. The *format-reading utility* is used to translate from the formats defined in this clause to a file system. The interface to these utilities, including their name or names, is implementation defined.

The headers of these formats are defined to use characters represented in ISO/IEC 646 {1}; however, no restrictions are placed on the contents of the files themselves. The data in a file may be binary data or text represented in any format available to the user. When these formats are used to transfer text at the source level, all characters shall be represented in ISO/IEC 646 {1} International Reference Version (IRV).

The media format and the frames on the media in which the data appear are unspecified by this part of ISO/IEC 9945.

26 NOTE: Guidelines are given in Annex B.

27 10.1.1 Extended tar Format

An extended tar archive tape or file contains a series of blocks. Each block is a fixed-size block of 512 bytes (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 is represented by a header block that describes the file, followed by zero or more blocks that give the contents of the file. At the end of the archive file are two blocks filled with binary zeroes, interpreted as an end-of-archive indicator.

I

I

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

42

The blocks may be grouped for physical I/O operations. Each group of n blocks (where n is set by the application utility creating the archive file) may be written with a single write() operation. On magnetic tape, the result of this write is a single tape record. The last group of blocks is always at the full size, so blocks after the two zero blocks contain undefined data.

The header block is structured as shown in Table 10-1. All lengths and offsets are in decimal.

44	Field Name	Byte Offset	Length (in bytes)
45	name	0	100
46	mode	100	8
47	uid	108	8
18	gid	116	8
49	size	124	12
50	mtime	136	12
51	chksum	148	8
52	typeflag	156	1
53	linkname	157	100
54	magic	257	6
55	version	263	2
56	uname	265	32
57	gname	297	32
58	devmajor	329	8
59	devminor	337	8
60	prefix	345	155
31			

Table 10-1 - tar Header Block

Symbolic constants used in the header block are defined in the header <tar.h>
 as follows:

```
#define TMAGIC
                       "ustar" /* ustar and a null */
64
      #define TMAGLEN
65
                       6
      #define TVERSION "00"
                                /* 00 and no null */
66
67
      #define TVERSLEN 2
68
      /* Values used in typeflag field */
      #define REGTYPE '0'
                                /* Regular file
69
                                                   */
70
      #define AREGTYPE '\0'
                                /* Regular file
                                                   */
      #define LNKTYPE
                       111
                                /* Link
                                                   */
71
      #define SYMTYPE
                       121
                                                   */
72
                                /* Reserved
                       131
      #define CHRTYPE
                                /* Character special */
73
                       141
74
      #define BLKTYPE
                                /* Block special */
75
      #define DIRTYPE
                       151
                                /* Directory
                                                   */
      #define FIFOTYPE '6'
                                                  */
76
                                /* FIFO special
      #define CONTTYPE '7'
                                /* Reserved
                                                   */
77
      /* Bits used in the mode field - values in octal */
78
79
      #define TSUID
                       04000
                              /* Set UID on execution */
      #define TSGID
                       02000
                                /* Set GID on execution */
80
      #define TSVTX
81
                       01000
                                /* Reserved */
82
                                /* File permissions */
```

83	#define	TUREAD	00 400	/*	Read by owner */
84	#define	TUWRITE	00200	/*	Write by owner */
85	#define	TUEXEC	00100	/*	Execute/Search by owner */
86	#define	TGREAD	00040	/*	Read by group */
87	#define	TGWRITE	00020	/*	Write by group */
88	#define	TGEXEC	00010	/*	Execute/Search by group */
89	#define	TOREAD	00004	/*	Read by other */
90	#define	TOWRITE	00002	/*	Write by other */
91	#define	TOEXEC	00001	/*	Execute/Search by other */

All characters are represented in the coded character set of ISO/IEC 646 {1}. For 92 maximum portability between implementations, names should be selected from 93 characters represented by the portable filename character set as 8-bit characters 94 with most significant bit zero. If an implementation supports the use of charac-95 ters outside the portable filename character set in names for files, users, and 96 groups, one or more implementation-defined encodings of these characters shall 97 be provided for interchange purposes. However, the format-reading utility shall 98 never create file names on the local system that cannot be accessed via the func-99 tions described previously in this part of ISO/IEC 9945; see 5.3.1, 5.6.2, 5.2.1, 100 6.5.2, and 5.1.2. If a file name is found on the medium that would create an 101invalid file name, the implementation shall define if the data from the file is 102 stored on the file hierarchy and under what name it is stored. A format-reading 103 utility may choose to ignore these files as long as it produces an error indicating 104 that the file is being ignored. 105

Each field within the header block is contiguous; that is, there is no padding used.
 Each character on the archive medium is stored contiguously.

The fields magic, uname, and gname are null-terminated character strings. The fields name, linkname, and prefix are null-terminated character strings except when all characters in the array contain nonnull characters including the last character. The version field is two bytes containing the characters "00" (zerozero). The typeflag contains a single character. All other fields are leading zerofilled octal numbers using digits from ISO/IEC 646 {1} IRV. Each numeric field is terminated by one or more space or null characters.

The name and the prefix fields produce the pathname of the file. The hierarchical 115 relationship of the file is retained by specifying the pathname as a path prefix, 116 and a slash character and filename as the suffix. A new pathname is formed, if 117 prefix is not an empty string (its first character is not null), by concatenating 118 prefix (up to the first null character), a slash character, and name; otherwise, 119 name is used alone. In either case, name is terminated at the first null character. 120 If *prefix* is an empty string, it is simply ignored. In this manner, pathnames of at 121 most 256 characters can be supported. If a pathname does not fit in the space 122provided, the format-creating utility shall notify the user of the error, and no 123 attempt shall be made by the format-creating utility to store any part of the file-124header or data-on the medium. 125

The *linkname* field, described below, does 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, the format-creating utility shall notify the user of the error, and the utility shall not attempt to store the link on the medium.

1

The mode field provides 9 bits specifying file permissions and 3 bits to specify the 130 set UID, set GID, and TSVTX modes. Values for these bits were defined previously. 131 When appropriate privilege is required to set one of these mode bits, and the user 132 restoring the files from the archive does not have the appropriate privilege, the 133 mode bits for which the user does not have appropriate privilege shall be ignored. 134 Some of the mode bits in the archive format are not mentioned elsewhere in this 135 part of ISO/IEC 9945. If the implementation does not support those bits, they may 136 be ignored. 137

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 bytes. If the *typeflag* field is set to specify a 140 file to be of type LNKTYPE or SYMTYPE, the size field shall be specified as zero. If 141 the typeflag field is set to specify a file of type DIRTYPE, the size field is inter-142 preted as described under the definition of that record type. No data blocks are 143 stored for LNKTYPE, SYMTYPE, or DIRTYPE. If the *typeflag* field is set to 144 CHRTYPE, BLKTYPE, or FIFOTYPE, the meaning of the *size* field is unspecified by 145 this part of ISO/IEC 9945, and no data blocks are stored on the medium. Addition-146 ally, for FIFOTYPE, the size field shall be ignored when reading. If the typeflag 147 field is set to any other value, the number of blocks written following the header 148 is (size+511)/512, ignoring any fraction in the result of the division. 149

The *mtime* field is the modification time of the file at the time it was archived. It is the ISO/IEC 646 {1} representation of the octal value of the modification time | obtained from the *stat*() function.

The *chksum* field is the ISO/IEC 646 {1} IRV representation of the octal value of the simple sum of all bytes in the header block. Each 8-bit byte in the header is treated as an unsigned value. These values are added to an unsigned integer, initialized to zero, the precision of which shall be no less than 17 bits. When calculating the checksum, the *chksum* field is treated as if it were all blanks.

The typeflag field specifies the type of file archived. If a particular implementa-158 tion does not recognize the type, or the user does not have appropriate privilege to 159 create that type, the file shall be extracted as if it were a regular file if the file 160 type is defined to have a meaning for the size field that could cause data blocks to 161 be written on the medium (see the previous description for size). If conversion to 162 an ordinary file occurs, the format-reading utility shall produce an error indicat-163 ing that the conversion took place. All of the typeflag fields are coded in 164 ISO/IEC 646 {1} IRV: 165

- 166'0'Represents a regular file. For backward compatibility, a typeflag167value of binary zero ('\0') should be recognized as meaning a regu-168lar file when extracting files from the archive. Archives written169with this version of the archive file format shall create regular files170with a typeflag value of ISO/IEC 646 {1} IRV '0'.
- 171'1'Represents a file linked to another file, of any type, previously172archived. Such files are identified by each file having the same dev-173ice and file serial number. The linked-to name is specified in the174linkname field with a null terminator if it is less than 100 bytes in175length.

1

1

I

I

1

- 176' 2'Reserved to represent a link to another file, of any type, whose dev-177ice or file serial number differs. This is provided for systems that178support linked files whose device or file serial numbers differ, and179should be treated as a type ' 1' file if this extension does not exist.
- '3','4' 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 part of ISO/IEC 9945. Implementations may map the device specifications to their own local specification or may ignore the entry.
- 186 '5' Specifies a directory or subdirectory. On systems where disk alloca187 tion is performed on a directory basis, the *size* field shall contain the
 188 maximum number of bytes (which may be rounded to the nearest
 189 disk block allocation unit) that the directory may hold. A *size* field
 190 of zero indicates no such limiting. Systems that do not support lim191 iting in this manner should ignore the *size* field.
- 192'6'Specifies a FIFO special file. Note that the archiving of a FIFO file193archives the existence of this file and not its contents.
- 194'7'Reserved to represent a file to which an implementation has associ-195ated some high performance attribute. Implementations without196such extensions should treat this file as a regular file (type '0').
- 197 'A'-'Z' The letters A through Z are reserved for custom implementations. |
 198 All other values are reserved for specification in future revisions of
 199 this part of ISO/IEC 9945.
- The magic field is the specification that this archive was output in this archive format. If this field contains TMAGIC, the *uname* and *gname* fields shall contain the ISO/IEC 646 {1} 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 password and group files 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.
- 207 The encoding of the header is designed to be portable across machines.
- 208 10.1.1.1 Cross-References

209 <grp.h>, 9.2.1; <pwd.h>, 9.2.2; <sys/stat.h>, 5.6.1; stat(), 5.6.2; 210 <unistd.h>, 2.9.

211 10.1.2 Extended cpio Format

The byte-oriented cpio archive format is a series of entries, each comprised of 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 bytes. This blocking shall be used only to make physical I/O more efficient. The last group of blocks is always at the full size.

For the byte-oriented cpio archive format, the individual entry information must be in the order indicated and described by Table 10-2.

222		Header	
223	Field Name	Length (in bytes)	Interpreted as
224	c_magic	6	Octal number
225	c_dev	6	Octal number
226	c_ino	6	Octal number
227	c_mode	6	Octal number
228	c_uid	6	Octal number
229	c_gid	6	Octal number
230	c_nlink	6	Octal number
231	c_rdev	6	Octal number
232	c_mtime	11	Octal number
233	c_namesize	6	Octal number
234	c_filesize	11	Octal number
235		File Name	
236	Field Name	Length	Interpreted as
237	c_name	c_namesize	Pathname string
238		File Data	
239	Field Name	Length	Interpreted as
240	c_filedata	c_filesize	Data
241			

Table 10-2 – Byte-Oriented cpio Archive Entry

242 10.1.2.1 cpio Header

For each file in the archive, a header as defined previously shall be written. The information in the header fields shall be written as streams of ISO/IEC 646 {1} characters interpreted as octal numbers. The octal numbers are extended to the necessary length by appending ISO/IEC 646 {1} IRV zeros at the most-significantdigit end of the number; the result is written to the stream of bytes mostsignificant-digit first. The fields shall be interpreted as follows:

- (1) *c_magic* shall identify the archive as being a transportable archive by containing the magic bytes as defined by MAGIC (070707).
- (2) c_dev and c_ino shall contain values that uniquely identify the file within
 the archive (i.e., no files shall 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.
- (3) c_{mode} shall contain the file type and access permissions as defined in Table 10-3.
- c_uid shall contain the user ID of the owner.
- c_{gid} shall contain the group ID of the group.

l

9 1		Table 10-6	\rightarrow values in	or cpio c_mode rield		
1			File Perm	issions		
2		Name	Value	Indicates		
3		C_IRUSR	000 400	Read by owner.		
ł		C_IWUSR	000 200	Write by owner.		
i i		C_IXUSR	000 100	Execute by owner.		
		C_IRGRP	000 040	Read by group.		
		C_IWGRP	000 020	Write by group.		
		C_IXGRP	000010	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	004000	Set <i>uid</i> .		
		C_ISGID	002 000	Set gid.		
		C_ISVTX	001000	Reserved.		
			File T	ype		
		Name	Value	Indicates		
		C_ISDIR	040 000	Directory.		
		C_ISFIFO	010000	FIFO.		
		C_ISREG	0100 000	Regular file.		
		C_ISBLK	060 000	Block special file.		
		C_ISCHR	020 000	Character special file.		
		C_ISCTG	0110 000	Reserved.		
		C_ISLNK	0120 000	Reserved.		
		C_ISSOCK	0140 000	Reserved.		
	(6)	a nlinh shall contai	in the number	of links referencing th	a file at the time	
	(0)	<i>c_nunr</i> shall contain the number of links referencing the file at the tim the archive was created.				
	(7)	c_rdev shall contain implementation-defined information for character or block special files.				
	(8)	<i>c_mtime</i> shall contain the latest time of modification of the file at the time the archive was created.				
	(9)	<i>c_namesize</i> shall contain the length of the pathname, including the ter- minating null byte.				
	(10)	<i>c_filesize</i> shall contain the length of the file in bytes. This is the length of the data section following the header structure.				

Table 10-3 - Values for cpio c_mode Field

296 10.1.2.2 cpio File Name

297 c_name shall contain the pathname of the file. The length of this field in bytes is 298 the value of $c_namesize$. If a file name is found on the medium that would create 299 an invalid pathname, the implementation shall define if the data from the file is 300 stored on the file hierarchy and under what name it is stored.

All characters are represented in ISO/IEC 646 {1} IRV. For maximum portability 301 between implementations, names should be selected from characters represented 302 by the portable filename character set as 8-bit characters most significant bit zero. 303 If an implementation supports the use of characters outside the portable filename 304 character set in names for files, users, and groups, one or more implementation-305 defined encodings of these characters shall be provided for interchange purposes. 306 However, the format-reading utility shall never create file names on the local sys-307 tem that cannot be accessed via the functions described previously in this part of 308 ISO/IEC 9945; see open(), stat(), chdir(), fcntl(), and opendir(). If a file name is 309 found on the medium that would create an invalid file name, the implementation 310 shall define if the data from the file is stored on the local file system and under 311 what name it is stored. A format-reading utility may choose to ignore these files 312 as long as it produces an error indicating that the file is being ignored. 313

314 **10.1.2.3 cpio File Data**

Following c_name , there shall be $c_filesize$ bytes of data. Interpretation of such data shall occur in a manner dependent on the file. If $c_filesize$ is zero, no data shall be contained in $c_filedata$.

318 10.1.2.4 cpio Special Entries

FIFO special files, directories, and the trailer are recorded with $c_{filesize}$ equal to zero. For other special files, $c_{filesize}$ is unspecified by this part of ISO/IEC 9945. The header for the next file entry in the archive shall be written directly after the last byte of the file entry preceding it. A header denoting the file name "TRAILER!!!" shall indicate the end of the archive; the contents of bytes in the last block of the archive following such a header are undefined.

325 **10.1.2.5 cpio Values**

Values needed by the cpio archive format are described in Table 10-3.

C_ISDIR, C_ISFIFO, and C_ISREG shall be supported on a system conforming to this part of ISO/IEC 9945; additional values defined previously are reserved for compatibility with existing systems. Additional file types may be supported; however, such files should not be written on archives intended for transport to portable systems.

- C_ISVTX, C_ISCTG, C_ISLNK, and C_ISSOCK have been reserved by this part of ISO/IEC 9945 to retain compatibility with some existing implementations.
- 334 When restoring from an archive:
- (1) If the user does not have the appropriate privilege to create a file of the
 specified type, the format-interpreting utility shall ignore the entry and
 issue an error to the standard error output.
- (2) 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.

(3) 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 part of ISO/IEC 9945. If the imple mentation does not support those flags, they may be ignored.

345 **10.1.2.6 Cross-References**

346 <grp.h>, 9.2.1; <pwd.h>, 9.2.2; <sys/stat.h>, 5.6.1; chmod(), 5.6.4; link(), 347 5.3.4; mkdir(), 5.4.1; read(), 6.4.1; stat(), 5.6.2.

348 **10.1.3 Multiple Volumes**

It shall be possible for data represented by the Archive/Interchange File Formatto reside in more than one file.

The format is considered a stream of bytes. An end-of-file (or equivalently an end-of-media) condition may occur between any two bytes of the logical byte stream. If this condition occurs, the byte following the end-of-file will be the first byte on the next file. The format-reading utility shall, in an implementationdefined manner, determine what file to read as the next file.



Annex A (informative)

Bibliography

1 This Annex contains lists of related open systems standards and suggested read-2 ing on historical implementations and application programming.

3 A.1 Related Open Systems Standards

4 A.1.1 Networking Standards

- {B1} ISO 7498: 1984, Information processing systems—Open Systems Interconnection—Basic Reference Model.¹⁾
- [B2] ISO 8072: 1986, Information processing systems—Open Systems Interconnection—Transport service definition.
- {B3} ISO/IEC 8073: 1988, Information processing systems—Open Systems Interconnection—Connection oriented transport protocol specification.²⁾
- 11 {B4} ISO 8326: 1987, Information processing systems—Open Systems Inter-12 connection—Basic connection oriented session service definition.
- 13 {B5} ISO 8327: 1987, Information processing systems—Open Systems Inter 14 connection—Basic connection oriented session protocol definition.
- ISO 8348: 1987, Information processing systems—Data communications—
 Network service definition.
- ISO 8473: 1988, Information processing systems—Data communications—
 Protocol for providing the connectionless-mode network service.
- 19 {B8} ISO 8571: 1988, Information processing systems—Open Systems Inter-20 connection—File Transfer, Access and Management.

ISO documents can be obtained from the ISO office, 1, rue de Varembé, Case Postale 56, CH-1211,
 Genève 20, Switzerland/Suisse.

 ²⁾ IEC documents can be obtained from the IEC office, 3, rue de Varembé, Case Postale 131, CH 1211, Genève 20, Switzerland/Suisse.

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

- {B9} ISO 8649: 1988, Information processing systems—Open Systems Interconnection—Service definition for the Association Control Service Element.
- {B10} ISO 8650: 1988, Information processing systems—Open Systems Interconnection—Protocol specification for the Association Control Service Element.
- 30 {B11} ISO 8802-2: 1989 [IEEE Std 802.2-1989 (ANSI)], Information processing
 31 systems—Local area networks—Part 2: Logical link control.
- {B12} ISO 8802-3: 1989 [IEEE Std 802.3-1988 (ANSI)], Information processing
 systems—Local area networks—Part 3: Carrier sense multiple access with
 collision detection (CSMA/CD) access method and physical layer
 specifications.
- 4B13} ISO/IEC 8802-4: 1990 [IEEE Std 802.4-1990 (ANSI)], Information
 technology—Local area networks—Part 4: Token-passing bus access method
 and physical layer specifications.
- 40 {B14} ISO 8802-5: ... (IEEE 802.5-1989), Information technology—Local area
 40 networks—Part 5: Token ring access method and physical layer
 41 specifications.
- 42 {B15} ISO 8822: 1988, Information processing systems—Open Systems Inter-43 connection—Connection oriented presentation service definition.
- 44 {B16} ISO 8823: 1988, Information processing systems—Open Systems Inter-45 connection—Connection oriented presentation protocol specification.
- 46 {B17} ISO 8831: 1989, Information processing systems—Open Systems Inter-47 connection—Job transfer and manipulation concepts and services.
- 48 {B18} ISO 8832: 1989, Information processing systems—Open Systems Inter 49 connection—Specification of the basic class protocol for job transfer and
 50 manipulation.
- {B19} CCITT Recommendation X.25, Interface between data terminal equipment
 (DTE) and data circuit-terminating equipment (DCT) for terminals operating
 in the packet mode and connected to public data networks by dedicated circuit.³⁾
- (B20) CCITT Recommendation X.212, Information processing systems—Data
 communication—Data link service definition for Open Systems Interconnec *tion*.

⁵⁸ 59 60

³⁾ CCITT documents can be obtained from the CCITT General Secretariat, International Telecommunications Union, Sales Section, Place des Nations, CH-1211, Genève 20, Switzerland/Suisse.

61 A.1.2 Language Standards

- 62 {B21} ISO 1539: 1980, Programming languages—FORTRAN.
- 63 {B22} ISO 1989: 1985, Programming Languages—COBOL.
- 64 {B23} ISO 8652: 1987, Programming Languages—Ada.
- (B24) ANSI X3.113-1987⁴, Information systems—Programming language—FULL
 BASIC.
- (B25) ANSI/IEEE 770X3.97-1983, Standard Pascal Computer Programming
 Language.
- 69 {B26} ANSI/MDC X11.1-1984, Programming Language MUMPS.

70 A.1.3 Graphics Standards

- {B27} ISO 7942: 1985, Information processing systems—Computer graphics—
 Graphical Kernel System (GKS) functional description.
- {B28} ISO 8632: 1987, Information processing systems—Computer graphics—
 Metafile for the storage and transfer of picture description information.
- {B29} ISO/IEC 9592: 1989 (ANSI X3.144-1988), Information processing systems—
 Computer graphics—Programmer's hierarchical interactive graphics system (PHIGS).

78 A.1.4 Database Standards

- 79 {B30} ISO 8907: 1987, Database Language—NDL.
- 80 {B31} ISO 9075: 1987, Database Language—SQL.

81 A.2 Other Standards

- 82 {B32} ISO 639: 1988, Code for the representation of names of languages.
- 83 {B33} ISO 3166: 1988, Code for the representation of names of countries.
- {B34} ISO 8859-1: 1987, Information Processing—8-bit single-byte coded graphic
 character sets—Part 1: Latin alphabet No. 1.
- {B35} ISO 9127: 1988, Information processing systems—User documentation and cover information for consumer software packages.
- {B36} ISO/IEC 9945-2:...,⁵⁾ Information technology—Portable operating system
 interface (POSIX)—Part 2: Shell and utilities.

 ⁴⁾ ANSI documents can be obtained from the Sales Department, American National Standards
 91 Institute, 1430 Broadway, New York, NY 10018.

⁵⁾ To be approved and published.

- 4837 ISO/IEC 10646: ...,⁶⁾ Information processing—Multiple octet coded charac ter set.
- {B38} IEEE Std 100-1988, IEEE Standard Dictionary of Electrical and Electronics
 Terms.

97 A.3 Historical Documentation and Introductory Texts

- (B39) American Telephone and Telegraph Company. System V Interface
 Definition (SVID), Issues 2 and 3. Morristown, NJ: UNIX Press, 1986,
 1989.⁷⁾
- 101 {B40} American Telephone and Telegraph Company. UNIX System III
 102 Programmer's Manual. Greensboro, NC: Western Electric Company,
 103 October 1981.
- (B41) American Telephone and Telegraph Company. UNIX Time Sharing System:
 UNIX Programmer's Manual. 7th ed. Murray Hill, NJ: Bell Telephone
 Laboratories, January 1979.
- (B42) "The UNIX System."⁸⁾ AT&T Bell Laboratories Technical Journal. vol. 63 (8
 Part 2), October 1984.
- (B43) "UNIX Time-Sharing System."⁹⁾ Bell System Technical Journal. vol. 57 (6
 Part 2), July-August 1978.
- (B44) Bach, Maurice J. *The Design of the UNIX Operating System*. Englewood
 Cliffs, NJ: Prentice-Hall, 1987.
- (B45) Harbison, Samuel P. and Steele, Guy L. C: A Reference Manual. Englewood Cliffs, NJ: Prentice-Hall, 1987.
- (B46) Kernighan, Brian W. and Ritchie, Dennis M. The C Programming
 Language. Englewood Cliffs, NJ: Prentice-Hall, 1978.
- (B47) Kernighan, Brian W. and Pike, Rob. *The UNIX Programming Environment*.
 Englewood Cliffs, NJ: Prentice-Hall, 1984.
- (B48) Leffler, Samuel J., McKusick, Marshall Kirk, Karels, Michael J., Quarter man, John S., and Stettner, Armando. *The Design and Implementation of the 4.3BSD UNIX Operating System.* Reading, MA: Addison-Wesley, 1988.
- (B49) McGilton, Henry and Morgan, Rachel. Introducing the UNIX System. New
 York: McGraw-Hill (BYTE Books), 1983.

^{124 6)} To be approved and published.

This is one of several documents that represent an industry specification in an area related to
 POSIX.1. The creators of such documents may be able to identify newer versions that may be
 interesting.

^{128 8)} This entire edition is devoted to the UNIX system.

^{129 9)} This entire edition is devoted to the UNIX time-sharing system.

- (B50) Organick, Elliot I. The Multics System: An Examination of Its Structure.
 Cambridge, MA: The MIT Press, 1972.
- (B51) Quarterman, John S., Silberschatz, Abraham, and Peterson, James L.
 "4.2BSD and 4.3BSD as Examples of the UNIX System." ACM Computing Surveys. vol. 17 (4), December 1985, pp. 379–418.
- (B52) Ritchie, Dennis M. "Reflections on Software Research." Communications of the ACM. vol. 27 (8), August 1984, pp. 758–760. ACM Turing Award Lecture.
- (B53) Ritchie, Dennis. "The Evolution of the UNIX Time-Sharing System." AT&T
 Bell Laboratories Technical Journal. vol. 63 (8), October 1984, pp.
 1577–1593.
- (B54) Ritchie, D. M. and Thompson, K. "The UNIX Time-Sharing System." Communications of the ACM. vol. 7 (7), July 1974, pp. 365–375. This is the original paper, which describes Version 6.
- [H4] {B55} Ritchie, D. M. and Thompson, K. "The UNIX Time-Sharing System." Bell
 [H45] System Technical Journal. vol. 57 (6 Part 2), July-August 1978, pp.
 [H66] 1905–1929. This is a revised version and describes Version 7.
- (B56) Ritchie, Dennis M. "Unix: A Dialectic." Winter 1987 USENIX Association *Conference Proceedings, Washington, D.C.*, pp. 29–34. Berkeley, CA:
 USENIX Association, January 1987.
- (B57) Rochkind, Marc J. Advanced UNIX Programming. Englewood Cliffs, NJ:
 Prentice-Hall, 1985.
- (B58) University of California at Berkeley—Computer Science Research Group.
 4.3 Berkeley Software Distribution, Virtual VAX-11 Version. Berkeley, CA:
 The Regents of the University of California, April 1986.
- (B59) /usr/group Standards Committee. 1984 /usr/group Standard. Santa
 Clara, CA: UniForum, 1984.
- (B60) X/Open Company, Ltd. X/Open Portability Guide, Issue 2. Amsterdam:
 Elsevier Science Publishers, 1987.
- (B61) X/Open Company, Ltd. X/Open Portability Guide, Issue 3. Englewood
 Cliffs, NJ: Prentice-Hall, 1989.



Annex B

(informative)

Rationale and Notes

The Annex is being published as an informative part of POSIX.1 to assist in the process of review. It contains historical information concerning the contents of POSIX.1 and why features were included or discarded. It also contains notes of interest to application programmers on recommended programming practices, emphasizing the consequences of some aspects of POSIX.1 that may not be immediately apparent.¹⁾

7 B.1 Scope and Normative References

8 **B.1.1 Scope**

9 This Rationale focuses primarily on additions, clarifications, and changes made to 10 the UNIX system, from which POSIX.1 was derived. It is not a rationale for the 11 UNIX system as a whole, since the goal of POSIX.1's developers was to codify exist-12 ing practice, not design a new operating system. No attempt is made in this 13 Rationale to defend the pre-existing structure of UNIX systems. It is primarily 14 deviations from existing practice, as codified in the base documents, that are 15 explained or justified here.

Material that is "outside the scope" or otherwise not addressed by this part of ISO/IEC 9945 is implicitly "unspecified." It may be included in an implementation, and thus the implementation does provide a specification for it. The term "implementation-defined" has a specific meaning in POSIX.1 and is not a synonym for "defined (or specified) by the implementation."

The Rationale discusses some UNIX system features that were *not* adopted into POSIX.1. Many of these are features that are popular in some UNIX system implementations, so that a user of those implementations might question why they do not appear in POSIX.1. This Rationale should provide the appropriate answers.

The material in this annex is derived in part from copyrighted draft documents developed under
 the sponsorship of UniForum, as part of an ongoing program of that association to support the
 POSIX standards program efforts.

T

I

I

I

There are choices allowed by POSIX.1 for some details of the interface specification; some of these are specifiable optional subsets of POSIX.1. See B.2.9.

Although the services POSIX.1 provides have been defined in the C language, the 30 concept of providing fundamental, standardized services should not be restricted 31 only to programs of a particular programming language. The possibility of imple-32menting interfaces in alternate programming languages inspired the term 33 POSIX.1 with the C Language Binding. The word Binding refers to the binding of 34 a conceptual set of services and a standardized C interface that establishes rules 35 and syntax for accessing them. Future international standards are expected to 36 separate the C language binding from the language-independent services of 37 POSIX.1 and to include bindings for other programming languages. 38

The C Standard {2} will be the basis for functional definitions of core services that are independent of programming languages. POSIX.1 as it stands now can be thought of as a C Language Binding. Sections 1 through 7, and 9, correspond roughly to the C language implementation of what will be defined in the programming language-independent core services portion of POSIX.1; Section 8 corresponds to the C language-specific portion.

The criteria used to choose the programming language-independent core services 45 may be different from those expected. The core services represent services that 46 are common to those programming languages likely to form language bindings to 47 POSIX.1—the greatest common denominator. They are not chosen to reflect the 48 most important system services of an ideal operating system. For this reason, 49 some fundamental system services are not included in the language-independent 50 core. As an example, memory management routines would at first seem to be a 51 core service-they are an absolutely fundamental system service. They must, 52 however, be included in language-specific portions of POSIX.1 because program-53 ming languages such as FORTRAN have traditionally not provided memory 54 management. Categorizing memory management as a core service would impose 55 unreasonable requirements for FORTRAN implementations. 56

Any programming language traditionally supporting memory management should include those routines in the language-dependent portions of their bindings. Work will be done at a later time to standardize the classes of functions that must be included in the language-dependent portions of language bindings if those functions have been traditionally implemented for that language. This will ensure that certain classes of critical functions, such as memory management, will not be excluded from any applicable language binding; see B.1.3.3.

POSIX.1 is not a tutorial on the use of the specified interface, nor is this Rationale.
 However, the Rationale includes some references to well-regarded historical documentation on the UNIX System in A.3.

67 **B.1.1.1 POSIX.1 and the C Standard**

Some C language functions and definitions were handled by POSIX.1, but most
were handled by the C Standard {2}. The general guideline is that POSIX.1
retained responsibility for operating-system specific functions, while the
C Standard {2} defined C library functions. See also B.2.7 and B.8.

90

91

92

93

94

95

96

97

98

99

100

101

102

- 72 There are several areas in which the two standards differ philosophically:
- (1) Function parameter type lists. These appear in the syntax of the 73 C Standard {2}. In this version of POSIX.1, the parameter lists were res-74 tated in terms of these function prototypes. There were two major rea-75 sons for making this change from IEEE Std 1003.1-1988: the use of the 76 C Standard {2} was rapidly becoming more widespread, and implemen-77 tors were experiencing difficulties with some of the function prototypes 78 where guidance was not provided in POSIX.1. (The modifier const pro-79 vided the most difficulty.) Specific guidance and permission remains in 80 POSIX.1 for translation to common-usage C. 81
- 82 (2) Single vs. multiple processes. The C Standard {2} specifies a language
 83 that can be used on single-process operating systems and as a freestand84 ing base for the implementation of operating systems or other stand85 alone programs. However, the POSIX.1 interface is that of a multiprocess
 86 timesharing system. Thus, POSIX.1 has to take multiple processes into
 87 account in places where the C Standard {2} does not mention processes at
 88 all, such as kill(). See also B.1.3.1.1.
 - (3) Single vs. multiple operating system environments. The C Standard {2} specifies a language that may be useful on more than one operating system and thus has means of tailoring itself to the particular current environment. POSIX.1 is an operating system interface specification and thus by definition is only concerned with one operating system environment, even though it has been carefully written to be broadly implementable (see Broadly Implementable in the Introduction) in terms of various underlying operating systems. See also B.1.3.1.1.
 - (4) Translation vs. execution environment. POSIX.1 is primarily concerned with the C Standard {2} execution environment, leaving the translation environment to the C Standard {2}. See also B.1.3.1.1.
 - (5) Hosted vs. freestanding implementations. All POSIX.1 implementations are hosted in the sense of the C Standard {2}. See also the remarks on conformance in the Introduction.
- Text vs. binary file modes. The C Standard (2) defines text and binary 103 (6) modes for a file. But the POSIX.1 interface and historical implementa-104 tions related to it make no such distinction, and all functions defined by 105 POSIX.1 treat files as if these modes were identical. (It should not be 106 stated that POSIX.1 files are either *text* or *binary*.) The definitions in the 107 C Standard {2} were written so that this interpretation is possible. In 108 particular, text mode files are not required to end with a line separator, 109 which also means that they are not required to include a line separator 110 at all. 111

Furthermore, there is a basic difference in approach between the Rationale accompanying the C Standard {2} and this Rationale Annex. The C Standard {2} Rationale, a separate document, addresses almost all changes as differences from the Base Documents of the C Standard {2}, usually either Kernighan and Ritchie {B46} or the 1984 /usr/group Standard {B59}. This Rationale cannot do that, since there are many more variants of (and Base Documents for) the operating system interface than for the C language. The most noticeable aspect of this

1

1

I

l

I

1

1

difference is that the C Standard {2} Rationale identifies "QUIET CHANGES" from

- the Base Documents. This Annex cannot include such markings, since a quiet
- change from one historical implementation may correspond exactly to another his-
- torical implementation, and may be very noticeable to an application written for vet another.
- The following subclauses justify the inclusion or omission of various C language functions in POSIX.1 or the C Standard {2}.

126 **B.1.1.1.1 Solely by POSIX.1**

- 127 These return parameters from the operating system environment: *ctermid()*, | 128 *ttyname()*, and *isatty()*.
- The *fileno()* and *fdopen()* functions map between C language stream pointers and POSIX.1 file descriptors.

131 B.1.1.1.2 Solely by the C Standard

There are many functions that are useful with the operating system interface and are required for conformance with POSIX.1, but that are properly part of the C Language. These are listed in 8.1, which also notes which functions are defined by both POSIX.1 and the C Standard {2}. Certain terms defined by the C Standard {2} are incorporated by POSIX.1 in 2.7.

Some routines were considered too specialized to be included in POSIX.1. These
 include *bsearch()* and *qsort()*.

139 B.1.1.1.3 By Neither POSIX.1 Nor the C Standard

Some functions were considered of marginal utility and problematical when international character sets were considered: _toupper(), _tolower(), toascii(), and isascii().

Although *malloc()* and *free()* are in the C Standard {2} and are required by 8.1 of POSIX.1, neither *brk()* nor *sbrk()* occur in either standard (although they were in the *1984 /usr/group Standard* {B59}), because POSIX.1 is designed to provide the basic set of functions required to write a Conforming POSIX.1 Application; the underlying implementation of *malloc()* or *free()* is not an appropriate concern for POSIX.1.

149 B.1.1.1.4 Base by POSIX.1, Additions by the C Standard

Since the C Standard {2} does not depend on POSIX.1 in any way, there are no items in this category.

152 B.1.1.1.5 Base by the C Standard, Additions by POSIX.1

The C Standard {2} has to define *errno* if only because examining that variable offers the only way to determine when some mathematics routines fail. But POSIX.1 uses it more extensively and adds some semantics to it in 2.4, which also defines some values for it.

157 Many numerical limits used by the C Standard {2} were incorporated by POSIX.1 | 158 in 2.8, and some new ones were added, all to be found in the header <limits.h>.

- 159 The C Standard {2} provides *signal()*, a minimal functionality for interrupts. The
- POSIX.1 definition replaces this with an elaborate mechanism that deals with multiple processes and is reliable when signals come from outside sources.
- 162 The *time()* function is used by the C Standard {2}, but POSIX.1 further specifies 163 the time value.
- The getenv() function is referenced in 2.6 and 3.1.2 and is also defined by the C Standard $\{2\}$.
- 166 The *rename()* function is extended to further specify its behavior when the new 167 filename already exists or either argument refers to a directory.
- 168 The *setlocale()* function and the handling of time zones were further specified to 169 take advantage of the POSIX environment.
- The standard-I/O functions were specified in terms of their relationship to file descriptors and the relationship between multiple processes.

172 B.1.1.1.6 Related Functions by Both

- The C Standard {2} definition of *compliance* and the POSIX.1 definition of *conformance* are similar, although the latter notes certain potential hardware limitations.
- POSIX.1 defined a portable filename character set in 2.2.2 that is like the C Standard {2} identifier character set. However, POSIX.1 did not allow upperand lowercase characters to be considered equivalent. See *filename portability* in 2.3.4.
- The exit() function is defined only by the C Standard {2} because it refers to closing streams, and that subject, as well as *fclose()* itself, is defined almost entirely by the C Standard {2}. But POSIX.1 defined _exit(), which also adds semantics to exit(). This allows POSIX.1 to omit references to the C Standard {2} atexit() function.
- POSIX.1 defined *kill()*, while the C Standard {2} defined *raise()*, which is similar except that it does not have a process ID argument, since the language defined by the C Standard {2} does not incorporate the idea of multiple processes.
- The new functions *sigsetjmp()* and *siglongjmp()* were added to provide similar functions to the C Standard {2} *setjmp()* and *longjmp()* that additionally save and restore signal state.

191 B.1.2 Normative References

192 There is no additional rationale provided for this subclause.

193 **B.1.3 Conformance**

These conformance definitions are descended from those of *conforming implementation*, *conforming application*, and *conforming portable application* of early drafts, but were changed to clarify

- 197 (1) Extensions, options, and limits;
- 198 (2) Relations among the three terms, and;
- (3) Relations between POSIX.1 and the C Standard {2}.

200 B.1.3.1 Implementation Conformance

- These definitions allow application developers to know what to depend on in an implementation.
- There is no definition of a *strictly conforming implementation*; that would be an implementation that provides *only* those facilities specified by POSIX.1 with no extensions whatsoever. This is because no actual operating system implementation can exist without system administration and initialization facilities that are beyond the scope of POSIX.1.

208 B.1.3.1.1 Requirements

The word "support" is used, rather than "provide," in order to allow an implementation that has no resident software development facilities, but that supports the execution of a *Strictly Conforming POSIX.1 Application*, to be a *conforming implementation*. See also B.1.1.1.

213 B.1.3.1.2 Documentation

Note that the use of "may" in terms of where conformance documents record where implementations may vary implies that it is not required to describe those features identified as undefined or unspecified.

Other aspects of systems must be evaluated by purchasers for suitability. Many 222 systems incorporate buffering facilities, maintaining updated data in volatile 223 storage and transferring such updates to nonvolatile storage asynchronously. 224 Various exception conditions, such as a power failure or a system crash, can cause 225 this data to be lost. The data may be associated with a file that is still open, with 226 one that has been closed, with a directory, or with any other internal system data 227 structures associated with permanent storage. This data can be lost, in whole or 228 part, so that only careful inspection of file contents could determine that an 229 update did not occur. 230

Also, interrelated file activities, where multiple files and/or directories are updated, or where space is allocated or released in the file system structures, can leave inconsistencies in the relationship between data in the various files and directories, or in the file system itself. Such inconsistencies can break applications that expect updates to occur in a specific sequence, so that updates in one place correspond with related updates in another place.

For example, if a user creates a file, places information in the file, and then 237 records this action in another file, a system or power failure at this point followed 238by restart may result in a state in which the record of the action is permanently 239 recorded, but the file created (or some of its information) has been lost. The 240 consequences of this to the user may be undesirable. For a user on such a system, 241 the only safe action may be to require the system administrator to have a policy 242 that requires, after any system or power failure, that the entire file system must 243 be restored from the most recent backup copy (causing all intervening work to be 244 lost). 245

The characteristics of each implementation will vary in this respect and may or may not meet the requirements of a given application or user. Enforcement of such requirements is beyond the scope of POSIX.1. It is up to the purchaser to determine what facilities are provided in an implementation that affect the exposure to possible data or sequence loss and also what underlying implementation techniques and/or facilities are provided that reduce or limit such loss or its consequences.

B.1.3.1.3 Conforming Implementation Options

Within POSIX.1 there are some symbolic constants that, if defined, indicate that a certain option is enabled. Other symbolic constants exist in POSIX.1 for other reasons. This clause helps clarify which constants are related to true "options" and which are related more to the behavior of differing systems.

To accommodate historical implementations where there were distinct semantics 258in certain situations, but where one was not clearly better or worse than another, 259 early drafts of POSIX.1 permitted either of (typically) two options using "may." At 260 the request of the working group developing test assertions, this was changed to 261 be specified by formal options with flags. It quickly became obvious that these 262 would be treated as options that could be selected by a purchaser, when the intent 263 of the developers of POSIX.1 was to allow either behavior (or both, in some cases) 264 to conform to the standard, and to constrain the application to accommodate 265 either. Thus, these options were removed and the phrase "An implementation 266may either" introduced to replace the option. Where this phrase is used, it indi-267 cates that an application shall tolerate either behavior. 268

It is intended that all conforming applications shall tolerate either behavior and that only in the most exceptional of circumstances (driven by technical need) should a purchaser specify only one behavior. Backwards compatibility is not considered exceptional, as this is not consistent with the intent of POSIX.1: to promote the portability of applications (and the development of portable applications).

An application can tolerate these behaviors either by ignoring the differences (if they are irrelevant to the application) or by taking an action to assure a known state. It might be that that action would be redundant on some implementations.

Validation programs, which are applications in this sense, could either report the actual result found or simply ignore the difference. In no case should either

acceptable behavior be treated as an error. This may complicate the validation
 slightly, but is more consistent with the intent of this permissible variation in
 behavior.

In certain circumstances, the behavior may vary for a given process. For example, in the presence of networked file systems, whether or not dot and dot-dot are present in the directory may vary with the directory being searched, and the program would only be portable if it tolerated, but did not require, the presence of these entries in a directory.

In situations like this, it is typically easier to simply ignore dot and dot-dot if they are found than to try to determine if they should be expected or not.

290 B.1.3.2 Application Conformance

These definitions guide users or adaptors of applications in determining on which implementations an application will run and how much adaptation would be required to make it run on others. These three definitions are modeled after related ones in the C Standard {2}.

POSIX.1 occasionally uses the expressions *portable application* or *conforming application*. As they are used, these are synonyms for any of these three terms. The differences between the three classes of application conformance relate to the requirements for other standards, or, in the case of the Conforming POSIX.1 Application Using Extensions, to implementation extensions. When one of the less explicit expressions is used, it should be apparent from the context of the discussion which of the more explicit names is appropriate.

302 B.1.3.2.1 Strictly Conforming POSIX.1 Application

- 303 This definition is analogous to that of a C Standard {2} conforming program.
- The major difference between a *Strictly Conforming POSIX.1 Application* and a C Standard (2) *strictly conforming program* is that the latter is not allowed to use features of POSIX.1 that are not in the C Standard (2).
- 307 B.1.3.2.2 Conforming POSIX.1 Application
- 308 Examples of *<National Bodies>* include ANSI, BSI, and AFNOR.

309 B.1.3.2.3 Conforming POSIX.1 Application Using Extensions

Due to possible requirements for configuration or implementation characteristics in excess of the specifications in 2.8 or related to the hardware (such as array size or file space), not every Conforming POSIX.1 Application Using Extensions will run on every conforming implementation.

B.1.3.3 Language-Dependent Services for the C Programming Language

POSIX.1 is, for historical reasons, both a specification of an operating system interface and a C binding for that specification. It is clear that these need to be separated into unique entities, but the urgency of getting the initial standard out, and the fact that C is the *de facto* primary language on systems similar to the

Ι

319 UNIX system, makes this a necessary and workable situation.

Nevertheless, work will be done on language bindings, beyond that for C before the specification and the current binding are separated. Language bindings for languages other than C should not model themselves too closely on the C binding and in the process pick up various idiosyncrasies of C.

Where functionality is duplicated in POSIX.1 [e.g., open() and creat()] there is no reason for that duplication to be carried forward into another language. On the other hand, some languages have functionality already in them that is essentially the same as that provided in POSIX.1. In this case, a mapping between the functionality in that language and the underlying functionality in POSIX.1 is a better choice than mimicking the C binding.

- Since C has no syntax for I/O, and I/O is a large fraction of POSIX.1, the paradigm 330 of functions has been used. This may not be appropriate to another language. 331 For example, FORTRAN's REWIND statement is a candidate to map onto a special 332 case of *lseek(*), and its SEEK statement may completely cover for *lseek(*). If this is 333 the case, there is no reason to provide SUBROUTINEs with the same functionality. 334 In the more general case, file descriptors and FORTRAN's logical unit numbers 335 may have a useful mapping. FORTRAN'S ERR= option in I/O operations might 336 replace returning -1; the whole concept of errors might be handled differently. 337
- As was done with C, it is not unreasonable for other language bindings to specify some areas that are undefined or unspecified by the underlying language standard or that are permissible as extensions. This may, in fact, solve some difficult problems.
- Using as much as possible of the target language in the binding enhances portability. If a program wishes to use some POSIX.1 capabilities, and these are bound to the language statements rather than appearing as additional procedure or function calls, and the program does in fact conform to the language standard while using those functions, it will port to a larger range of systems than one that is obligated to use procedure or function calls introduced specifically for the binding to POSIX.1 to do the same thing.
- A program that requires the POSIX.1 capabilities that are not bound to the stan-349 dard language directly (as above) has no chance to be portable outside the POSIX.1 350 environment. It does not matter whether the extension is syntactic or a new func-351 tion; it still will not port without effort. Given this, it seems unreasonable not to 352 consider language extensions when determining how best to map the functionality 353 of POSIX.1 into a particular language binding. For example, a new statement 354 similar to READ, which loads the values from a call like stat(), might be the best 355 solution for reading the data lists returned as structures in C into a list of FOR-356 TRAN variables. 357
- No attempt to mimic *printf()* or *scanf()* (or the rest of the C Standard {2} functions) should be made; the equivalent functions in the language should be used. (Formatted READ and WRITE in FORTRAN, read/readln and write/writeln in Pascal, for example.)
- There is an inherent special relationship between an operating system standard and a language standard. It is unlikely that standards for other kinds of features (such as graphics) will bind directly to statements in a general purpose language.

However, an operating system standard should provide the services required by a language. This is an unusual situation, and the tendency to use only new functions and procedures when creating a binding should be examined carefully. (A one-to-one binding in all cases is probably not possible, but bindings such as those for standard I/O in Section 8 may be possible.)

Binding directly to the language, where possible, should be encouraged both by making maximal use of the mapping between the operating system and the language that naturally exists and, where appropriate, by having the languages request changes to the operating system to facilitate such a mapping. (A future inclusion of a truncate function, specifically for the FORTRAN ENDFILE statement, but that is also generally useful, is a good example.)

Part of the job of creating a binding is choosing names for functions that are intro-376 duced, and these will need to be appropriate for that language. It is possible to 377 use other than the most restrictive form of a name, since, as discussed previously, 378 using these functions inherently makes the application not portable to systems 379 that are not POSIX.1, and if POSIX.1 conformant systems typically accept names 380 that the lowest-common-denominator system will not, there is no reason to a381 priori exclude such names. (The specific example is C, where it is typically "non-382 UNIX" systems that limit external identifiers to six characters.) 383

384 See B.1.1 for additional information about C bindings.

385 B.1.3.3.1 Types of Conformance

- ³⁸⁶ There is no additional rationale provided for this subclause.
- 387 B.1.3.3.2 C Standard Language-Dependent System Support
- The issue of "namespace pollution" needs to be understood in this context. See B.2.7.2.
- 390 B.1.3.3.3 Common-Usage C Language-Dependent System Support
- The issue of "namespace pollution" needs to be understood in this context. See B.2.7.2.

393 B.1.3.4 Other C Language-Related Specifications

The information concerning the use of library functions was adapted from a description in the C Standard {2}. Here is an example of how an application program can protect itself from library functions that may or may not be macros, rather than true functions:

- ³⁹⁸ The *atoi*() function may be used in any of several ways:
- (1) By use of its associated header (possibly generating a macro expansion)
- 400 #include <stdlib.h>
- 401 /* ... */
- 402 i = atoi(str);
- 403 (2) By use of its associated header (assuredly generating a true function call)
| 404
405
406
407 | | <pre>#include <stdlib.h> #undef atoi /* */ i = atoi(str);</stdlib.h></pre> |
|--------------------------|-----|--|
| 408 | | or |
| 409
410
411 | | <pre>#include <stdlib.h> /* */ i = (atoi) (str);</stdlib.h></pre> |
| 412 | (3) | By explicit declaration |
| 413
414
415 | | <pre>extern int atoi (const char *); /* */ i = atoi(str);</pre> |
| 416 | (4) | By implicit declaration |
| 417
418 | | <pre>/* */ i = atoi(str);</pre> |

419(Assuming no function prototype is in scope. This is not allowed by the420C Standard {2} for functions with variable arguments; furthermore,421parameter type conversion "widening" is subject to different rules in this422case.)

Note that the C Standard {2} reserves names starting with '_' for the compiler.
Therefore, the compiler could, for example, implement an intrinsic, built-in function _asm_builtin_atoi(), which it recognized and expanded into inline assembly
code. Then, in <stdlib.h>, there could be the following:

The user's "normal" call to *atoi*() would then be expanded inline, but the implementor would also be required to provide a callable function named *atoi*() for use when the application requires it; for example, if its address is to be stored in a function pointer variable.

432 **B.1.3.5 Other Language-Related Specifications**

It is intended that "long" identifiers and multicase linkage would be supported on POSIX.1 systems for all languages, including C. This is where that condition is stated. The portion of the sentence about "if such extensions are" is included to permit languages that have an absolute maximum, or an absolute requirement of case folding, to be conformant.

- 438 The requirement for longer names is included for several reasons:
- (1) Most systems similar to POSIX.1 are already conformant.
- 440 (2) Many existing language standards restrict the length of names to accommodate existing systems that cannot be modified to allow longer names.
 442 However, those systems are not expected to be POSIX.1 conformant, for other reasons.

- (3) Many historical applications rely on such long names.
- 445
- (4) Future languages (such as FORTRAN 88) are likely to require it.

Specific to FORTRAN 77 {B21}, that standard permits long names, and this part of ISO/IEC 9945 requires that FORTRAN implementations running on POSIX.1 support long names. The requirements of case distinction and length are considered orthogonal, but both are required if both are permitted by the language. Note that a language can be conformant to POSIX.1 even though a binding does not exist, because an application need not step outside the language standard to write a useful program.

This requirement permits the use of reasonable-length names in a POSIX.1 binding to a language such as FORTRAN. Clearly nothing prohibits a program that does conform to the FORTRAN minima to compile and run on POSIX.1.

It is within the constraints of POSIX.1 to specify the behavior of the language processors and linker, consistent with the language, as it is a specification for an execution environment. This is different than a package such as GKS {B27}, which can reasonably be expected to be ported to a system that enforces the language minima.

461 It might be argued that this specification is appropriate to the language binding committees for POSIX generally, rather than specifically to POSIX.1. That argu-462 ment misses the intent. The intent is to require that the linker and other code 463 that handles "object code" (a concept not formally defined in POSIX.1) are able to 464 support long names. This requirement, being one that spans all languages, 465 belongs in the specification standard, rather than tied to any one language. Note 466 that it is also somewhat permissive, in that if the language is unable to deal with 467 long names it is permitted not to require them, but it does remove the argument 468 that "the loader might not permit long names, so [a specific] language binding 469 should not force the issue." 470

A strictly conforming application for a given language could not use any extensions outside of POSIX.1 for that language (regardless of the underlying operating system). An application will strictly conform to POSIX.1 if it conforms to the language using additional interfaces from that language's binding to POSIX.1.

475 **B.2 Definitions and General Requirements**

476 **B.2.1 Conventions**

477 There is no additional rationale provided for this subclause.

478 **B.2.2 Definitions**

479 **B.2.2.1 Terminology**

The meanings specified in POSIX.1 for the words *shall*, *should*, and *may* are mandated by ISO/IEC directives.

In this Rationale, the words *shall*, *should*, and *may* are sometimes used to illustrate similar usages in the standard. However, the Rationale itself does not specify anything regarding implementations or applications.

conformance document: As a practical matter, the conformance document is
effectively part of the system documentation. They are distinguished by POSIX.1
so that they can be referred to distinctly.

implementation defined: This definition is analogous to that of the
 C Standard {2} and, together with *undefined* and *unspecified*, provides a range of
 specification of freedom allowed to the interface implementor.

may: The use of may has been limited as much as possible, due both to confusion stemming from its ordinary English meaning and to objections regarding the
desirability of having as few options as possible and those as clearly specified as
possible.

shall: Declarative sentences are sometimes used in POSIX.1 as if they included
the word *shall*, and facilities thus specified are no less required. For example, the
two statements:

- 498 (1) The *foo*() function shall return zero
- 499 (2) The foo() function returns zero

are meant to be exactly equivalent. It is expected that a future version of POSIX.1
 will be rewritten to use the "shall" form more consistently.

should: In POSIX.1, the word *should* does not usually apply to the implementation, but rather to the application. Thus, the important words regarding implementations are *shall*, which indicates requirements, and *may*, which indicates options.

506 **obsolescent:** The term *obsolescent* was preferred over *deprecated* to represent 507 functionality that should not be used in new work. The term *obsolescent* is more 508 intuitive and reduced the possibility of misunderstanding in the intended context.

supported: An example of this concept is the *setpgid()* function. If the imple mentation does not support the optional job control feature, it nevertheless has to
 provide a function named *setpgid()*, even though its only ability is that of return ing [ENOSYS].

513 **system documentation:** The system documentation should normally describe 514 the whole of the implementation, including any extensions provided by the imple-515 mentation. Such documents normally contain information at least as detailed as 516 the POSIX.1 specifications. Few requirements are made on the system documen-517 tation, but the term is needed to avoid a dangling pointer where the conformance 518 document is permitted to point to the system documentation.

- ⁵¹⁹ **undefined:** See *implementation defined*.
- 520 **unspecified:** See *implementation defined*.

The definitions for *unspecified* and *undefined* appear nearly identical at first examination, but are not. *Unspecified* means that a conforming program may deal with the unspecified behavior, and it should not care what the outcome is. *Undefined* says that a conforming program should not do it because no definition is provided for what it does (and implicitly it would care what the outcome was if it tried it). It is important to remember, however, that if the syntax permits the statement at all, it must have some outcome in a real implementation.

- Thus, the terms *undefined* and *unspecified* apply to the way the application should think about the feature. In terms of the implementation it is always "defined"—there is always some result, even if it is an error. The implementation is free to choose the behavior it prefers.
- This also implies that an implementation, or another standard, could specify or define the result in a useful fashion. The terms apply to POSIX.1 specifically.

The term *implementation defined* implies requirements for documentation that 534 are not required for undefined (or unspecified). Where there is no need for a con-535 536 forming program to know the definition, the term *undefined* is used, even though *implementation defined* could also have been used in this context. There could be 537 a fourth term, specifying "POSIX.1 does not say what this does; it is acceptable to 538 define it in an implementation, but it does not need to be documented," and 539 undefined would then be used very rarely for the few things for which any 540 definition is not useful. 541

In many places POSIX.1 is silent about the behavior of some possible construct. 542 For example, a variable may be defined for a specified range of values and 543 behaviors are described for those values; nothing is said about what happens if 544 the variable has any other value. That kind of silence can imply an error in the 545 standard, but it may also imply that the standard was intentionally silent and 546 that any behavior is permitted. There is a natural tendency to infer that if the 547 standard is silent, a behavior is prohibited. That is not the intent. Silence is 548 intended to be equivalent to the term *unspecified*. 549

550 B.2.2.2 General Terms

Many of these definitions are necessarily circular, and some of the terms (such as *process*) are variants of basic computing science terms that are inherently hard to define. Some are defined by context in the prose topic descriptions of the general concepts in 2.3, but most appear in the alphabetical glossary format of the terms in 2.2.2.

Some definitions must allow extension to cover terms or facilities that are not explicitly mentioned in POSIX.1. For example, the definition of *file* must permit interpretation to include streams, as found in the Eighth Edition (a research version of the UNIX system). The use of abstract intermediate terms (such as *object* in place of, or in addition to, *file*) has mostly been avoided in favor of careful definition of more traditional terms.

1

562 Some terms in the following list of notes do not appear in POSIX.1; these are 563 marked prefixed with a asterisk (*). Many of them have been specifically 564 excluded from POSIX.1 because they concern system administration, implementa-565 tion, or other issues that are not specific to the programming interface. Those are 566 marked with a reason, such as "implementation defined."

appropriate privileges: One of the fundamental security problems with many 567 historical UNIX systems has been that the privilege mechanism is monolithic—a 568 user has either no privileges or all privileges. Thus, a successful "trojan horse" 569 attack on a privileged process defeats all security provisions. Therefore, POSIX.1 570 allows more granular privilege mechanisms to be defined. For many historical 571 implementations of the UNIX system, the presence of the term appropriate 572 privileges in POSIX.1 may be understood as a synonym for super-user (UID 0). 573 However, future systems will undoubtedly emerge where this is not the case and 574 each discrete controllable action will have appropriate privileges associated with 575 it. Because this mechanism is *implementation defined*, it must be described in 576 the conformance document. Although that description affects several parts of 577 POSIX.1 where the term *appropriate privilege* is used, because the term *implemen*-578 tation defined only appears here, the description of the entire mechanism and its 579 effects on these other sections belongs in clause 2.3 of the conformance document. 580 This is especially convenient for implementations with a single mechanism that 581 applies in all areas, since it only needs to be described once. 582

clock tick: The C Standard {2} defines a similar interval for use by the *clock()* function. There is no requirement that these intervals be the same. In historical implementations these intervals are different. Currently only the *times()* function uses values stated in terms of clock ticks, although other functions might use them in the future.

- **controlling terminal:** The question of which of possibly several special files referring to the terminal is meant is not addressed in POSIX.1.
- 590
- *device number: The concept is handled in *stat()* as *ID of device*.

directory: The format of the directory file is implementation defined and differs radically between System V and 4.3BSD. However, routines (derived from 4.3BSD) for accessing directories are provided in 5.1.2 and certain constraints on the format of the information returned by those routines are made in 5.1.1.

- directory entry: Throughout the document, the term link is used [about the link() function, for example] in describing the objects that point to files from directories.
- **dot:** The symbolic name *dot* is carefully used in POSIX.1 to distinguish the working directory filename from a period or a decimal point.

dot-dot: Historical implementations permit the use of these filenames without
their special meanings. Such use precludes any meaningful use of these
filenames by a Conforming POSIX.1 Application. Therefore, such use is considered
an extension, the use of which makes an implementation nonconforming. See also
B.2.3.7.

Epoch: Historically, the origin of UNIX system time was referred to as "00:00:00 606 GMT, January 1, 1970." Greenwich Mean Time is actually not a term ack-607 nowledged by the international standards community; therefore, this term, 608 *Epoch*, is used to abbreviate the reference to the actual standard, Coordinated 609 Universal Time. The concept of leap seconds is added for precision; at the time 610 POSIX.1 was published, 14 leap seconds had been added since January 1, 1970. 611 These 14 seconds are ignored to provide an easy and compatible method of com-612 puting time differences. 613

Most systems' notion of "time" is that of a continuously increasing value, so this value should increase even during leap seconds. However, not only do most systems not keep track of leap seconds, but most systems are probably not synchronized to any standard time reference. Therefore, it is inappropriate to require that a time represented as seconds since the Epoch precisely represent the number of seconds between the referenced time and the Epoch.

- It is sufficient to require that applications be allowed to treat this time as if it represented the number of seconds between the referenced time and the Epoch. It is the responsibility of the vendor of the system, and the administrator of the system, to ensure that this value represents the number of seconds between the referenced time and the Epoch as closely as necessary for the application being run on that system.
- It is important that the interpretation of time names and seconds since the Epoch 626 values be consistent across conforming systems. That is, it is important that all 627 conforming systems interpret "536457599 seconds since the Epoch" as 59 628 seconds, 59 minutes, 23 hours 31 December 1986, regardless of the accuracy of 629 the system's idea of the current time. The expression is given to assure a con-630 sistent interpretation, not to attempt to specify the calendar. The relationship 631 between tm yday and the day of week, day of month, and month is presumed to 632 be specified elsewhere and is not given in POSIX.1. 633
- Consistent interpretation of *seconds since the Epoch* can be critical to certain types of distributed applications that rely on such timestamps to synchronize events. The accrual of leap seconds in a time standard is not predictable. The number of leap seconds since the Epoch will likely increase. POSIX.1 is more concerned about the synchronization of time between applications of astronomically short duration. These concerns are expected to become more critical in the future.
- Note that *tm_yday* is zero-based, not one-based, so the day number in the example above is 364. Note also that the division is an integer division (discarding remainder) as in the C language.
- Note also that in Section 8, the meaning of gmtime(), localtime(), and mktime() is specified in terms of this expression. However, the C Standard {2} computes tm_yday from tm_mday , tm_mon , and tm_year in mktime(). Because it is stated as a (bidirectional) relationship, not a function, and because the conversion between month-day-year and day-of-year dates is presumed well known and is also a relationship, this is not a problem.

Note that the expression given will fail after the year 2099. Since the issue of *time_t* overflowing a 32-bit integer occurs well before that time, both of these will have to be addressed in revisions to POSIX.1.

- **FIFO special file:** See *pipe* in B.2.2.2.
- 653 **file:** It is permissible for an implementation-defined file type to be nonreadable 654 or nonwritable.

file classes: These classes correspond to the historical sets of permission bits.
The classes are general to allow implementations flexibility in expanding the
access mechanism for more stringent security environments. Note that a process
is in one and only one class, so there is no ambiguity.

659 **filename:** At the present time, the primary responsibility for truncating 660 filenames containing multibyte characters must reside with the application. 661 Some industry groups involved in internationalization believe that in the future 662 the responsibility must reside with the kernel. For the moment, a clearer under-663 standing of the implications of making the kernel responsible for truncation of 664 multibyte file names is needed.

- 665 Character level truncation was not adopted because there is no support in 666 POSIX.1 that advises how the kernel distinguishes between single and multibyte 667 characters. Until that time, it must be incumbent upon application writers to 668 determine where multibyte characters must be truncated.
- **file system:** Historically the meaning of this term has been overloaded with two meanings: that of the complete file hierarchy and that of a mountable subset of that hierarchy; i.e., a mounted file system. POSIX.1 uses the term *file system* in the second sense, except that it is limited to the scope of a process (and a process's root directory). This usage also clarifies the domain in which a file serial number is unique.
- 675 ***group file:** Implementation defined; see B.9.
- *historical implementations: This refers to previously existing implementations of programming interfaces and operating systems that are related to the
 interface specified by POSIX.1. See also "Minimal Changes to Historical Implementations" in the Introduction.
- ***hosted implementation:** This refers to a POSIX.1 implementation that is 680 accomplished through interfaces from the POSIX.1 services to some alternate form 681 of operating system kernel services. Note that the line between a hosted imple-682 mentation and a native implementation is blurred, since most implementations 683 will provide some services directly from the kernel and others through some 684 indirect path. [For example, fopen() might use open(); or mkfifo() might use 685 *mknod().*] There is no necessary relationship between the type of implementation 686 and its correctness, performance, and/or reliability. 687
- *implementation: The term is generally used instead of its synonym, system,
 to emphasize the consequences of decisions to be made by system implementors.
 Perhaps if no options or extensions to POSIX.1 were allowed, this usage would not
 have occurred.
- The term specific implementation is sometimes used as a synonym for implementation. This should not be interpreted too narrowly; both terms can represent a relatively broad group of systems. For example, a hardware vendor could market a very wide selection of systems that all used the same instruction set, with some systems desktop models and others large multiuser minicomputers. This wide

range would probably share a common POSIX.1 operating system, allowing an
 application compiled for one to be used on any of the others; this is a [specific]
 implementation.

However, that wide range of machines probably has some differences between the models. Some may have different clock rates, different file systems, different resource limits, different network connections, etc., depending on their sizes or intended usages. Even on two identical machines, the system administrators may configure them differently. Each of these different systems is known by the term *a specific instance of a specific implementation*. This term is only used in the portions of POSIX.1 dealing with run-time queries: sysconf() and pathconf().

*incomplete pathname: Absolute pathname has been adequately defined.

job control: In order to understand the job-control facilities in POSIX.1 it is useful to understand how they are used by a job-control-cognizant shell to create the user interface effect of job control.

711 While the job-control facilities supplied by POSIX.1 can, in theory, support dif-712 ferent types of interactive job-control interfaces supplied by different types of 713 shells, there is historically one particular interface that is most common (provided 714 by BSD C Shell). This discussion describes that interface as a means of illustrat-715 ing how the POSIX.1 job-control facilities can be used.

- Job control allows users to selectively stop (suspend) the execution of processes and continue (resume) their execution at a later point. The user typically employs this facility via the interactive interface jointly supplied by the terminal I/O driver and a command interpreter (shell).
- The user can launch jobs (command pipelines) in either the foreground or background. When launched in the foreground, the shell waits for the job to complete before prompting for additional commands. When launched in the background, the shell does not wait, but immediately prompts for new commands.
- If the user launches a job in the foreground and subsequently regrets this, the user can type the suspend character (typically set to control-Z), which causes the foreground job to stop and the shell to begin prompting for new commands. The stopped job can be continued by the user (via special shell commands) either as a foreground job or as a background job. Background jobs can also be moved into the foreground via shell commands.
- If a background job attempts to access the login terminal (controlling terminal), it is stopped by the terminal driver and the shell is notified, which, in turn, notifies the user. [Terminal access includes *read()* and certain terminal control functions and conditionally includes *write()*.] The user can continue the stopped job in the foreground, thus allowing the terminal access to succeed in an orderly fashion. After the terminal access succeeds, the user can optionally move the job into the background via the suspend character and shell commands.
- 737 Implementing Job Control Shells
- The interactive interface described previously can be accomplished using the POSIX.1 job-control facilities in the following way.
- The key feature necessary to provide job control is a way to group processes into jobs. This grouping is necessary in order to direct signals to a single job and also

to identify which job is in the foreground. (There is at most one job that is in theforeground on any controlling terminal at a time.)

The concept of process groups is used to provide this grouping. The shell places 744 each job in a separate process group via the setpgid() function. To do this, the 745 setpgid() function is invoked by the shell for each process in the job. It is actually 746 useful to invoke *setpgid()* twice for each process: once in the child process, after 747 calling fork() to create the process, but before calling one of the exec functions to 748 begin execution of the program, and once in the parent shell process, after calling 749 fork() to create the child. The redundant invocation avoids a race condition by 750 ensuring that the child process is placed into the new process group before either 751 the parent or the child relies on this being the case. The process group ID for the 752 job is selected by the shell to be equal to the process ID of one of the processes in 753 the job. Some shells choose to make one process in the job be the parent of the 754 other processes in the job (if any). Other shells (e.g., the C Shell) choose to make 755 themselves the parent of all processes in the pipeline (job). In order to support 756 this latter case, the *setpgid()* function accepts a process group ID parameter since 757 the correct process group ID cannot be inherited from the shell. The shell itself is 758 considered to be a job and is the sole process in its own process group. 759

- The shell also controls which job is currently in the foreground. A foreground and 760 background job differ in two ways: the shell waits for a foreground command to 761 complete (or stop) before continuing to read new commands, and the terminal I/O 762 driver inhibits terminal access by background jobs (causing the processes to stop). 763 Thus, the shell must work cooperatively with the terminal I/O driver and have a 764 common understanding of which job is currently in the foreground. It is the user 765 who decides which command should be currently in the foreground, and the user 766 informs the shell via shell commands. The shell, in turn, informs the terminal I/O 767 driver via the *tcsetpgrp()* function. This indicates to the terminal I/O driver the 768 process group ID of the foreground process group (job). When the current fore-769 ground job either stops or terminates, the shell places itself in the foreground via 770 tcsetpgrp() before prompting for additional commands. Note that when a job is 771 created the new process group begins as a background process group. It requires 772 an explicit act of the shell via *tcsetpgrp()* to move a process group (job) into the 773 foreground. 774
- When a process in a job stops or terminates, its parent (e.g., the shell) receives synchronous notification by calling the *waitpid()* function with the WUNTRACED flag set. Asynchronous notification is also provided when the parent establishes a signal handler for SIGCHLD and does not specify the SA_NOCLDSTOP flag. Usually all processes in a job stop as a unit since the terminal I/O driver always sends job-control stop signals to all processes in the process group.
- To continue a stopped job, the shell sends the SIGCONT signal to the process group of the job. In addition, if the job is being continued in the foreground, the shell invokes *tcsetpgrp()* to place the job in the foreground before sending SIGCONT. Otherwise, the shell leaves itself in the foreground and reads additional commands.
- There is additional flexibility in the POSIX.1 job-control facilities that allows deviations from the typical interface. Clearing the TOSTOP terminal flag (see 7.1.2.5) allows background jobs to perform *write()* functions without stopping. The same effect can be achieved on a per-process basis by having a process set the signal action for SIGTTOU to SIG_IGN.

- Note that the terms job and process group can be used interchangeably. A login 791 session that is not using the job control facilities can be thought of as a large col-792 lection of processes that are all in the same job (process group). Such a login ses-793 sion may have a partial distinction between foreground and background 794 processes; that is, the shell may choose to wait for some processes before continu-795 ing to read new commands and may not wait for other processes. However, the 796 terminal I/O driver will consider all these processes to be in the foreground since 797 they are all members of the same process group. 798
- In addition to the basic job-control operations already mentioned, a job-controlcognizant shell needs to perform the following actions:
- When a foreground (not background) job stops, the shell must sample and remember the current terminal settings so that it can restore them later when it continues the stopped job in the foreground [via the *tcgetattr()* and *tcsetattr()* functions].
- Because a shell itself can be spawned from a shell, it must take special action to ensure that subshells interact well with their parent shells.
- A subshell can be spawned to perform an interactive function (prompting the ter-807 minal for commands) or a noninteractive function (reading commands from a file). 808 When operating noninteractively, the job-control shell will refrain from perform-809 ing the job-control specific actions described above. It will behave as a shell that 810 does not support job control. For example, all jobs will be left in the same process 811 group as the shell, which itself remains in the process group established for it by 812 its parent. This allows the shell and its children to be treated as a single job by a 813 parent shell, and they can be affected as a unit by terminal keyboard signals. 814
- 815 An interactive subshell can be spawned from another job-control-cognizant shell in either the foreground or background. (For example, from the C Shell, the user 816 can execute the command, csh &.) Before the subshell activates job control by 817 calling setpgid() to place itself in its own process group and tcsetpgrp() to place its 818 new process group in the foreground, it needs to ensure that it has already been 819 placed in the foreground by its parent. (Otherwise, there could be multiple job-820 control shells that simultaneously attempt to control mediation of the terminal.) 821 To determine this, the shell retrieves its own process group via getpgrp() and the 822 process group of the current foreground job via tcgetpgrp(). If these are not equal, 823 the shell sends SIGTTIN to its own process group, causing itself to stop. When 824 continued later by its parent, the shell repeats the process-group check. When 825 the process groups finally match, the shell is in the foreground and it can proceed 826 to take control. After this point, the shell ignores all the job-control stop signals 827 so that it does not inadvertently stop itself. 828
- 829 Implementing Job Control Applications
- Most applications do not need to be aware of job-control signals and operations; the intuitively correct behavior happens by default. However, sometimes an application can inadvertently interfere with normal job-control processing, or an application may choose to overtly effect job control in cooperation with normal shell procedures.
- An application can inadvertently subvert job-control processing by "blindly" altering the handling of signals. A common application error is to learn how many

signals the system supports and to ignore or catch them all. Such an application 837 makes the assumption that it does not know what this signal is, but knows the 838 right handling action for it. The system may initialize the handling of job-control 839 stop signals so that they are being ignored. This allows shells that do not support 840 job control to inherit and propagate these settings and hence to be immune to stop 841 signals. A job-control shell will set the handling to the default action and pro-842 pagate this, allowing processes to stop. In doing so, the job-control shell is taking 843 responsibility for restarting the stopped applications. If an application wishes to 844 catch the stop signals itself, it should first determine their inherited handling 845 states. If a stop signal is being ignored, the application should continue to ignore 846 it. This is directly analogous to the recommended handling of SIGINT described in 847 the UNIX Programmer's Manual (B41). 848

- If an application is reading the terminal and has disabled the interpretation of 849 special characters (by clearing the ISIG flag), the terminal I/O driver will not send 850 SIGTSTP when the suspend character is typed. Such an application can simulate 851 the effect of the suspend character by recognizing it and sending SIGTSTP to its 852 process group as the terminal driver would have done. Note that the signal is 853 sent to the process group, not just to the application itself; this ensures that other 854 processes in the job also stop. (Note also that other processes in the job could be 855 children, siblings, or even ancestors.) Applications should not assume that the 856 suspend character is control-Z (or any particular value); they should retrieve the 857 current setting at startup. 858
- 859 Implementing Job Control Systems

The intent in adding 4.2BSD-style job control functionality was to adopt the necessary 4.2BSD programmatic interface with only minimal changes to resolve syntactic or semantic conflicts with System V or to close recognized security holes. The goal was to maximize the ease of providing both conforming implementations and Conforming POSIX.1 Applications.

 865
 Discussions of the changes can be found in the clauses that discuss the specific

 866
 interfaces. See B.3.2.1, B.3.2.2, B.3.3.1.1, B.3.3.2, B.3.3.4, B.4.3.1, B.4.3.3,

 867
 B.7.1.1.4, and B.7.2.4.

It is only useful for a process to be affected by job-control signals if it is the des-868 cendant of a job-control shell. Otherwise, there will be nothing that continues the 869 stopped process. Because a job-control shell is allowed, but not required, by 870 POSIX.1, an implementation must provide a mechanism that shields processes 871 from job-control signals when there is no job-control shell. The usual method is 872 for the system initialization process (typically called init), which is the ancestor 873 of all processes, to launch its children with the signal handling action set to 874 SIG_IGN for the signals SIGTSTP, SIGTTIN, and SIGTTOU. Thus, all login shells 875 start with these signals ignored. If the shell is not job-control cognizant, then it 876 should not alter this setting and all its descendants should inherit the same 877 ignored settings. At the point where a job-control shell is launched, it resets the 878 signal handling action for these signals to be SIG_DFL for its children and (by 879 inheritance) their descendants. Also, shells that are not job-control cognizant will 880 not alter the process group of their descendants or of their controlling terminal; 881 this has the effect of making all processes be in the foreground (assuming the 882 shell is in the foreground). While this approach is valid, POSIX.1 added the con-883 cept of orphaned process groups to provide a more robust solution to this problem. 884

All processes in a session managed by a shell that is not job-control cognizant are in an orphaned process group and are protected from stopping.

POSIX.1 does not specify how controlling terminal access is affected by a user log-887 ging out (that is, by a controlling process terminating). 4.2BSD uses the 888 *vhangup()* function to prevent any access to the controlling terminal through file 889 descriptors opened prior to logout. System V does not prevent controlling termi-890 nal access through file descriptors opened prior to logout (except for the case of 891 the special file, /dev/tty). Some implementations choose to make processes 892 immune from job control after logout (that is, such processes are always treated 893 as if in the foreground); other implementations continue to enforce 894 foreground/background checks after logout. Therefore, a Conforming POSIX.1 895 Application should not attempt to access the controlling terminal after logout 896 since such access is unreliable. If an implementation chooses to deny access to a 897 controlling terminal after its controlling process exits, POSIX.1 requires a certain 898 type of behavior (see 7.1.1.3). 899

- 900 ***kernel:** See system call.
- 901 ***library routine:** See system call.
- 902 ***logical device:** Implementation defined.
- *mount point: The directory on which a mounted file system is mounted. This term, like mount() and umount(), was not included because it was implementation defined.
- 906 ***mounted file system:** See *file system*.

*native implementation: This refers to an implementation of POSIX.1 that
interfaces directly to an operating-system kernel. See also *hosted implementation*and *cooperating implementation*. A similar concept is a native UNIX system,
which would be a kernel derived from one of the original UNIX system products.

open file description: An open file description, as it is currently named,
 describes how a file is being accessed. What is currently called a *file descriptor* is
 actually just an identifier or "handle"; it does not actually describe anything.

- 914 The following alternate names were discussed:
- 915 For open file description:
- open instance, file access description, open file information, and file
 access information.
- 918 For file descriptor:
- *file handle, file number* [c.f., *fileno()*]. Some historical implementations use the term *file table entry*.

orphaned process group: Historical implementations have a concept of an 921 orphaned process, which is a process whose parent process has exited. When job 922 control is in use, it is necessary to prevent processes from being stopped in 923 response to interactions with the terminal after they no longer are controlled by a 924 job-control-cognizant program. Because signals generated by the terminal are 925 sent to a process group and not to individual processes, and because a signal may 926 be provoked by a process that is not orphaned, but sent to another process that is 927 orphaned, it is necessary to define an orphaned process group. The definition 928

- assumes that a process group will be manipulated as a group and that the jobcontrol-cognizant process controlling the group is outside of the group and is the parent of at least one process in the group [so that state changes may be reported via *waitpid()*]. Therefore, a group is considered to be controlled as long as at least one process in the group has a parent that is outside of the process group, but within the session.
- This definition of orphaned process groups ensures that a session leader's process group is always considered to be orphaned, and thus it is prevented from stopping in response to terminal signals.
- 938 ***passwd file:** Implementation defined; see B.9.

parent directory: There may be more than one directory entry pointing to a given directory in some implementations. The wording here identifies that exactly one of those is the parent directory. In 2.3.6, *dot-dot* is identified as the way that the unique directory is identified. (That is, the parent directory is the one to which dot-dot points.) In the case of a remote file system, if the same file system is mounted several times, it would appear as if they were distinct file systems (with interesting synchronization properties).

- pipe: It proved convenient to define a *pipe* as a special case of a *FIFO* even
 though historically the latter was not introduced until System III and does not
 exist at all in 4.3BSD.
- portable filename character set: The encoding of this character set is not
 specified—specifically, ASCII is not required. But the implementation must provide a unique character code for each of the printable graphics specified by
 POSIX.1. See also B.2.3.5.
- Situations where characters beyond the portable filename character set (or histor-953 ically ASCII or ISO/IEC 646 {1}) would be used (in a context where the portable 954 filename character set or ISO/IEC 646 {1} is required by POSIX.1) are expected to 955 be common. Although such a situation renders the use technically noncompliant, 956 mutual agreement among the users of an extended character set will make such 957 use portable between those users. Such a mutual agreement could be formalized 958 as an optional extension to POSIX.1. (Making it required would eliminate too 959 many possible systems, as even those systems using ISO/IEC 646 {1} as a base 960 character set extend their character sets for Western Europe and the rest of the 961 world in different ways.) 962
- Nothing in POSIX.1 is intended to preclude the use of extended characters where interchange is not required or where mutual agreement is obtained. It has been suggested that in several places "should" be used instead of "shall." Because (in the worst case) use of any character beyond the portable filename character set would render the program or data not portable to all possible systems, no extensions are permitted in this context.
- **regular file:** POSIX.1 does not intend to preclude the addition of structuring data (e.g., record lengths) in the file, as long as such data is not visible to an application that uses the features described in POSIX.1.
- root directory: This definition permits the operation of chroot(), even though
 that function is not in POSIX.1. See also file hierarchy.

ł

- 974 ***root file system:** Implementation defined.
- *root of a file system: Implementation defined. See mount point.
- **seconds since the Epoch:** The formula here is not precisely correct for leap centuries. See the discussion for *Epoch* for further details.
- signal: The definition implies a double meaning for the term. Although a signal
 is an event, common usage implies that a signal is an identifier of the class of
 event.
- *system call: The distinction between a system call and a library routine is an
 implementation detail that may differ between implementations and has thus
 been excluded from POSIX.1. See "Interface, Not Implementation" in the Introduction.
- *super-user: This concept, with great historical significance to UNIX system
 users, has been replaced with the notion of *appropriate privileges*.
- 987 B.2.2.3 Abbreviations
- 988 There is no additional rationale provided for this subclause.
- 989 B.2.3 General Concepts
- B.2.3.1 extended security controls: Allowing an implementation to define 990 extended security controls enables the use of POSIX.1 in environments that 991 require different or more rigorous security than that provided in POSIX.1. Exten-992 sions are allowed in two areas: privilege and file access permissions. The seman-993 tics of these areas have been defined to permit extensions with reasonable, but 994 not exact, compatibility with all existing practices. For example, the elimination 995 of the super-user definition precludes identifying a process as privileged or not by 996 virtue of its effective user ID. 997

B.2.3.2 file access permissions: A process should not try to anticipate the 998 result of an attempt to access data by a priori use of these rules. Rather, it should 999 make the attempt to access data and examine the return value (and possibly 1000 errno as well), or use access(). An implementation may include other security 1001 mechanisms in addition to those specified in POSIX.1, and an access attempt may 1002 fail because of those additional mechanisms, even though it would succeed accord-1003 ing to the rules given in this subclause. (For example, the user's security level 1004 might be lower than that of the object of the access attempt.) The optional supple-1005 mentary group IDs provide another reason for a process to not attempt to antici-1006 pate the result of an access attempt. 1007

1008 **B.2.3.3 file hierarchy:** Though the file hierarchy is commonly regarded to be a 1009 tree, POSIX.1 does not define it as such for three reasons:

- 1010 (1) Links may join branches.
- 1011 (2) In some network implementations, there may be no single absolute root 1012 directory. See *pathname resolution*.

- (3) With symbolic links (found in 4.3BSD), the file system need not be a tree
 or even a directed acyclic graph.
- 1015 **B.2.3.4 file permissions:** Examples of implementation-defined constraints that 1016 may deny access are mandatory labels and access control lists.

1017 **B.2.3.5 filename portability:** Historically, certain filenames have been 1018 reserved. This list includes core, /etc/passwd, etc. Portable applications 1019 should avoid these.

- Most historical implementations prohibit case folding in filenames; i.e., treating | upper- and lowercase alphabetic characters as identical. However, some consider case folding desirable:
- 1023 For user convenience
- For ease of implementation of the POSIX.1 interface as a hosted system on
 some popular operating systems, which is compatible with the goal of mak ing the POSIX.1 interface broadly implementable (see "Broadly Implement able" in the Introduction)

Variants such as maintaining case distinctions in filenames, but ignoring them in
 comparisons, have been suggested. Methods of allowing escaped characters of the
 case opposite the default have been proposed

- 1031 Many reasons have been expressed for not allowing case folding, including:
- 1032 (1) No solid evidence has been produced as to whether case sensitivity or 1033 case insensitivity is more convenient for users.
- 1034 (2) Making case insensitivity a POSIX.1 implementation option would be 1035 worse than either having it or not having it, because
- 1036 (a) More confusion would be caused among users.
- 1037 (b) Application developers would have to account for both cases in their 1038 code.
- 1039(c)POSIX.1 implementors would still have other problems with native1040file systems, such as short or otherwise constrained filenames or1041pathnames, and the lack of hierarchical directory structure.
- (3) Case folding is not easily defined in many European languages, both
 because many of them use characters outside the USASCII alphabetic set,
 and because
- 1045(a) In Spanish, the digraph 11 is considered to be a single letter, the
capitalized form of which may be either L1 or LL, depending on con-
text.
- 1048(b) In French, the capitalized form of a letter with an accent may or1049may not retain the accent depending on the country in which it is1050written.
- 1051(c) In German, the sharp ess may be represented as a single character1052resembling a Greek beta (β) in lowercase, but as the digraph ss in1053uppercase.

- 1054 1055
- (d) In Greek, there are several lowercase forms of some letters; the one to use depends on its position in the word. Arabic has similar rules.
- Many East Asian languages, including Japanese, Chinese, and Korean,
 do not distinguish case and are sometimes encoded in character sets that
 use more than one byte per character.
- (5) Multiple character codes may be used on the same machine simultaneously. There are several ISO character sets for European alphabets. In Japan, several Japanese character codes are commonly used together, sometimes even in filenames; this is evidently also the case in China. To handle case insensitivity, the kernel would have to at least be able to distinguish for which character sets the concept made sense.
- 1065 (6) The file system implementation historically deals only with bytes, not 1066 with characters, except for slash and the null byte.
- (7) The purpose of POSIX.1 is to standardize the common, existing definition
 (see "Application Oriented" in the Introduction) of the UNIX system pro gramming interface, not to change it. Mandating case insensitivity
 would make all historical implementations nonstandard.
- 1071 (8) Not only the interface, but also application programs would need to
 1072 change, counter to the purpose of having minimal changes to existing
 1073 application code.
- At least one of the original developers of the UNIX system has expressed objection in the strongest terms to either requiring case insensitivity or making it an option, mostly on the basis that POSIX.1 should not hinder portability of application programs across related implementations in order to allow compatibility with unrelated operating systems.
- 1079 Two proposals were entertained regarding case folding in filenames:
- 1080 Remove all wording that previously permitted case folding.
- 1081Rationale: Case folding is inconsistent with portable filename character set1082definition and filename definition (all characters except slash and null). No1083known implementations allowing all characters except slash and null also1084do case folding.
- 1085 Change "though this practice is not recommended:" to "although this practice is strongly discouraged."
- 1087 Rationale: If case folding must be included in POSIX.1, the wording should 1088 be stronger to discourage the practice.

1089 The consensus selected the first proposal. Otherwise, a portable application 1090 would have to assume that case folding would occur when it was not wanted, but 1091 that it would not occur when it was wanted.

B.2.3.6 file times update: This subclause reflects the actions of historical | implementations. The times are not updated immediately, but are only marked for update by the functions. An implementation may update these times immediately.

1096 The accuracy of the time update values is intentionally left unspecified so that 1097 systems can control the bandwidth of a possible covert channel.

1098 The wording was carefully chosen to make it clear that there is no requirement 1099 that the conformance document contain information that might incidentally affect 1100 file update times. Any function that performs pathname resolution might update 1101 several st_atime fields. Functions such as getpwnam() and getgrnam() might 1102 update the st_atime field of some specific file or files. It is intended that these are 1103 not required to be documented in the conformance document, but they should 1104 appear in the system documentation.

B.2.3.7 pathname resolution: What the filename dot-dot refers to relative to the root directory is implementation defined. In Version 7 it refers to the root directory itself; this is the behavior mentioned in the standard. In some networked systems the construction /../hostname/ is used to refer to the root directory of another host, and POSIX.1 permits this behavior.

Other networked systems use the construct //hostname for the same purpose; i.e., a double initial slash is used. There is a potential problem with existing applications that create full pathnames by taking a trunk and a relative pathname and making them into a single string separated by /, because they can accidentally create networked pathnames when the trunk is /. This practice is not prohibited because such applications can be made to conform by simply changing to use // as a separator instead of /:

- (1) If the trunk is /, the full path name will begin with /// (the initial / and the separator //). This is the same as /, which is what is desired. (This is the general case of making a relative pathname into an absolute one by prefixing with /// instead of /.)
- (2) If the trunk is /A, the result is /A//...; since nonleading sequences of two or more slashes are treated as a single slash, this is equivalent to the desired /A/....
- (3) If the trunk is //A, the implementation-defined semantics will apply.
 (The multiple slash rule would apply.)

Application developers should avoid generating pathnames that start with "//". Implementations are strongly encouraged to avoid using this special interpretation since a number of applications currently do not follow this practice and may inadvertently generate "//...".

1130 The term root directory is only defined in POSIX.1 relative to the process. In some 1131 implementations, there may be no absolute root directory. The initialization of 1132 the root directory of a process is implementation defined.

1133 B.2.4 Error Numbers

1134 The definition of *errno* in POSIX.1 is stricter than that in the C Standard {2}. The 1135 C Standard {2} merely requires that it be an assignable *lvalue*. The POSIX.1 1136 *extern int errno* meets that requirement and supports historical usage as well.

l

1137 Checking the value of *errno* alone is not sufficient to determine the existence or 1138 type of an error, since it is not required that a successful function call clear *errno*. 1139 The variable *errno* should only be examined when the return value of a function 1140 indicates that the value of *errno* is meaningful. In that case, the function is 1141 required to set the variable to something other than zero.

A successful function call may set the value of *errno* to zero, or to any other value 1142 (except where specifically prohibited; see B.5.4.1). But it is meaningless to do so, 1143 since the value of *errno* is undefined except when the description of a function 1144 explicitly states that it is set, and no function description states that it should be 1145 set on a successful call. Most functions in most implementations do not change 1146 errno on successful completion. Exceptions are *isatty()* and *ptrace()*. The latter is 1147 not in POSIX.1, but is widely implemented and clears errno when called. The 1148 value of *errno* is not defined unless all signal handlers that use functions that 1149 could change errno save and restore it. 1150

POSIX.1 requires (in the Errors subclauses of function descriptions) certain error values to be set in certain conditions because many existing applications depend on them. Some error numbers, such as [EFAULT], are entirely implementation defined and are noted as such in their description in 2.4. This subclause otherwise allows wide latitude to the implementation in handling error reporting.

1156 Some of the Errors clauses in POSIX.1 have two subclauses. The first:

- "If any of the following conditions occur, the *foo()* function shall
 return -1 and set *errno* to the corresponding value:"
- 1159 could be called the "mandatory" subclause. The second:

"For each of the following conditions, when the condition is detected,
the *foo()* function shall return -1 and set *errno* to the corresponding
value:"

could be informally known as the "optional" subclause. This latter subclause has evolved in meaning over time. In early drafts, it was only used for error conditions that could not be detected by certain hardware configurations, such as the [EFAULT] error, as described below. The subclause recently has also added conditions associated with optional system behavior, such as job control errors. Attempting to infer the quality of an implementation based on whether it detects such conditions is not useful.

Following each one-word symbolic name for an error, there is a one-line tag, which is followed by a description of the error. The one-line tag is merely a mnemonic or historical referent and is not part of the specification of the error. Many programs print these tags on the standard error stream [often by using the C Standard {2} perror() function] when the corresponding errors are detected, but POSIX.1 does not require this action.

^{1176[}EFAULT]Most historical implementations do not catch an error and set1177errno when an invalid address is given to the functions wait(),1178time(), or times(). Some implementations cannot reliably detect1179an invalid address. And most systems that detect invalid1180addresses will do so only for a system call, not for a library1181routine.

- 1182[EINTR]POSIX.1 prohibits conforming implementations from restarting1183interrupted system calls. However, it does not require that1184[EINTR] be returned when another legitimate value may be1185substituted; e.g., a partial transfer count when read() or write()1186are interrupted. This is only given when the signal catching1187function returns normally as opposed to returns by mechanisms1188like longjmp() or siglongjmp().
- 1189 [ENOMEM] The term *main memory* is not used in POSIX.1 because it is 1190 implementation defined.
- 1191[ENOTTY]The symbolic name for this error is derived from a time when1192device control was done by *ioctl()* and that operation was only1193permitted on a terminal interface. The term "TTY" is derived1194from *teletypewriter*, the devices to which this error originally1195applied.
- 1196[EPIPE]This condition normally generates the signal SIGPIPE; the error1197is returned if the signal does not terminate the process.
- [EROFS] In historical implementations, attempting to unlink() or 1198 rmdir() a mount point would generate an [EBUSY] error. An 1199 implementation could be envisioned where such an operation 1200 could be performed without error. In this case, if either the 1201 directory entry or the actual data structures reside on a read-1202 only file system, [EROFS] is the appropriate error to generate. 1203 1204 (For example, changing the link count of a file on a read-only file system could not be done, as is required by unlink(), and 1205 thus an error should be reported.) 1206

Two error numbers, [EDOM] and [ERANGE], were added to this subclause primarily for consistency with the C Standard {2}.

1209 **B.2.5 Primitive System Data Types**

The requirement that additional types defined in this subclause end in "_t" was 1210 prompted by the problem of namespace pollution (see B.2.7.2). It is difficult to 1211 define a type (where that type is not one defined by POSIX.1) in one header file 1212 and use it in another without adding symbols to the namespace of the program. 1213 To allow implementors to provide their own types, all POSIX.1 conforming applica-1214 tions are required to avoid symbols ending in "_t", which permits the implementor 1215 to provide additional types. Because a major use of types is in the definition of 1216 structure members, which can (and in many cases must) be added to the struc-1217 tures defined in POSIX.1, the need for additional types is compelling. 1218

1219 The types such as *ushort* and *ulong*, which are in common usage, are not defined 1220 in POSIX.1 (although *ushort_t* would be permitted as an extension). They can be 1221 added to <sys/types.h> using a feature test macro (see 2.7.2). A suggested 1222 symbol for these is _SYSIII. Similarly, the types like *u_short* would probably be 1223 best controlled by _BSD.

1224 Some of these symbols may appear in other headers; see 2.7.

dev_t

- 12251226
- 1227 1228

1229

locality considerations of networked systems. This type must be arithmetic. Earlier drafts allowed this to be nonarithmetic (such as a structure) and provided a *samefile()* function for comparison.

This type may be made large enough to accommodate host-

1230gid_tSome implementations had separated gid_t from uid_t before1231POSIX.1 was completed. It would be difficult for them to1232coalesce them when it was unnecessary. Additionally, it is1233quite possible that user IDs might be different than group IDs1234because the user ID might wish to span a heterogeneous net-1235work, where the group ID might not.

1236For current implementations, the cost of having a separate1237 gid_t will be only lexical.

1238mode_tThis type was chosen so that implementations could choose the1239appropriate integral type, and for compatibility with the1240C Standard {2}. 4.3BSD uses unsigned short and the SVID uses1241ushort, which is the same. Historically, only the low-order six-1242teen bits are significant.

1243nlink_tThis type was introduced in place of short for st_nlink (see12445.6.1) in response to an objection that short was too small.

1245off_tThis type is used only in lseek(), fcntl(), and <sys/stat.h>.1246Many implementations would have difficulties if it were defined1247as anything other than long. Requiring an integral type limits1248the capabilities of lseek() to four gigabytes. See the description1249of lread() in B.6.4. Also, the C Standard {2} supplies routines1250that use larger types: see fgetpos() and fsetpos() in B.6.5.3.

1251 pid_t The inclusion of this symbol was controversial because it is tied1252to the issue of the representation of a process ID as a number.1253From the point of view of a portable application, process IDs1254should be "magic cookies"²⁾ that are produced by calls such as1255fork(), used by calls such as waitpid() or kill(), and not other-1256wise analyzed (except that the sign is used as a flag for certain1257operations).

1258The concept of a {PID_MAX} value interacted with this in early1259drafts. Treating process IDs as an opaque type both removes1260the requirement for {PID_MAX} and allows systems to be more1261flexible in providing process IDs that span a large range of1262values, or a small one.

^{1263 2)} An historical term meaning: "An opaque object, or token, of determinate size, whose significance
1264 is known only to the entity which created it. An entity receiving such a token from the
1265 generating entity may only make such use of the 'cookie' as is defined and permitted by the
1266 supplying entity."

1286

1287

1288

1289 1290

1298

1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

- 1267Since the values in uid_t, gid_t, and pid_t will be numbers gen-1268erally, and potentially both large in magnitude and sparse,1269applications that are based on arrays of objects of this type are1270unlikely to be fully portable in any case. Solutions that treat1271them as magic cookies will be portable.
- 1272{CHILD_MAX} precludes the possibility of a "toy implementa-1273tion," where there would only be one process.
- ssize t This is intended to be a signed analog of *size_t*. The wording is 1274 such that an implementation may either choose to use a longer 1275 type or simply to use the signed version of the type that under-1276 1277 lies *size_t*. All functions that return *ssize_t* [*read()* and *write()*] describe as "implementation defined" the result of an input 1278 exceeding {SSIZE_MAX}. It is recognized that some implemen-1279 tations might have ints that are smaller than size_t. A port-1280 able application would be constrained not to perform I/O in 1281 pieces larger than {SSIZE_MAX}, but a portable application 1282 using extensions would be able to use the full range if the 1283 implementation provided an extended range, while still having 1284 a single type-compatible interface. 1285
 - The symbols *size_t* and *ssize_t* are also required in <unistd.h> to minimize the changes needed for calls to *read()* and *write()*. Implementors are reminded that it must be possible to include both <sys/types.h> and <unistd.h> in the same program (in either order) without error.
- 1291uid_tBefore the addition of this type, the data types used to1292represent these values varied throughout early drafts. The1293<\$ys/stat.h> header defined these values as type short, the1294<passwd.h> file (now <pwd.h> and <grp.h>) used an int,1295and getuid() returned an int. In response to a strong objection1296to the inconsistent definitions, all the types to were switched to1297uid_t.
 - In practice, those historical implementations that use varying types of this sort can typedef *uid_t* to *short* with no serious consequences.

The problem associated with this change concerns object compatibility after structure size changes. Since most implementations will define uid_t as a short, the only substantive change will be a reduction in the size of the *passwd* structure. Consequently, implementations with an overriding concern for object compatibility can pad the structure back to its current size. For that reason, this problem was not considered critical enough to warrant the addition of a separate type to POSIX.1.

1309The types uid_t and gid_t are magic cookies. There is no1310{UID_MAX} defined by POSIX.1, and no structure imposed on1311 uid_t and gid_t other than that they be positive arithmetic1312types. (In fact, they could be unsigned char.) There is no max-1313imum or minimum specified for the number of distinct user or1314group IDs.

I

1315 **B.2.6 Environment Description**

The variable *environ* is not intended to be declared in any header, but rather to be declared by the user for accessing the array of strings that is the environment. This is the traditional usage of the symbol. Putting it into a header could break some programs that use the symbol for their own purposes.

- 1320LC_*The description of the environment variable names starting1321with the characters "LC_" acknowledges the fact that the inter-1322faces presented in the current version of POSIX.1 are not com-1323plete and may be extended as new international functionality is1324required. In the C Standard {2}, names preceded by "LC_" are1325reserved in the name space for future categories.
- 1326To avoid name clashes, new categories and environment vari-1327ables are divided into two classifications: implementation1328independent and implementation dependent.
- 1329Implementation-independent names will have the following1330format:

LC_NAME

- where *NAME* is the name of the new category and environment variable. Capital letters must be used for implementation-independent names.
- 1335Implementation-dependent names must be in lowercase letters,1336as below:

LC_name

- 1338PATHMany historical implementations of the Bourne shell do not1339interpret a trailing colon to represent the current working1340directory and are thus nonconforming. The C Shell and the1341KornShell conform to POSIX.1 on this point. The usual name of1342dot may also be used to refer to the current working directory.
- 1343 TZ See 8.1.1 for an explanation of the format.

LOGNAME 4.3BSD uses the environment variable USER for this purpose. 1344 In most implementations, the value of such a variable is easily 1345 forged, so security-critical applications should rely on other 1346 means of determining user identity. LOGNAME is required to 1347 be constructed from the portable filename character set for rea-1348 sons of interchange. No diagnostic condition is specified for 1349 violating this rule, and no requirement for enforcement exists. 1350 The intent of the requirement is that if extended characters are 1351 used, the "guarantee" of portability implied by a standard is 1352 voided. (See also B.2.2.2.) 1353

The following environment variables have been used historically as indicated. However, such use was either so variant as to not be amenable to standardization or to be relevant only to other facilities not specified in POSIX.1, and they have therefore been excluded. They may or may not be included in future POSIX standards. Until then, writers of conforming applications should be aware that

1331

1332

1333

1334

L

details of the use of these variables are likely to vary in different contexts.

- 1360 **IFS** Characters used as field separators.
- 1361 MAIL System mailer information.
- 1362 **PS1** Prompting string for interactive programs.
- 1363 **PS2** Prompting string for interactive programs.
- 1364 SHELL The shell command interpreter name.

1365 B.2.7 C Language Definitions

1366 The construct <name. h> for headers is also taken from the C Standard {2}.

1367 B.2.7.1 Symbols From the C Standard

1368The reservation of identifiers is paraphrased from the C Standard {2}. The text is1369included because it needs to be part of POSIX.1, regardless of possible changes in1370future versions of the C Standard {2}. The reservation of other namespaces is par-1371ticularly for <errno.h>.

1372 These identifiers may be used by implementations, particularly for feature test 1373 macros. Implementations should not use feature test macro names that might be 1374 reasonably used by a standard.

1375 The requirement for representing the number of clock ticks in 24 h refers to the 1376 interval defined by POSIX.1, not to the interval defined by the C Standard {2}.

Including headers more than once is a reasonably common practice, and it should 1377 be carried forward from the C Standard {2}. More significantly, having definitions 1378 in more than one header is explicitly permitted. Where the potential declaration 1379 is "benign" (the same definition twice) the declaration can be repeated, if that is 1380 permitted by the compiler. (This is usually true of macros, for example.) In those 1381 situations where a repetition is not benign (e.g., typedefs), conditional compilation 1382 must be used. The situation actually occurs both within the C Standard (2) and 1383 within POSIX.1: time_t should be in <sys/types.h>, and the C Standard {2} 1384 mandates that it be in <time.h>. POSIX.1 requires using <sys/types.h> with 1385 <time.h> because of the common-usage environment. 1386

1387 B.2.7.2 POSIX.1 Symbols

This subclause addresses the issue of "namespace pollution." The C Standard {2} requires that the namespace beyond what it reserves not be altered except by explicit action of the application writer. This subclause defines the actions to add the POSIX.1 symbols for those headers where both the C Standard {2} and POSIX.1 | need to define symbols. Where there are nonoverlapping uses of headers, there is no problem.

The list of symbols defined in the C Standard (2) is summarized in the rationale associated with Annex C.

1396 Implementors should note that the requirement on type conversion disallows

using an older declaration as a prototype and in effect requires that the number of
 arguments in the prototype match that given in POSIX.1.

When headers are used to provide symbols, there is a potential for introducing 1399 symbols that the application writer cannot predict. Ideally, each header should 1400 only contain one set of symbols, but this is not practical for historical reasons. 1401 Thus, the concept of feature test macros is included. This is done in a general 1402 manner because it is expected that future additions to POSIX.1 and other related 1403 standards will have this same problem. (Future standards not constrained by 1404 historical practice should avoid the problem by using new header files rather than 1405 1406 using ones already extant.)

This idea is split into two subclauses: 2.7.2.1 covers the case of the C Standard | {2} conformant systems, where the requirements of the C Standard {2} are that unless specifically requested the application will not see any other symbols, and "Common Usage," where the default set of symbols is not well controlled and backwards compatibility is an issue.

The common usage case is the more difficult to define. In the C Standard (2) case, 1412each feature test macro simply adds to the possible symbols. In common usage, 1413 _POSIX_SOURCE is a special case in that it reduces the set to the sum of the 1414 C Standard (2) and POSIX.1. (The developers of the C Standard (2) will determine 1415 if they want a similar macro to limit the features to just the C Standard (2); the 1416 wording permits this because under those circumstances _POSIX_SOURCE would 1417 be just another ordinary feature test macro. The only order requirement is 1418 "before headers.") 1419

1420 If _POSIX_SOURCE is not defined in a common-usage environment, the user 1421 presumably gets the same results as in previous releases. Some applications may 1422 today be conformant without change, so they would continue to compile as long as 1423 common usage is provided. When the C Standard {2} is the default they will have 1424 to change (unless they are already C Standard {2} conformant), but this can be 1425 done gradually.

Note that the net result of defining _POSIX_SOURCE at the beginning of a program is in either case the same: the implementation-defined symbols are only visible if they are requested. (But if _POSIX_SOURCE is not used, the implementation default, which is probably backwards compatible, determines their visibility.)

The area of namespace pollution versus additions to structures is difficult because
of the macro structure of C. The following discussion summarizes all the various
problems with and objections to the issue.

Note the phrase "user defined macro." Users are not permitted to define macro names (or any other name) beginning with _[A-Z_]. Thus, the conflict cannot occur for symbols reserved to the vendor's namespace, and the permission to add fields automatically applies, without qualification, to those symbols.

- 1438 (1) Data structures (and unions) need to be defined in headers by implementations to meet certain requirements of POSIX.1 and the C Standard {2}.
- (2) The structures defined by POSIX.1 are typically minimal, and any practical implementation would wish to add fields to these structures either to hold additional related information or for backwards compatibility (or

- 1443both). Future standards (and de facto standards) would also wish to add1444to these structures. Issues of field alignment make it impractical (at1445least in the general case) to simply omit fields when they are not defined1446by the particular standard involved.
- 1447Struct dirent is an example of such a minimal structure (although one1448could argue about whether the other fields need visible names). The1449st_rdev field of most implementations' stat structure is a common exam-1450ple where extension is needed and where a conflict could occur.
- (3) Fields in structures are in an independent namespace, so the addition of
 such fields presents no problem to the C language itself in that such
 names cannot interact with identically named user symbols because
 access is qualified by the specific structure name.
- (4) There is an exception to this: macro processing is done at a lexical level. 1455 Thus, symbols added to a structure might be recognized as user-provided 1456 macro names at the location where the structure is declared. This only 1457 can occur if the user-provided name is declared as a macro before the 1458 header declaring the structure is included. The user's use of the name 1459 after the declaration cannot interfere with the structure because the sym-1460 bol is hidden and only accessible through access to the structure. 1461 Presumably, the user would not declare such a macro if there was an 1462 intention to use that field name. 1463
- 1464 (5) Macros from the same or a related header might use the additional fields
 1465 in the structure, and those field names might also collide with user mac1466 ros. Although this is a less frequent occurrence, since macros are
 1467 expanded at the point of use, no constraint on the order of use of names
 1468 can apply.
- (6) An "obvious" solution of using names in the reserved namespace and then redefining them as macros when they should be visible does not work because this has the effect of exporting the symbol into the general namespace. For example, given a (hypothetical) system-provided header
 (h.h>, and two parts of a C program in a.c and b.c:
- 1474 In header <h.h>:

1475	struct foo {
1476	inti;
1477	}
1478	<pre>#ifdef FEATURE_TEST</pre>
1479	#define i i;
1480	#endif
1481	In file a.c:
1482	#include h.h
1483	extern int i;
1484	
1485	In file b.c:

1486 1487		extern int i;	
1488		The symbol that the user thinks of as i in both files has an external	
1/89		name of " i " in a c: the same symbol i in b c has an external name	
1400		"i " (ignoring any hidden manipulations the compiler might perform on	
1490		the normed). This would cause a mustarious norme resolution mehlem	
1491		the names). This would cause a mysterious name resolution problem	
1492		when a.o and b.o are linked.	
1493		Simply avoiding definition then causes alignment problems in the	
1494		structure.	
1495		A structure of the form	
1496		struct foo {	
1497		union {	
1498		inti;	
1499		#ifdef FEATURE TEST	
1500		int i;	
1501		#endif	
1502		}ii;	
1503		}	
1504		does not work because the name of the logical field i is " ii.i", and	
1505		introduction of a macro to restore the logical name immediately reintro-	
1506		duces the problem discussed previously (although its manifestation	
1507		might be more immediate because a syntax error would result if a recur-	
1508		sive macro did not cause it to fail first).	
1509	(7)	A more workable solution would be to declare the structure:	
1510		struct foo {	
1511		#ifdef FEATURE TEST	
1512		int i;	
1513		#else	
1514		int i;	
1515		#endif	
1516		}	
1517		However, if a macro (particularly one required by a standard) is to be	
1518		defined that uses this field, two must be defined: one that uses i. the	
1519		other that uses i. If more than one additional field is used in a macro	
1520		and they are conditional on distinct combinations of features the com-	
1591		navity goes up as 2^n	
1971		piently gues up as 2.	
1522	All this	leaves a difficult situation: vendors must provide very complex headers to	
1523	deal with what is conceptually simple and safe: adding a field to a structure. It is		
1524	the possibility of user-provided macros with the same name that makes this		
1525	difficult.		

Several alternatives were proposed that involved constraining the user's access to 1526part of the namespace available to the user (as specified by the C Standard {2}). 1527In some cases, this was only until all the headers had been included. There were 1528two proposals discussed that failed to achieve consensus: 1529

- 1530 Limiting it for the whole program.
- Restricting the use of identifiers containing only uppercase letters until after all system headers had been included. It was also pointed out that because macros might wish to access fields of a structure (and macro expansion occurs totally at point of use) restricting names in this way would not protect the macro expansion, and thus the solution was inadequate.
- 1537 It was finally decided that reservation of symbols would occur, but as constrained.
- The current wording also allows the addition of fields to a structure, but requires that user macros of the same name not interfere. This allows vendors to either:
- Not create the situation [do not extend the structures with user-accessible
 names or use the solution in (7) above] or
- Extend their compilers to allow some way of adding names to structures
 and macros safely.
- There are at least two ways that the compiler might be extended: add new preprocessor directives that turn off and on macro expansion for certain symbols (without changing the value of the macro) and a function or lexical operation that suppresses expansion of a word. The latter seems more flexible, particularly because it addresses the problem in macros as well as in declarations.
- The following seems to be a possible implementation extension to the C language that will do this: any token that during macro expansion is found to be preceded by three # symbols shall not be further expanded in exactly the same way as described for macros that expand to their own name as in section 3.8.3.4 of the C Standard {2}. A vendor may also wish to implement this as an operation that is lexically a function, which might be implemented as
- 1555 #define safe name(x) ###x
- Using a function notation would insulate vendors from changes in standards until such a functionality is standardized (if ever). Standardization of such a function would be valuable because it would then permit third parties to take advantage of it portably in software they may supply.
- The symbols that are "explicitly permitted, but not required by this part of ISO/IEC 9945" include those classified below. (That is, the symbols classified below might, but are not required to, be present when _POSIX_SOURCE is defined.)
- Symbols in 2.8 and 2.9 that are defined to indicate support for options or limits that are constant at compile-time.
- 1566 Symbols in the namespace reserved for the implementation by the 1567 C Standard {2}.
- Symbols in a namespace reserved for a particular type of extension (e.g.,
 type names ending with _t in <sys/types.h>).
- Additional members of structures or unions whose names do not reduce the
 namespace reserved for applications (see B.2.7.2).

1200

The phrase "when that header is included" was chosen to allow any fine structure of auxiliary headers the implementor may choose to use, as long as the net result is as required.

There are several common environments available today where a feature test macro would be useful to applications programmers during the transition to standard-conforming environments from certain common historical environments. The symbols in Table B-1, derived from common porting bases and industry specifications are suggested.

1582	Symbol	Description	
1583	_V7	Version 7	
1584	_BSD	General BSD systems	
1585	_BSD4_2	4.2BSD	
1586	_BSD4_3	4.3BSD	
1587	_SYSIII	System III	
1588	_SYSV	System V.1, V.2	
1589	_SYSV3	System V.3	
1590	_XPGn	X/Open Portability Guide, Issue n	
1591	_USR_GROUP	The 1984 /usr/group standard	
1592			

Table B-1 - Suggested Feature Test Macros

Only symbols that are actually in the porting base or industry specification should be enabled by these symbols.

Feature test macros for implementation extensions will also probably be required. Quite a few of these are traditionally available, but are in violation of the intent of namespace pollution control. These can be made conforming simply by prefixing them with an underscore. Symbols beginning with "_POSIX" are strongly discouraged, as they will probably be used by later revisions of POSIX.1.

The environment for compilation has traditionally been fairly portable in historical systems, but during the transition to the C Standard (2) there will be confusion about how to specify that a C Standard (2) compiler is expected, as considerations of backwards compatibility will constrain many implementors from providing a conformant environment replacing the traditional one. This concern has more to do with the issues of namespace than with the syntax of the language accepted, which is highly compatible.

For systems that are sufficiently similar to traditional UNIX systems for this to make sense, it is suggested that if a compilation line of the form

1609 cc -D__STDC__...

is provided, that the system provide an environment that is conformant with theC Standard {2}, at least with respect to namespace.

1612 It was decided to use feature test macros, rather than the inclusion of a header,
1613 both because <unistd.h> was already in use and would itself have this problem,
1614 and because the underlying mechanism would probably have been this anyway,
1615 but in a less flexible fashion.

L

I

I

ł

POSIX.1 requires that headers be included in all cases, although it is not directly clear from the text at this point in the standard. If a function does not need any special types, then it must be declared in <unistd.h>, as stated here. If it does require something special, then it has an associated header, and the program will not compile without that header.

1621 B.2.7.3 Headers and Function Prototypes

The statement that names need not be carried forward literally exists for several reasons. These include the fact that some vendors may historically use other names and that the names are irrelevant to application portability. More importantly, because of the pervasive nature of C macros, a declaration of the form:

1626 kill (pid_t pid, int sig);

1627 could be seriously undermined by a (perfectly valid) user declaration of the form:

1628 #define pid statusstruct.pidinfo

1629 B.2.8 Numerical Limits

1630 This subclause clarifies the scope and mutability of several classes of limits.

1631 B.2.8.1 C Language Limits

- 1632 See also 2.7 and B.1.1.1.
- 1633{CHAR_MIN}It is possible to tell if the implementation supports native char-
acter comparison as signed or unsigned by comparing this limit
to zero.
- 1636 {WORD_BIT} This limit has been omitted, as it is not referenced elsewhere in POSIX.1.

No limits are given in <limits.h> for floating point values because none of the functions in POSIX.1 use floating point values, and all the functions that do that are imported from the C Standard {2} by 8.1, as are the limits that apply to the floating point values associated with them.

Though limits to the addresses to system calls were proposed, they were not included in POSIX.1 because it is not clear how to implement them for the range of systems being considered, and no complete proposal was ever received. Limits | regarding hardware register characteristics were similarly proposed and not attempted.

1647 B.2.8.2 Minimum Values

1648 There has been confusion about the minimum maxima, and when that is under-1649 stood there is still a concern about providing ways to allocate storage based on the 1650 symbols. This is particularly true for those in 2.8.4 where an indeterminate value 1651 will leave the programmer with no symbol upon which to fall back.

1652 Providing explicit symbols for the minima (from the implementor's point of view, 1653 or maxima from the the application's point of view) helps to resolve possible

confusion. Symbols are still provided for the actual value, and it is expected that many applications will take advantage of these larger values, but they need not do so unless it is to their advantage. Where the values in this subclause are adequate for the application, it should use them. These are given symbolically both because it is easier to understand and because the values of these symbols could change between revisions of POSIX.1. Arguments to "good programming practice" also apply.

1661 B.2.8.3 Run-Time Increasable Values

The heading of the far-right column of the table is given as "Minimum Value" 1662 rather than "Value" in order to emphasize that the numbers given in that column 1663 are minimal for the actual values a specific implementation is permitted to define 1664 in its <limits.h>. The values in the actual <limits.h> define, in turn, the 1665 maximum amount of a given resource that a Conforming POSIX.1 Application can 1666 depend on finding when translated to execute on that implementation. A Con-1667 forming POSIX.1 Application Using Extensions must function correctly even if the 1668 value given in <limits.h> is the minimum that is specified in POSIX.1. (The 1669 application may still be written so that it performs more efficiently when a larger 1670 value is found in <limits.h>.) A conforming implementation must provide at 1671 least as much of a particular resource as that given by the value in POSIX.1. An 1672 implementation that cannot meet this requirement (a "toy implementation") can-1673 not be a conforming implementation. 1674

1675 B.2.8.4 Run-Time Invariant Values (Possibly Indeterminate)

- 1676 {CHILD_MAX} This name can be misleading. This limit applies to all
 1677 processes in the system with the same user ID, regardless of
 1678 ancestry.
- 1679 B.2.8.5 Pathname Variable Values

[MAX_INPUT] Since the only use of this limit is in relation to terminal input 1680 queues, it mentions them specifically. This limit was originally 1681 named {MAX_CHAR}. Application writers should use 1682 {MAX_INPUT} primarily as an indication of the number of bytes 1683 that can be written as a single unit by one Conforming POSIX.1 1684 Application Using Extensions communicating with another via 1685 a terminal device. It is not implied that input lines received 1686 from terminal devices always contain {MAX_INPUT} bytes or 1687 fewer: an application that attempts to read more than 1688 {MAX_INPUT} bytes from a terminal may receive more than 1689 {MAX_INPUT} bytes. 1690

1691It is not obvious that {MAX_INPUT} is of direct value to the1692application writer. The existence of such a value (whatever it1693may be) is directly of use in understanding how the tty driver1694works (particularly with respect to flow control and dropped1695characters). The value can be determined by finding out when1696flow control takes effect (see the description of IXOFF in16977.1.2.2).

1698 1699 1700 1701		Understanding that the limit exists and knowing its magnitude is important to making certain classes of applications work correctly. It is unlikely to be used in an application, but its presence makes POSIX.1 clearer.
1702 1703 1704 1705	{PATH_MAX}	A Conforming POSIX.1 Application or Conforming POSIX.1 Application Using Extensions that, for example, compiles to use different algorithms depending on the value of {PATH_MAX} should use code such as:
1706 1707 1708 1709 1710 1711 1712 1713 1714		<pre>#if defined(PATH_MAX) && PATH_MAX < 512 #else #if defined(PATH_MAX) /* PATH_MAX >= 512 */ #else /* PATH_MAX indeterminate */ #endif #endif</pre>
1715 1716 1717 1718 1719		This is because the value tends to be very large or indeter- minate on most historical implementations (it is arbitrarily large on System V). On such systems there is no way to quan- tify the limit, and it seems counterproductive to include an artificially small fixed value in <limits.h> in such cases.</limits.h>

- 1720 B.2.9 Symbolic Constants
- 1721 B.2.9.1 Symbolic Constants for the access() Function
- 1722 There is no additional rationale provided for this subclause.
- 1723 **B.2.9.2 Symbolic Constants for the** *lseek*() Function
- 1724 There is no additional rationale provided for this subclause.

1725 B.2.9.3 Compile-Time Symbolic Constants for Portability Specifications

The purpose of this material is to allow an application developer to have a chance | to determine whether a given application would run (or run well) on a given implementation. To this purpose has been added that of simplifying development of verification suites for POSIX.1. The constants given here were originally proposed for a separate file, <posix.h>, but it was decided that they should appear | in <unistd.h> along with other symbolic constants.

1732 **B.2.9.4 Execution-Time Symbolic Constants for Portability Specifications**

1733 Without the addition of {_POSIX_NO_TRUNC} and {_PC_NO_TRUNC} to this list, 1734 POSIX.1 says nothing about the effect of a pathname component longer than 1735 {NAME_MAX}. There are only two effects in common use in implementations: 1736 truncation or an error. It is desirable to limit allowable behavior to these two 1737 cases. It is also desirable to permit applications to determine what an 1738 implementation's behavior is because services that are available with one 1739 behavior may be impractical to provide with the other. However, since the 1740 behavior may vary from one file system to another, it may be necessary to use 1741 pathconf() to resolve it.

1742 **B.3 Process Primitives**

Consideration was given to enumerating all characteristics of a process defined by 1743 POSIX.1 and describing each function in terms of its effects on those characteris-1744 tics, rather than English text. This is quite different from any known descriptions 1745 of historical implementations, and it was not certain that this could be done ade-1746 quately and completely enough to produce a usable standard. Providing such 1747 descriptions in addition to the text was also considered. This was not done 1748 because it would provide at best two redundant descriptions, and more likely two 1749 descriptions with subtle inconsistencies. 1750

1751 B.3.1 Process Creation and Execution

Running a new program takes two steps. First the existing process (the parent) calls the fork() function, producing a new process (the child), which is a copy of itself. One of these processes (normally, but not necessarily, the child) then calls one of the *exec* functions to overlay itself with a copy of the new process image.

If the new program is to be run synchronously (the parent suspends execution 1756 until the child completes), the parent process then uses either the wait() or wait-1757 *pid()* function. If the new program is to be run asynchronously, it does not suffice 1758 to simply omit the wait() or waitpid() call, because after the child terminates it 1759 continues to hold some resources until it is waited for. A common way to produce 1760 ("spawn") a descendant process that does not need to be waited on is to fork() to 1761 produce a child and wait() on the child. The child fork()s again to produce a 1762 grandchild. The child then exits and the parent's wait() returns. The grandchild 1763 is thus disinherited by its grandparent. 1764

- 1765 A simpler method (from the programmer's point of view) of spawning is to do
- 1766 system("something &");

However, this depends on features of a process (the shell) that are outside the
scope of POSIX.1, although they are currently being addressed by the working
group preparing ISO/IEC 9945-2 {B36}.

1770 B.3.1.1 Process Creation

1771 Many historical implementations have timing windows where a signal sent to a 1772 process group (e.g., an interactive SIGINT) just prior to or during execution of 1773 fork() is delivered to the parent following the fork() but not to the child because 1774 the fork() code clears the child's set of pending signals. POSIX.1 does not require, 1775 or even permit, this behavior. However, it is pragmatic to expect that problems of 1776 this nature may continue to exist in implementations that appear to conform to 1777 POSIX.1 and pass available verification suites. This behavior is only a consequence of the implementation failing to make the interval between signal genera-1778 tion and delivery totally invisible. From the application's perspective, a fork() call 1779 should appear atomic. A signal that is generated prior to the fork() should be 1780 delivered prior to the fork(). A signal sent to the process group after the fork()1781 should be delivered to both parent and child. The implementation might actually 1782 initialize internal data structures corresponding to the child's set of pending sig-1783 nals to include signals sent to the process group during the fork(). Since the 1784 fork() call can be considered as atomic from the application's perspective, the set 1785 would be initialized as empty and such signals would have arrived after the 1786 fork(). See also B.3.3.1.2. 1787

One approach that has been suggested to address the problem of signal inheritance across fork() is to add an [EINTR] error, which would be returned when a signal is detected during the call. While this is preferable to losing signals, it was not considered an optimal solution. Although it is not recommended for this purpose, such an error would be an allowable extension for an implementation.

The [ENOMEM] error value is reserved for those implementations that detect and 1793 distinguish such a condition. This condition occurs when an implementation 1794 detects that there is not enough memory to create the process. This is intended to 1795 be returned when [EAGAIN] is inappropriate because there can never be enough 1796 memory (either primary or secondary storage) to perform the operation. Because 1797 fork() duplicates an existing process, this must be a condition where there is 1798 sufficient memory for one such process, but not for two. Many historical imple-1799 mentations actually return [ENOMEM] due to temporary lack of memory, a case 1800 that is not generally distinct from [EAGAIN] from the perspective of a portable 1801 application. 1802

Part of the reason for including the optional error [ENOMEM] is because the SVID
(B39) specifies it and it should be reserved for the error condition specified there.
The condition is not applicable on many implementations.

1806 IEEE Std 1003.1-1988 neglected to require concurrent execution of the parent and 1807 child of fork(). A system that single-threads processes was clearly not intended 1808 and is considered an unacceptable, "toy implementation" of POSIX.1. The only 1809 objection anticipated to the phrase "executing independently" is testability, but 1810 this assertion should be testable. Such tests require that both the parent and 1811 child can block on a detectable action of the other, such as a write to a pipe or a 1812 signal. An interactive exchange of such actions should be possible for the system 1813 to conform to the intent of POSIX.1.

The [EAGAIN] error exists to warn applications that such a condition might occur. Whether it will occur or not is not in any practical sense under the control of the application because the condition is usually a consequence of the user's use of the system, not of the application's code. Thus, no application can or should rely upon its occurrence under any circumstances, nor should the exact semantics of what concept of "user" is used be of concern to the application writer. Validation writers should be cognizant of this limitation.

I

l

1821 B.3.1.2 Execute a File

Early drafts of POSIX.1 required that the value of argc passed to main() be "one or 1822 greater." This was driven by the same requirement in drafts of the C Standard 1823 {2}. In fact, historical implementations have passed a value of zero when no argu-1824 ments are supplied to the caller of the *exec* functions. This requirement was 1825 removed from the C Standard (2) and subsequently removed from POSIX.1 as well. 1826 The POSIX.1 wording, in particular the use of the word "should," requires a 1827 Strictly Conforming POSIX.1 Application (see 1.3.3) to pass at least one argument 1828 to the *exec* function, thus guaranteeing that *argc* be one or greater when invoked 1829 by such an application. In fact, this is good practice, since many existing applica-1830 tions reference argv[0] without first checking the value of argc. 1831

The requirement on a Strictly Conforming POSIX.1 Application also states that 1832 the value passed as the first argument be a filename associated with the process 1833 being started. Although some existing applications pass a pathname rather than 1834 a filename in some circumstances, a filename is more generally useful, since the 1835 common usage of argv[0] is in printing diagnostics. In some cases the filename 1836 passed is not the actual filename of the file; for example, many implementations 1837 of the login utility use a convention of prefixing a hyphen (-) to the actual 1838 filename, which indicates to the command interpreter being invoked that it is a 1839 "login shell." 1840

Some systems can exec shell scripts. This functionality is outside the scope of 1841 POSIX.1, since it requires standardization of the command interpreter language of 1842 the script and/or where to find a command interpreter. These fall in the domain 1843 of the shell and utilities standard, currently under development as ISO/IEC 9945-2 1844 (B36). However, it is important that POSIX.1 neither require nor preclude any 1845 reasonable implementation of this behavior. In particular, the description of the 1846 [ENOEXEC] error is intended to permit discretion to implementations on whether 1847 to give this error for shell scripts. 1848

One common historical implementation is that the execl(), execv(), execle(), and 1849 1850 execue() functions return an [ENOEXEC] error for any file not recognizable as executable, including a shell script. When the execlp() and execvp() functions 1851 encounter such a file, they assume the file to be a shell script and invoke a known 1852 command interpreter to interpret such files. These implementations of execup() 1853 and *execlp()* only give the [ENOEXEC] error in the rare case of a problem with the 1854 command interpreter's executable file. Because of these implementations the 1855 [ENOEXEC] error is not mentioned for execlp() or execup(), although implementa-1856 tions can still give it. 1857

Another way that some historical implementations handle shell scripts is by recognizing the first two bytes of the file as the character string #! and using the remainder of the first line of the file as the name of the command interpreter to execute.

Some implementations provide a third argument to main() called envp. This is defined as a pointer to the environment. The C Standard {2} specifies invoking main() with two arguments, so implementations must support applications written this way. Since POSIX.1 defines the global variable environ, which is also provided by historical implementations and can be used anywhere envp could be used, there is no functional need for the envp argument. Applications should use the getenv() function rather than accessing the environment directly via either
envp or environ. Implementations are required to support the two-argument calling sequence, but this does not prohibit an implementation from supporting envp
as an optional, third argument.

POSIX.1 specifies that signals set to SIG_IGN remain set to SIG_IGN and that the 1872 process signal mask be unchanged across an exec. This is consistent with histori-1873 cal implementations, and it permits some useful functionality, such as the nohup 1874 command. However, it should be noted that many existing applications wrongly 1875 assume that they start with certain signals set to the default action and/or 1876 unblocked. In particular, applications written with a simpler signal model that 1877 does not include blocking of signals, such as the one in the C Standard {2}, may 1878 not behave properly if invoked with some signals blocked. Therefore, it is best not 1879 to block or ignore signals across execs without explicit reason to do so, and espe-1880 cially not to block signals across execs of arbitrary (not closely co-operating) 1881 programs. 1882

1883 If {_POSIX_SAVED_IDS} is defined, the *exec* functions always save the value of the 1884 effective user ID and effective group ID of the process at the completion of the 1885 *exec*, whether or not the set-user-ID or the set-group-ID bit of the process image 1886 file is set.

The statement about argv[] and envp[] being constants is included to make expli-1887 cit to future writers of language bindings that these objects are completely con-1888 stant. Due to a limitation of the C Standard {2}, it is not possible to state that 1889 idea in Standard C. Specifying two levels of const-qualification for the argv[] 1890 and envp[] parameters for the exec functions may seem to be the natural choice, 1891 given that these functions do not modify either the array of pointers or the charac-1892 ters to which the function points, but this would disallow existing correct code. 1893 Instead, only the array of pointers is noted as constant. The table of assignment 1894 compatibility for dst = src, derived from the C Standard {2}, summarizes the 1895 compatibility: 1896

1897		dst:			
1898			const	char	const
1899		<u>char *[]</u>	<pre>char*[]</pre>	*const[]	char*const[]
1900	src:				
1901	char *[]	VALID		VALID	
1902	<pre>const char *[]</pre>		VALID		VALID
1903	char * const []			VALID	
1904	<pre>const char *const[]</pre>				VALID

Since all existing code has a source type matching the first row, the column that gives the most valid combinations is the third column. The only other possibility is the fourth column, but using it would require a cast on the *argv* or *envp* arguments. It is unfortunate that the fourth column cannot be used, because the declaration a nonexpert would naturally use would be that in the second row.

The C Standard $\{2\}$ and POSIX.1 do not conflict on the use of *environ*, but some historical implementations of *environ* may cause a conflict. As long as *environ* is treated in the same way as an entry point [e.g., fork()], it conforms to both standards. A library can contain fork(), but if there is a user-provided fork(), that fork() is given precedence and no problem ensues. The situation is similar for
 environ—the POSIX.1 definition is to be used if there is no user-provided environ
 to take precedence. At least three implementations are known to exist that solve
 this problem.

1918[E2BIG]The limit {ARG_MAX} applies not just to the size of the argument list, but to the sum of that and the size of the environment list.1920ment list.

1921[EFAULT]Some historical systems return [EFAULT] rather than1922[ENOEXEC] when the new process image file is corrupted. They1923are nonconforming.

1924 [ENAMETOOLONG]

- Since the file pathname may be constructed by taking elements in the **PATH** variable and putting them together with the filename, the [ENAMETOOLONG] condition could also be reached this way.
- 1929[ETXTBSY]The error [ETXTBSY] was considered too implementation1930dependent to include. System V returns this error when the1931executable file is currently open for writing by some process.1932POSIX.1 neither requires nor prohibits this behavior.

Other systems (such as System V) may return [EINTR] from *exec*. This is not addressed by POSIX.1, but implementations may have a window between the call to *exec* and the time that a signal could cause one of the *exec* calls to return with [EINTR].

1937 B.3.2 Process Termination

Early drafts drew a different distinction between normal and abnormal process 1938 termination. Abnormal termination was caused only by certain signals and 1939 resulted in implementation-defined "actions," as discussed below. Subsequent 1940 drafts of POSIX.1 distinguished three types of termination: normal termination 1941 (as in the current POSIX.1), "simple abnormal termination," and "abnormal termi-1942 nation with actions." Again the distinction between the two types of abnormal 1943 termination was that they were caused by different signals and that 1944 implementation-defined actions would result in the latter case. Given that these 1945 actions were completely implementation defined, the early drafts were only saying 1946 when the actions could occur and how their occurrence could be detected, but not 1947 what they were. This was of little or no use to portable applications, and thus the 1948 distinction was dropped from POSIX.1. 1949

The implementation-defined actions usually include, in most historical implementations, the creation of a file named core in the current working directory of the process. This file contains an image of the memory of the process, together with descriptive information about the process, perhaps sufficient to reconstruct the state of the process at the receipt of the signal.

There is a potential security problem in creating a core file if the process was set-user-ID and the current user is not the owner of the program, if the process was set-group-ID and none of the user's groups match the group of the program,

L
or if the user does not have permission to write in the current directory. In this
situation, an implementation either should not create a core file or should make
it unreadable by the user.

Despite the silence of POSIX.1 on this feature, applications are advised not to
create files named core because of potential conflicts in many implementations.
Some historical implementations use a different name than core for the file, such
as by appending the process ID to the filename.

1965 B.3.2.1 Wait for Process Termination

A call to the wait() or waitpid() function only returns status on an immediate 1966 child process of the calling process; i.e., a child that was produced by a single 1967 fork() call (perhaps followed by an exec or other function calls) from the parent. If 1968 a child produces grandchildren by further use of fork(), none of those grandchil-1969 dren nor any of their descendants will affect the behavior of a wait() from the ori-1970 ginal parent process. Nothing in POSIX.1 prevents an implementation from pro-1971 viding extensions that permit a process to get status from a grandchild or any 1972 other process, but a process that does not use such extensions must be guaranteed 1973 to see status from only its direct children. 1974

- 1975 The *waitpid()* function is provided for three reasons:
- 1976 To support job control (see B.3.3).
- 1977 To permit a nonblocking version of the *wait()* function.
- 1978 To permit a library routine, such as system() or pclose(), to wait for its chil 1979 dren without interfering with other terminated children for which the pro 1980 cess has not waited.

The first two of these facilities are based on the wait3() function provided by 1981 4.3BSD. The interface uses the options argument, which is identical to an argu-1982 ment to wait3(). The WUNTRACED flag is used only in conjunction with job con-1983 trol on systems supporting that option. Its name comes from 4.3BSD and refers to 1984 the fact that there are two types of stopped processes in that implementation: 1985 processes being traced via the ptrace() debugging facility and (untraced) processes 1986 stopped by job-control signals. Since *ptrace()* is not part of POSIX.1, only the 1987 second type is relevant. The name WUNTRACED was retained because its usage 1988 is the same, even though the name is not intuitively meaningful in this context. 1989

The third reason for the *waitpid()* function is to permit independent sections of a process to spawn and wait for children without interfering with each other. For example, the following problem occurs in developing a portable shell, or command interpreter:

```
1994stream = popen("/bin/true");1995(void) system("sleep 100");1996(void) pclose(stream);
```

1997 On all historical implementations, the final *pclose()* will fail to reap the wait 1998 status of the *popen()*.

1999 The status values are retrieved by macros, rather than given as specific bit encod-2000 ings as they are in most historical implementations (and thus expected by

existing programs). This was necessary to eliminate a limitation on the number of signals an implementation can support that was inherent in the traditional encodings. POSIX.1 does require that a status value of zero corresponds to a process calling $_exit(0)$, as this is the most common encoding expected by existing programs. Some of the macro names were adopted from 4.3BSD.

These macros syntactically operate on an arbitrary integer value. The behavior is 2006 undefined unless that value is one stored by a successful call to *wait()* or *wait-*2007 pid() in the location pointed to by the stat_loc argument. An earlier draft 2008 attempted to make this clearer by specifying each argument as *stat_loc rather 2009 than stat val. However, that did not follow the conventions of other specifications 2010 in POSIX.1 or traditional usage. It also could have implied that the argument to 2011 the macro must literally be **stat loc*; in fact, that value can be stored or passed as 2012 an argument to other functions before being interpreted by these macros. 2013

The extension that affects wait() and waitpid() and is common in historical imple-2014 mentations is the *ptrace()* function. It is called by a child process and causes that 2015 child to stop and return a status that appears identical to the status indicated by 2016 WIFSTOPPED. The status of *ptraced* children is traditionally returned regardless 2017 of the WUNTRACED flag [or by the *wait()* function]. Most applications do not need 2018 to concern themselves with such extensions because they have control over what 2019 extensions they or their children use. However, applications, such as command 2020 interpreters, that invoke arbitrary processes may see this behavior when those 2021 arbitrary processes misuse such extensions. 2022

Implementations that support core file creation or other implementation-defined actions on termination of some processes traditionally provide a bit in the status returned by *wait()* to indicate that such actions have occurred.

2026 B.3.2.2 Terminate a Process

Most C language programs should use the *exit()* function rather than _*exit()*. The _*exit()* function is defined here instead of *exit()* because the C Standard {2} defines the latter to have certain characteristics that are beyond the scope of POSIX.1, specifically the flushing of buffers on open files and the use of *atexit()*. See "The C Language" in the Introduction. There are several public-domain implementations of *atexit()* that may be of use to interface implementors who wish to incorporate it.

It is important that the consequences of process termination as described in this subclause occur regardless of whether the process called $_exit()$ [perhaps indirectly through exit()] or instead was terminated due to a signal or for some other reason. Note that in the specific case of exit() this means that the status argument to exit() is treated the same as the status argument to $_exit()$. See also B.3.2.

A language other than C may have other termination primitives than the C language exit() function, and programs written in such a language should use its native termination primitives, but those should have as part of their function the behavior of $_{exit}()$ as described in this subclause. Implementations in languages other than C are outside the scope of the present version of POSIX.1, however.

As required by the C Standard {2}, using return from main() is equivalent to calling exit() with the same argument value. Also, reaching the end of the main()

2046 function is equivalent to using *exit()* with an unspecified value.

A value of zero (or EXIT_SUCCESS, which is required by 8.1 to be zero) for the 2047 argument status conventionally indicates successful termination. 2048 This corresponds to the specification for exit() in the C Standard {2}. The convention is 2049 followed by utilities such as make and various shells, which interpret a zero 2050 status from a child process as success. For this reason, applications should not 2051 call exit(0) or _exit(0) when they terminate unsuccessfully, for example in signal-2052 catching functions. 2053

- Historically, the implementation-dependent process that inherits children whose parents have terminated without waiting on them is called init and has a process ID of 1.
- The sending of a SIGHUP to the foreground process group when a controlling pro-2057 cess terminates corresponds to somewhat different historical implementations. In 2058 System V, the kernel sends a SIGHUP on termination of (essentially) a controlling 2059 process. In 4.2BSD, the kernel does not send SIGHUP in a case like this, but the 2060 termination of a controlling process is usually noticed by a system daemon, which 2061 arranges to send a SIGHUP to the foreground process group with the *vhangup()* 2062 function. However, in 4.2BSD, due to the behavior of the shells that support job 2063 control, the controlling process is usually a shell with no other processes in its 2064 process group. Thus, a change to make exit() behave this way in such systems 2065 should not cause problems with existing applications. 2066
- The termination of a process may cause a process group to become orphaned in 2067 either of two ways. The connection of a process group to its parent(s) outside of 2068 the group depends on both the parents and their children. Thus, a process group 2069 may be orphaned by the termination of the last connecting parent process outside 2070 of the group or by the termination of the last direct descendant of the parent 2071 process(es). In either case, if the termination of a process causes a process group 2072 to become orphaned, processes within the group are disconnected from their job 2073 control shell, which no longer has any information on the existence of the process 2074 group. Stopped processes within the group would languish forever. In order to 2075 avoid this problem, newly orphaned process groups that contain stopped processes 2076 are sent a SIGHUP signal and a SIGCONT signal to indicate that they have been 2077 disconnected from their session. The SIGHUP signal causes the process group 2078 members to terminate unless they are catching or ignoring SIGHUP. Under most 2079 circumstances, all of the members of the process group are stopped if any of them 2080 are stopped. 2081
- The action of sending a SIGHUP and a SIGCONT signal to members of a newly 2082 orphaned process group is similar to the action of 4.2BSD, which sends SIGHUP 2083 and SIGCONT to each stopped child of an exiting process. If such children exit in 2084 response to the SIGHUP, any additional descendants will receive similar treat-2085 ment at that time. In POSIX.1, the signals will be sent to the entire process group 2086 at the same time. Also, in POSIX.1, but not in 4.2BSD, stopped processes may be 2087 orphaned, but may be members of a process group that is not orphaned; therefore, 2088 the action taken at *exit()* must consider processes other than child processes. 2089
- It is possible for a process group to be orphaned by a call to *setpgid()* or *setsid()*, as well as by process termination. POSIX.1 does not require sending SIGHUP and SIGCONT in those cases, because, unlike process termination, those cases will not

1

be caused accidentally by applications that are unaware of job control. An implementation can choose to send SIGHUP and SIGCONT in those cases as an extension; such an extension must be documented as required in 3.3.1.2.

2096 **B.3.3 Signals**

Signals, as defined in Version 7, System III, the 1984 /usr/group Standard (B59), 2097 and System V (except very recent releases), have shortcomings that make them 2098 unreliable for many application uses. Several objections were raised against early 2099 drafts of POSIX.1 because of this. Therefore, a new signal mechanism, based very 2100 closely on the one of 4.2BSD and 4.3BSD, was added to POSIX.1. With the excep-2101 tion of one feature [see item (4) below and also sigpending()], it is possible to 2102 implement the POSIX.1 interface as a simple library veneer on top of 4.3BSD. 2103 There are also a few minor aspects of the underlying 4.3BSD implementation (as 2104 opposed to the interface) that would also need to change to conform to POSIX.1. 2105

- 2106 The major differences from the BSD mechanism are:
- 2107 (1) Signal mask type. BSD uses the type int to represent a signal mask, thus
 2108 limiting the number of signals to the number of bits in an int (typically
 2109 32). The new standard instead uses a defined type for signal masks.
 2110 Because of this change, the interface is significantly different than it is in
 2111 BSD implementations, although the functionality, and potentially the
 2112 implementation, are very similar.
- (2) Restarting system calls. Unlike all previous historical implementations, 2113 4.2BSD restarts some interrupted system calls rather than returning an 2114error with *errno* set to [EINTR] after the signal-catching function returns. 2115 This change caused problems for some existing application code. 4.3BSD 2116 and other systems derived from 4.2BSD allow the application to choose 2117 whether system calls are to be restarted. POSIX.1 (in 3.3.4) does not 2118 require restart of functions because it was not clear that the semantics of 2119 system-call restart in any historical implementation were useful enough 2120 to be of value in a standard. Implementors are free to add such mechan-2121 2122 isms as extensions.
- (3) Signal stacks. The 4.2BSD mechanism includes a function sigstack().
 The 4.3BSD mechanism includes this and a function sigreturn(). No
 equivalent is included in POSIX.1 because these functions are not portable, and no sufficiently portable and useful equivalent has been identified. See also 8.3.1.
- (4) Pending signals. The sigpending() function is the sole new signal opera tion introduced in POSIX.1.

A proposal was considered for making reliable signals optional. However, the 2130consensus was that this would hurt application portability, as a large percentage 2131 of applications using signals can be hurt by the unreliable aspects of historical 2132 implementations of the signal() mechanism defined by the C Standard (2). This 2133 unreliability stems from the fact that the signal action is reset to SIG_DFL before 2134 the user's signal-catching routine is entered. The C Standard (2) does not require 2135 2136 this behavior, but does explicitly permit it, and most historical implementations behave this way. 2137

I

1

For example, an application that catches the SIGINT signal using signal() could 2138 be terminated with no chance to recover when two such signals arrive sufficiently 2139 close in time (e.g., when an impatient user types the INTR character twice in a 2140 row on a busy system). Although the C Standard (2) no longer requires this 2141 unreliable behavior, many historical implementations, including System V, will 2142 reset the signal action to SIG_DFL. For this reason, it is strongly recommended 2143 that the signal() function not be used by POSIX.1 conforming applications. Imple-2144 mentations should also consider blocking signals during the execution of the 2145 signal-catching function instead of resetting the action to SIG_DFL, but backward 2146 compatibility considerations will most likely prevent this from becoming 2147 universal. 2148

Most historical implementations do not queue signals; i.e., a process's signal handler is invoked once, even if the signal has been generated multiple times before it is delivered. A notable exception to this is SIGCLD, which, in System V, is queued. The queueing of signals is neither required nor prohibited by POSIX.1. See 3.3.1.2. It is expected that a future realtime extension to POSIX.1 will address the issue of reliable queueing of event notification.

2155 B.3.3.1 Signal Concepts

2156 B.3.3.1.1 Signal Names

The restriction on the actual type used for *sigset_t* is intended to guarantee that these objects can always be assigned, have their address taken, and be passed as parameters by value. It is not intended that this type be a structure including pointers to other data structures, as that could impact the portability of applications performing such operations. A reasonable implementation could be a structure containing an array of some integer type.

The signals described in POSIX.1 must have unique values so that they may be named as parameters of case statements in the body of a C language switch clause. However, implementation-defined signals may have values that overlap with each other or with signals specified in this document. An example of this is SIGABRT, which traditionally overlaps some other signal, such as SIGIOT.

2168 SIGKILL, SIGTERM, SIGUSR1, and SIGUSR2 are ordinarily generated only through 2169 the explicit use of the *kill()* function, although some implementations generate 2170 SIGKILL under extraordinary circumstances. SIGTERM is traditionally the 2171 default signal sent by the kill command.

The signals SIGBUS, SIGEMT, SIGIOT, SIGTRAP, and SIGSYS were omitted from 2172 POSIX.1 because their behavior is implementation dependent and could not be 2173 adequately categorized. Conforming implementations may deliver these signals, 2174 but must document the circumstances under which they are delivered and note 2175any restrictions concerning their delivery. The signals SIGFPE, SIGILL, and SIG-2176 SEGV are similar in that they also generally result only from programming errors. 2177 They were included in POSIX.1 because they do indicate three relatively well-2178 categorized conditions. They are all defined by the C Standard {2} and thus would 2179 have to be defined by any system with a C Standard {2} binding, even if not expli-2180 citly included in POSIX.1. 2181

There is very little that a Conforming POSIX.1 Application can do by catching, 2182 ignoring, or masking any of the signals SIGILL, SIGTRAP, SIGIOT, SIGEMT, 2183 SIGBUS, SIGSEGV, SIGSYS, or SIGFPE. They will generally be generated by the 2184 system only in cases of programming errors. While it may be desirable for some 2185 robust code (e.g., a library routine) to be able to detect and recover from program-2186 ming errors in other code, these signals are not nearly sufficient for that purpose. 2187 One portable use that does exist for these signals is that a command interpreter 2188 can recognize them as the cause of a process's termination [with wait()] and print 2189 an appropriate message. The mnemonic tags for these signals are derived from 2190 their PDP-11 origin. 2191

The signals SIGSTOP, SIGTSTP, SIGTTIN, SIGTTOU, and SIGCONT are provided for job control and are unchanged from 4.2BSD. The signal SIGCHLD is also typically used by job control shells to detect children that have terminated or, as in 4.2BSD, stopped. See also B.3.3.4.

Some implementations, including System V, have a signal named SIGCLD, which 2196 is similar to SIGCHLD in 4.2BSD. POSIX.1 permits implementations to have a sin-2197 gle signal with both names. POSIX.1 carefully specifies ways in which portable 2198 applications can avoid the semantic differences between the two different imple-2199 mentations. The name SIGCHLD was chosen for POSIX.1 because most current 2200 application usages of it can remain unchanged in conforming applications. 2201 SIGCLD in System V has more cases of semantics that POSIX.1 does not specify, 2202 and thus applications using it are more likely to require changes in addition to 2203 the name change. 2204

2205 Some implementations that do not support job control may nonetheless implement SIGCHLD. Similarly, such an implementation may choose to implement SIG-2206 STOP. Since POSIX.1 requires that symbolic names always be defined (with the 2207 exception of certain names in <limits.h> and <unistd.h>), a portable method 2208 of determining, at run-time, whether an optional signal is supported is to call the 2209 sigaction() function with NULL act and oact arguments. A successful return indi-2210 cates that the signal is supported. Note that if sysconf() shows that job control is 2211 present, then all of the optional signals shall also be supported. 2212

The signals SIGUSR1 and SIGUSR2 are commonly used by applications for 2213 notification of exceptional behavior and are described as "reserved as application 2214 defined" so that such use is not prohibited. Implementations should not generate 2215 SIGUSR1 or SIGUSR2, except when explicitly requested by kill(). It is recom-2216 mended that libraries not use these two signals, as such use in libraries could 2217 interfere with their use by applications calling the libraries. If such use is una-2218 voidable, it should be documented. It is prudent for nonportable libraries to use 2219 2220 nonstandard signals to avoid conflicts with use of standard signals by portable libraries. 2221

There is no portable way for an application to catch or ignore nonstandard sig-2222 nals. Some implementations define the range of signal numbers, so applications 2223 signal-catching functions for all of them. can install Unfortunately, 2224 implementation-defined signals often cause problems when caught or ignored by 2225 applications that do not understand the reason for the signal. While the desire 2226 exists for an application to be more robust by handling all possible signals [even 2227 those only generated by *kill()*, no existing mechanism was found to be sufficiently 2228 portable to include in POSIX.1. The value of such a mechanism, if included, would 2229

2230 be diminished given that SIGKILL would still not be catchable.

2231 B.3.3.1.2 Signal Generation and Delivery

The terms defined in this subclause are not used consistently in documentation of historical systems. Each signal can be considered to have a lifetime beginning with *generation* and ending with *delivery*. The POSIX.1 definition of *delivery* does not exclude ignored signals; this is considered a more consistent definition.

Implementations should deliver unblocked signals as soon after they are generated as possible. However, it is difficult for POSIX.1 to make specific requirements about this, beyond those in *kill()* and *sigprocmask()*. Even on systems with prompt delivery, scheduling of higher priority processes is always likely to cause delays.

In general, the interval between the generation and delivery of unblocked signals cannot be detected by an application. Thus, references to pending signals generally apply to blocked, pending signals.

In the 4.3BSD system, signals that are blocked and set to SIG_IGN are discarded 2244 immediately upon generation. For a signal that is ignored as its default action, if 2245 the action is SIG_DFL and the signal is blocked, a generated signal remains pend-2246 ing. In the 4.1BSD system and in System V Release 3, two other implementations 2247 that support a somewhat similar signal mechanism, all ignored, blocked signals 2248 remain pending if generated. Because it is not normally useful for an application 2249 to simultaneously ignore and block the same signal, it was unnecessary for 2250 POSIX.1 to specify behavior that would invalidate any of the historical 2251 implementations. 2252

There is one case in some historical implementations where an unblocked, pend-2253 ing signal does not remain pending until it is delivered. In the System V imple-2254 mentation of signal(), pending signals are discarded when the action is set to 2255SIG_DFL or a signal-catching routine (as well as to SIG_IGN). Except in the case 2256 of setting SIGCHLD to SIG_DFL, implementations that do this do not conform com-2257 pletely to POSIX.1. Some earlier drafts of POSIX.1 explicitly stated this, but these 2258 statements were redundant due to the requirement that functions defined by 2259 POSIX.1 not change attributes of processes defined by POSIX.1 except as explicitly 2260 stated (see Section 3). 2261

POSIX.1 specifically states that the order in which multiple, simultaneously pending signals are delivered is unspecified. This order has not been explicitly specified in historical implementations, but has remained quite consistent and been known to those familiar with the implementations. Thus, there have been cases where applications (usually system utilities) have been written with explicit or implicit dependencies on this order. Implementors and others porting existing applications may need to be aware of such dependencies.

When there are multiple pending signals that are not blocked, implementations should arrange for the delivery of all signals at once, if possible. Some implementations stack calls to all pending signal-catching routines, making it appear that each signal-catcher was interrupted by the next signal. In this case, the implementation should ensure that this stacking of signals does not violate the semantics of the signal masks established by *sigaction()*. Other implementations process at most one signal when the operating system is entered, with remaining

I

signals saved for later delivery. Although this practice is widespread, this
behavior is neither standardized nor endorsed. In either case, implementations
should attempt to deliver signals associated with the current state of the process
(e.g., SIGFPE) before other signals, if possible.

In 4.2BSD and 4.3BSD, it is not permissible to ignore or explicitly block SIGCONT 2280 because if blocking or ignoring this signal prevented it from continuing a stopped 2281 process, such a process could never be continued (only killed by SIGKILL). How-22.82 ever, 4.2BSD and 4.3BSD do block SIGCONT during execution of its signal-catching 2283 function when it is caught, creating exactly this problem. A proposal was con-2284 sidered to disallow catching SIGCONT in addition to ignoring and blocking it, but 2285 this limitation led to objections. The consensus was to require that SIGCONT 2286 always continue a stopped process when generated. This removed the need to 2287 disallow ignoring or explicit blocking of the signal; note that SIG_IGN and 2288 SIG_DFL are equivalent for SIGCONT. 2289

- 2290 B.3.3.1.3 Signal Actions
- Earlier drafts of POSIX.1 mentioned SIGCONT as a second exception to the rule that signals are not delivered to stopped processes until continued. Because POSIX.1 now specifies that SIGCONT causes the stopped process to continue when it is generated, delivery of SIGCONT is not prevented because a process is stopped, even without an explicit exception to this rule.
- Ignoring a signal by setting the action to SIG_IGN (or SIG_DFL for signals whose default action is to ignore) is not the same as installing a signal-catching function that simply returns. Invoking such a function will interrupt certain system functions that block processes [e.g., *wait()*, *sigsuspend()*, *pause()*, *read()*, *write()*] while ignoring a signal has no such effect on the process.
- Historical implementations discard pending signals when the action is set to SIG_IGN. However, they do not always do the same when the action is set to SIG_DFL and the default action is to ignore the signal. POSIX.1 requires this for the sake of consistency and also for completeness, since the only signal this applies to is SIGCHLD, and POSIX.1 disallows setting its action to SIG_IGN.
- The specification of the effects of SIG_IGN on SIGCHLD as implementation defined permits, but does not require, the System V effect of causing terminating children to be ignored by *wait()*. Yet it permits SIGCHLD to be effectively ignored in an implementation-independent manner by use of SIG_DFL.
- Some implementations (System V, for example) assign different semantics for SIGCLD depending on whether the action is set to SIG_IGN or SIG_DFL. Since POSIX.1 requires that the default action for SIGCHLD be to ignore the signal, applications should always set the action to SIG_DFL in order to avoid SIGCHLD.
- Some implementations (System V, for example) will deliver a SIGCLD signal 2314 immediately when a process establishes a signal-catching function for SIGCLD 2315 when that process has a child that has already terminated. Other implementa-2316 tions, such as 4.3BSD, do not generate a new SIGCHLD signal in this way. In gen-2317 eral, a process should not attempt to alter the signal action for the SIGCHLD sig-2318nal while it has any outstanding children. However, it is not always possible for a 2319 process to avoid this; for example, shells sometimes start up processes in pipe-2320 lines with other processes from the pipeline as children. Processes that cannot 2321

ensure that they have no children when altering the signal action for SIGCHLD thus need to be prepared for, but not depend on, generation of an immediate SIGCHLD signal.

The default action of the stop signals (SIGSTOP, SIGTSTP, SIGTTIN, SIGTTOU) is to stop a process that is executing. If a stop signal is delivered to a process that is already stopped, it has no effect. In fact, if a stop signal is generated for a stopped process whose signal mask blocks the signal, the signal will never be delivered to the process since the process must receive a SIGCONT, which discards all pending stop signals, in order to continue executing.

The SIGCONT signal shall continue a stopped process even if SIGCONT is blocked (or ignored). However, if a signal-catching routine has been established for SIGCONT, it will not be entered until SIGCONT is unblocked.

If a process in an orphaned process group stops, it is no longer under the control 2334 of a job-control shell and hence would not normally ever be continued. Because of 2335 this, orphaned processes that receive terminal-related stop signals (SIGTSTP, 2336 SIGTTIN, SIGTTOU, but not SIGSTOP) must not be allowed to stop. The goal is to 2337 prevent stopped processes from languishing forever. [As SIGSTOP is sent only via 2338 kill(), it is assumed that the process or user sending a SIGSTOP can send a 2339 SIGCONT when desired.] Instead, the system must discard the stop signal. As an 2340 extension, it may also deliver another signal in its place. 4.3BSD sends a SIG-2341 KILL, which is overly effective because SIGKILL is not catchable. Another possible 2342 choice is SIGHUP. 4.3BSD also does this for orphaned processes (processes whose 2343 parent has terminated) rather than for members of orphaned process groups; this 2344 is less desirable because job-control shells manage process groups. POSIX.1 also 2345 prevents SIGTTIN and SIGTTOU signals from being generated for processes in 2346 orphaned process groups as a direct result of activity on a terminal, preventing 2347 infinite loops when read() and write() calls generate signals that are discarded. 2348(See B.7.1.1.4.) A similar restriction on the generation of SIGTSTP was con-2349 sidered, but that would be unnecessary and more difficult to implement due to its 2350 asynchronous nature. 2351

Although POSIX.1 requires that signal-catching functions be called with only one argument, there is nothing to prevent conforming implementations from extending POSIX.1 to pass additional arguments, as long as Strictly Conforming POSIX.1 Applications continue to compile and execute correctly. Most historical implementations do, in fact, pass additional, signal-specific arguments to certain signalcatching routines.

2358 There was a proposal to change the declared type of the signal handler to:

The usage of ellipses (", ...") is C Standard {2} syntax to indicate a variable number of arguments. Its use was intended to allow the implementation to pass additional information to the signal handler in a standard manner.

2363 Unfortunately, this construct would require all signal handlers to be defined with 2364 this syntax because the C Standard {2} allows implementations to use a different 2365 parameter passing mechanism for variable parameter lists than for nonvariable 2366 parameter lists. Thus, all existing signal handlers in all existing applications 2367 would have to be changed to use the variable syntax in order to be standard and

portable. This is in conflict with the goal of Minimal Changes to Existing Applica-tion Code.

When terminating a process from a signal-catching function, processes should be aware of any interpretation that their parent may make of the status returned by wait() or waitpid(). In particular, a signal-catching function should not call exit(0) or $_exit(0)$ unless it wants to indicate successful termination. A nonzero argument to exit() or $_exit()$ can be used to indicate unsuccessful termination. Alternatively, the process can use kill() to send itself a fatal signal (first ensuring that the signal is set to the default action and not blocked). (See also B.3.2.2).

- The behavior of *unsafe* functions, as defined by this subclause, is undefined when 2377 they are invoked from signal-catching functions in certain circumstances. The 2378 behavior of reentrant functions, as defined by this subclause, is as specified by 2379 POSIX.1, regardless of invocation from a signal-catching function. This is the only 2380 intended meaning of the statement that reentrant functions may be used in 2381 signal-catching functions without restriction. Applications must still consider all 2382 effects of such functions on such things as data structures, files, and process state. 2383 In particular, application writers need to consider the restrictions on interactions 2384 when interrupting *sleep()* [see *sleep()* and B.3.4.3] and interactions among multi-2385 ple handles for a file description (see 8.2.3 and B.8.2.3). The fact that any specific 2386 function is listed as reentrant does not necessarily mean that invocation of that 2387 function from a signal-catching function is recommended. 2388
- In order to prevent errors arising from interrupting nonreentrant function calls, 2389 applications should protect calls to these functions either by blocking the 2390 appropriate signals or through the use of some programmatic semaphore. 2391 POSIX.1 does not address the more general problem of synchronizing access to 2392 shared data structures. Note in particular that even the "safe" functions may 2393 modify the global variable errno; the signal-catching function may want to save 2394 and restore its value. The same principles apply to the reentrancy of application 2395 routines and asynchronous data access. 2396
- Note that longjmp() and siglongjmp() are not in the list of reentrant functions. 2397 This is because the code executing after $long_{imp}()$ or $siglong_{imp}()$ can call any 2398 unsafe functions with the same danger as calling those unsafe functions directly 2399 from the signal handler. Applications that use longimp() or siglongimp() out of 2400 signal handlers require rigorous protection in order to be portable. Many of the 2401 other functions that are excluded from the list are traditionally implemented 2402 using either the C language malloc() or free() functions or the C language stan-2403 dard I/O library, both of which traditionally use data structures in a nonreentrant 2404 manner. Because any combination of different functions using a common data 2405 structure can cause reentrancy problems, POSIX.1 does not define the behavior 2406 when any unsafe function is called in a signal handler that interrupts any unsafe 2407 function. 2408

2409 **B.3.3.1.4 Signal Effects on Other Functions**

The most common behavior of an interrupted function after a signal-catching | function returns is for the interrupted function to give an [EINTR] error. However, there are a number of specific exceptions, including *sleep()* and certain situations with *read()* and *write()*.

The historical implementations of many functions defined by POSIX.1 are not 2414 interruptible, but delay delivery of signals generated during their execution until 2415 after they complete. This is never a problem for functions that are guaranteed to 2416 complete in a short (imperceptible to a human) period of time. It is normally 2417 those functions that can suspend a process indefinitely or for long periods of time 2418 [e.g., wait(), pause(), sigsuspend(), sleep(), or read()/write() on a slow device like a 2419 terminal] that are interruptible. This permits applications to respond to interac-2420 tive signals or to set timeouts on calls to most such functions with alarm(). 2421 Therefore, implementations should generally make such functions (including ones 2422 defined as extensions) interruptible. 2423

Functions not mentioned explicitly as interruptible may be so on some implementations, possibly as an extension where the function gives an [EINTR] error. There are several functions [e.g., getpid(), getuid()] that are specified as never returning an error, which can thus never be extended in this way.

2428 B.3.3.2 Send a Signal to a Process

The semantics for permission checking for *kill()* differ between System V and most other implementations, such as Version 7 or 4.3BSD. The semantics chosen for POSIX.1 agree with System V. Specifically, a set-user-ID process cannot protect itself against signals (or at least not against SIGKILL) unless it changes its real user ID. This choice allows the user who starts an application to send it signals even if it changes its effective user ID. The other semantics give more power to an application that wants to protect itself from the user who ran it.

Some implementations provide semantic extensions to the *kill()* function when the absolute value of *pid* is greater than some maximum, or otherwise special, value. Negative values are a flag to *kill()*. Since most implementations return [ESRCH] in this case, this behavior is not included in POSIX.1, although a conforming implementation could provide such an extension.

The implementation-defined processes to which a signal cannot be sent may include the scheduler or init.

Most historical implementations use kill (-1, sig) from a super-user process to 2443 send a signal to all processes (excluding system processes like init). This use of 2444 the kill() function is for administrative purposes only; portable applications 2445 should not send signals to processes about which they have no knowledge. In 2446 addition, there are semantic variations among different implementations that, 2447 because of the limited use of this feature, were not necessary to resolve by stan-2448 System V implementations also use kill(-1, sig) from a dardization. 2449 nonsuper-user process to send a signal to all processes with matching user IDs. 2450 This use was considered neither sufficiently widespread nor necessary for applica-2451 tion portability to warrant inclusion in POSIX.1. 2452

There was initially strong sentiment to specify that, if *pid* specifies that a signal be sent to the calling process and that signal is not blocked, that signal would be delivered before *kill()* returns. This would permit a process to call *kill()* and be guaranteed that the call never return. However, historical implementations that provide only the *signal()* interface make only the weaker guarantee in POSIX.1, because they only deliver one signal each time a process enters the kernel. Modifications to such implementations to support the *sigaction()* interface

generally require entry to the kernel following return from a signal-catching func-2460 tion, in order to restore the signal mask. Such modifications have the effect of 2461 satisfying the stronger requirement, at least when signation() is used, but not 2462 necessarily when *signal()* is used. The developers of POSIX.1 considered making 2463 the stronger requirement except when *signal()* is used, but felt this would be 2464 unnecessarily complex. Implementors are encouraged to meet the stronger 2465 requirement whenever possible. In practice, the weaker requirement is the same, 2466 except in the rare case when two signals arrive during a very short window. This 2467 reasoning also applies to a similar requirement for *sigprocmask()*. 2468

In 4.2BSD, the SIGCONT signal can be sent to any descendant process regardless 2469 of user-ID security checks. This allows a job-control shell to continue a job even if 2470 processes in the job have altered their user IDs (as in the su command). In keep-2471 ing with the addition of the concept of sessions, similar functionality is provided 2472 by allowing the SIGCONT signal to be sent to any process in the same session, 9473 regardless of user-ID security checks. This is less restrictive than BSD in the 2474 sense that ancestor processes (in the same session) can now be the recipient. It is 2475 more restrictive than BSD in the sense that descendant processes that form new 2476 sessions are now subject to the user-ID checks. A similar relaxation of security is 2477 not necessary for the other job-control signals since those signals are typically 2478 sent by the terminal driver in recognition of special characters being typed; the 2479 terminal driver bypasses all security checks. 2480

In secure implementations, a process may be restricted from sending a signal to a process having a different security label. In order to prevent the existence or nonexistence of a process from being used as a covert channel, such processes should appear nonexistent to the sender; i.e., [ESRCH] should be returned, rather than [EPERM], if *pid* refers only to such processes.

Existing implementations vary on the result of a kill() with pid indicating an 2486 inactive process (a terminated process that has not been waited for by its parent). 2487 Some indicate success on such a call (subject to permission checking), while others 2488 give an error of [ESRCH]. Since POSIX.1's definition of process lifetime covers inac-2489 tive processes, the [ESRCH] error as described is inappropriate in this case. In 2490 particular, this means that an application cannot have a parent process check for 2491 termination of a particular child with kill() [usually this is done with the null sig-2492 nal; this can be done reliably with *waitpid()*]. 2493

There is some belief that the name *kill()* is misleading, since the function is not always intended to cause process termination. However, the name is common to all historical implementations, and any change would be in conflict with the goal of Minimal Changes to Existing Application Code.

2498 B.3.3.3 Manipulate Signal Sets

The implementation of the *sigemptyset()* [or *sigfillset()*] functions could quite trivially clear (or set) all the bits in the signal set. Alternatively, it would be reasonable to initialize part of the structure, such as a version field, to permit binary compatibility between releases where the size of the set varies. For such reasons, either *sigemptyset()* or *sigfillset()* must be called prior to any other use of the signal set, even if such use is read-only [e.g., as an argument to *sigpending()*]. This function is not intended for dynamic allocation. 2506 The *sigfillset()* and *sigemptyset()* functions require that the resulting signal set include (or exclude) all the signals defined in POSIX.1. Although it is outside the 2507 scope of POSIX.1 to place this requirement on signals that are implemented as 2508 extensions, it is recommended that implementation-defined signals also be 2509 affected by these functions. However, there may be a good reason for a particular 2510 signal not to be affected. For example, blocking or ignoring an implementation-2511 defined signal may have undesirable side effects, whereas the default action for 2512 that signal is harmless. In such a case, it would be preferable for such a signal to 2513 be excluded from the signal set returned by *sigfillset()*. 2514

In earlier drafts of POSIX.1 there was no distinction between invalid and unsup-2515 ported signals (the names of optional signals that were not supported by an 2516 implementation were not defined by that implementation). The [EINVAL] error 2517 was thus specified as a required error for invalid signals. With that distinction, it 2518 is not necessary to require implementations of these functions to determine 2519 whether an optional signal is actually supported, as that could have a significant 2520 performance impact for little value. The error could have been required for 2521 invalid signals and optional for unsupported signals, but this seemed unneces-2522 sarily complex. Thus, the error is optional in both cases. 2523

2524 B.3.3.4 Examine and Change Signal Action

Although POSIX.1 requires that signals that cannot be ignored shall not be added to the signal mask when a signal-catching function is entered, there is no explicit requirement that subsequent calls to *sigaction()* reflect this in the information returned in the *oact* argument. In other words, if SIGKILL is included in the *sa_mask* field of *act*, it is unspecified whether or not a subsequent call to *sigaction()* will return with SIGKILL included in the *sa_mask* field of *oact*.

The SA_NOCLDSTOP flag, when supplied in the act->sa_flags parameter, allows 2531 overloading SIGCHLD with the System V semantics that each SIGCLD signal indi-2532 cates a single terminated child. Most portable applications that catch SIGCHLD 2533 are expected to install signal-catching functions that repeatedly call the waitpid() 2534 function with the WNOHANG flag set, acting on each child for which status is 2535 returned, until waitpid() returns zero. If stopped children are not of interest, the 2536 use of the SA NOCLDSTOP flag can prevent the overhead from invoking the 2537 signal-catching routine when they stop. 2538

Some historical implementations also define other mechanisms for stopping processes, such as the *ptrace()* function. These implementations usually do not generate a SIGCHLD signal when processes stop due to this mechanism; however, that is beyond the scope of POSIX.1.

POSIX.1 requires that calls to *sigaction()* that supply a NULL *act* argument succeed, even in the case of signals that cannot be caught or ignored (i.e., SIGKILL or SIGSTOP). The System V *signal()* and BSD *sigvec()* functions return [EINVAL] in these cases and, in this respect, their behavior varies from *sigaction()*.

POSIX.1 requires that sigaction() properly save and restore a signal action set up by the C Standard {2} signal() function. However, there is no guarantee that the reverse is true, nor could there be given the greater amount of information conveyed by the sigaction structure. Because of this, applications should avoid using both functions for the same signal in the same process. Since this cannot always

- ²⁵⁵² be avoided in case of general-purpose library routines, they should always be ²⁵⁵³ implemented with *sigaction()*.
- It was intended that the *signal()* function should be implementable as a library routine using *sigaction()*.

2556 **B.3.3.5 Examine and Change Blocked Signals**

When a process's signal mask is changed in a signal-catching function that is installed by *sigaction()*, the restoration of the signal mask on return from the signal-catching function overrides that change [see *sigaction()*]. If the signalcatching function was installed with *signal()*, it is unspecified whether this occurs.

- 2562 See B.3.3.2 for a discussion of the requirement on delivery of signals.
- 2563 **B.3.3.6 Examine Pending Signals**
- 2564 There is no additional rationale provided for this subclause.
- 2565 **B.3.3.7 Wait for a Signal**

Normally, at the beginning of a critical code section, a specified set of signals is blocked using the *sigprocmask()* function. When the process has completed the critical section and needs to wait for the previously blocked signal(s), it pauses by calling *sigsuspend()* with the mask that was returned by the *sigprocmask()* call.

2570 B.3.4 Timer Operations

2571 B.3.4.1 Schedule Alarm

Many historical implementations (including Version 7 and System V) allow an alarm to occur up to a second early. Other implementations allow alarms up to half a second or one clock tick early or do not allow them to occur early at all. The latter is considered most appropriate, since it gives the most predictable behavior, especially since the signal can always be delayed for an indefinite amount of time due to scheduling. Applications can thus choose the *seconds* argument as the minimum amount of time they wish to have elapse before the signal.

The term "real time" here and elsewhere [*sleep()*, *times()*] is intended to mean "wall clock" time as common English usage, and has nothing to do with "realtime operating systems." It is in contrast to "virtual time," which could be misinterpreted if just "time" were used.

In some implementations, including 4.3BSD, very large values of the *seconds* argument are silently rounded down to an implementation-defined maximum value. This maximum is large enough (on the order of several months) that the effect is not noticeable.

Application writers should note that the type of the argument seconds and the return value of alarm() is unsigned int. That means that a Strictly Conforming POSIX.1 Application cannot pass a value greater than the minimum guaranteed

I

I

value for {UINT_MAX}, which the C Standard {2} sets as 65535, and any application passing a larger value is restricting its portability. A different type was considered, but historical implementations, including those with a 16-bit *int* type, consistently use either *unsigned int* or *int*.

Application writers should be aware of possible interactions when the same process uses both the *alarm()* and *sleep()* functions [see *sleep()* and B.3.4.3].

2596 B.3.4.2 Suspend Process Execution

Many common uses of *pause()* have timing windows. The scenario involves checking a condition related to a signal and, if the signal has not occurred, calling *pause()*. When the signal occurs between the check and the call to *pause()*, the process often blocks indefinitely. The *sigprocmask()* and *sigsuspend()* functions can be used to avoid this type of problem.

2602 B.3.4.3 Delay Process Execution

There are two general approaches to the implementation of the *sleep()* function. One is to use the *alarm()* function to schedule a SIGALRM signal and then suspend the process waiting for that signal. The other is to implement an independent facility. POSIX.1 permits either approach.

In order to comply with the wording of the introduction to Section 3, that no prim-2607 itive shall change a process attribute unless explicitly described by POSIX.1, an 2608 implementation using SIGALRM must carefully take into account any SIGALRM 2609 signal scheduled by previous *alarm()* calls, the action previously established for 2610SIGALRM, and whether SIGALRM was blocked. If a SIGALRM has been scheduled 2611before the *sleep()* would ordinarily complete, the *sleep()* must be shortened to that 2612time and a SIGALRM generated (possibly simulated by direct invocation of the 2613 signal-catching function) before *sleep()* returns. If a SIGALRM has been scheduled 2614 after the *sleep()* would ordinarily complete, it must be rescheduled for the same 2615 time before *sleep()* returns. The action and blocking for SIGALRM must be saved 2616 and restored. 2617

Historical implementations often implement the SIGALRM-based version using 2618 alarm() and pause(). One such implementation is prone to infinite hangups, as 2619 described in B.3.4.2. Another such implementation uses the C language $set_{jmp}()$ 2620 and longimp() functions to avoid that window. That implementation introduces a 2621 different problem: when the SIGALRM signal interrupts a signal-catching function 2622 installed by the user to catch a different signal, the longjmp() aborts that signal-2623 catching function. An implementation based on sigprocmask(), alarm(), and sig-2624 suspend() can avoid these problems. 2625

Despite all reasonable care, there are several very subtle, but detectable and unavoidable, differences between the two types of implementations. These are the cases mentioned in POSIX.1 where some other activity relating to SIGALRM takes place, and the results are stated to be unspecified. All of these cases are sufficiently unusual as not to be of concern to most applications.

2631 (See also the discussion of the term "real time" in B.3.4.1.)

1

Because *sleep()* can be implemented using *alarm()*, the discussion about alarms occurring early under B.3.4.1 applies to *sleep()* as well.

Application writers should note that the type of the argument seconds and the return value of sleep() is unsigned int. That means that a Strictly Conforming POSIX.1 Application cannot pass a value greater than the minimum guaranteed value for {UINT_MAX}, which the C Standard {2} sets as 65535, and any application passing a larger value is restricting its portability. A different type was considered, but historical implementations, including those with a 16-bit *int* type, consistently use either *unsigned int* or *int*.

Scheduling delays may cause the process to return from the *sleep()* function significantly after the requested time. In such cases, the return value should be set to zero, since the formula (requested time minus the time actually spent) yields a negative number and *sleep()* returns an *unsigned int*.

- 2645 **B.4 Process Environment**
- 2646 B.4.1 Process Identification
- 2647 B.4.1.1 Get Process and Parent Process IDs
- 2648 There is no additional rationale provided for this subclause.
- 2649 B.4.2 User Identification
- B.4.2.1 Get Real User, Effective User, Real Group, and Effective Group
 IDs
- 2652 There is no additional rationale provided for this subclause.
- 2653 B.4.2.2 Set User and Group IDs

The saved set-user-ID capability allows a program to regain the effective user ID established at the last *exec* call. Similarly, the saved set-group-ID capability allows a program to regain the effective group ID established at the last *exec* call.

These two capabilities are derived from System V. Without them, a program may have to run as super-user in order to perform the same functions, because superuser can write on the user's files. This is a problem because such a program can write on *any* user's files, and so must be carefully written to emulate the permissions of the calling process properly.

A process with appropriate privilege on a system with this saved ID capability establishes all relevant IDs to the new value, since this function is used to establish the identity of the user during login or su. Any change to this behavior would be dangerous since it involves programs that need to be trusted. The behavior of 4.2BSD and 4.3BSD that allows setting the real ID to the effective ID is viewed as a value-dependent special case of appropriate privilege.

2668 B.4.2.3 Get Supplementary Group IDs

The related function *setgroups()* is a privileged operation and therefore is not covered by POSIX.1.

As implied by the definition of supplementary groups, the effective group ID may 2671 appear in the array returned by getgroups() or it may be returned only by 2672 getegid(). Duplication may exist, but the application needs to call getegid() to be 2673 sure of getting all of the information. Various implementation variations and 2674 administrative sequences will cause the set of groups appearing in the result of 2675 getgroups() to vary in order and as to whether the effective group ID is included, 2676 even when the set of groups is the same (in the mathematical sense of "set"). (The 2677 history of a process and its parents could affect the details of result.) 2678

Applications writers should note that {NGROUPS_MAX} is not necessarily a constant on all implementations.

2681 B.4.2.4 Get User Name

The getlogin() function returns a pointer to the user's login name. The same user ID may be shared by several login names. If it is desired to get the user database entry that is used during login, the result of getlogin() should be used to provide the argument to the getpwnam() function. (This might be used to determine the user's login shell, particularly where a single user has multiple login shells with distinct login names, but the same user ID.)

The information provided by the *cuserid()* function, which was originally defined in IEEE Std 1003.1-1990 and subsequently removed, can be obtained by the following:

2691 getpwuid(geteuid())

while the information provided by historical implementations of *cuserid()* can be obtained by:

2694 getpwuid(getuid())

2695 **B.4.3 Process Groups**

2696 B.4.3.1 Get Process Group ID

4.3BSD provides a getpgrp() function that returns the process group ID for a specified process. Although this function is used to support job control, all known job-control shells always specify the calling process with this function. Thus, the simpler System V getpgrp() suffices, and the added complexity of the 4.3BSD getpgrp() has been omitted from POSIX.1.

2702 B.4.3.2 Create Session and Set Process Group ID

The *setsid()* function is similar to the *setpgrp()* function of System V. System V, without job control, groups processes into process groups and creates new process groups via *setpgrp()*; only one process group may be part of a login session.

Job control allows multiple process groups within a login session. In order to 2706 limit job-control actions so that they can only affect processes in the same login 2707 session, POSIX.1 adds the concept of a session that is created via setsid(). The set-2708 sid() function also creates the initial process group contained in the session. 2709 Additional process groups can be created via the setpgid() function. A System V 2710 process group would correspond to a POSIX.1 session containing a single POSIX.1 2711 process group. Note that this function requires that the calling process not be a 2712 process group leader. The usual way to ensure this is true is to create a new pro-2713 cess with fork() and have it call setsid(). The fork() function guarantees that the 2714 process ID of the new process does not match any existing process group ID. 2715

2716 B.4.3.3 Set Process Group ID for Job Control

The *setpgid()* function is used to group processes together for the purpose of signaling, placement in foreground or background, and other job-control actions. See B.2.2.2.

- The *setpgid()* function is similar to the *setpgrp()* function of 4.2BSD, except that 4.2BSD allowed the specified new process group to assume any value. This presents certain security problems and is more flexible than necessary to support job control.
- To provide tighter security, *setpgid()* only allows the calling process to join a process group already in use inside its session or create a new process group whose process group ID was equal to its process ID.
- When a job-control shell spawns a new job, the processes in the job must be placed into a new process group via *setpgid()*. There are two timing constraints involved in this action:
- (1) The new process must be placed in the new process group before the appropriate program is launched via one of the *exec* functions.
- (2) The new process must be placed in the new process group before the shell
 can correctly send signals to the new process group.

To address these constraints, the following actions are performed: The new 2734 processes call setpgid() to alter their own process groups after fork() but before 2735 exec. This satisfies the first constraint. Under 4.3BSD, the second constraint is 2736 satisfied by the synchronization property of vfork(); that is, the shell is suspended 2737 until the child has completed the *exec*, thus ensuring that the child has completed 2738 the set pgid(). A new version of fork() with this same synchronization property 2739 was considered, but it was decided instead to merely allow the parent shell pro-2740 cess to adjust the process group of its child processes via setpgid(). Both timing 2741constraints are now satisfied by having both the parent shell and the child 2742 attempt to adjust the process group of the child process; it does not matter which 2743 succeeds first. 2744

Because it would be confusing to an application to have its process group change 2745 2746 after it began executing (i.e., after exec) and because the child process would already have adjusted its process group before this, the [EACCES] error was added 2747 to disallow this. 2748

One nonobvious use of *setpgid()* is to allow a job-control shell to return itself to its 2749 original process group (the one in effect when the job-control shell was executed). 2750 A job-control shell does this before returning control back to its parent when it is 2751 terminating or suspending itself as a way of restoring its job control "state" back 2752 to what its parent would expect. (Note that the original process group of the job-2753 control shell typically matches the process group of its parent, but this is not 2754 necessarily always the case.) See also B.7.2.4. 2755

B.4.4 System Identification 2756

B.4.4.1 System Name 2757

The values of the structure members are not constrained to have any relation to 2758 the version of POSIX.1 implemented in the operating system. An application 2759 should instead depend on {_POSIX_VERSION} and related constants defined in 2.9. 2760

POSIX.1 does not define the sizes of the members of the structure and permits 2761 them to be of different sizes, although most implementations define them all to be 2762 the same size: eight bytes plus one byte for the string terminator. That size for 2763 nodename is not enough for use with many networks. 2764

The uname() function is specific to System III, System V, and related implementa-2765 tions, and it does not exist in Version 7 or 4.3BSD. The values it returns are set 2766 at system compile time in those historical implementations. 2767

4.3BSD has gethostname() and gethostid(), which return a symbolic name and a 2768 numeric value, respectively. There are related sethostname() and sethostid() 2769 functions that are used to set the values the other two functions return. The 2770 length of the host name is limited to 31 characters in most implementations and 2771 the host ID is a 32-bit integer. 2772

B.4.5 Time 2773

The *time()* function returns a value in seconds (type *time_t*) while *times()* returns 2774 a set of values in clock ticks (type *clock_t*). Some historical implementations, such 1 2775 as 4.3BSD, have mechanisms capable of returning more precise times [see the 2776 description of gettimeofday() in B.4.5.1]. A generalized timing scheme to unify 2777 these various timing mechanisms has been proposed but not adopted in POSIX.1. 2778

B.4.5.1 Get System Time 2779

Implementations in which *time_t* is a 32-bit signed integer (most historical imple-2780 mentations) will fail in the year 2038. This version of POSIX.1 does not address 2781 this problem. However, the use of the new time_t type is mandated in order to 2782 ease the eventual fix. 2783

Ĺ

I

The use of the header <time.h>, instead of <sys/types.h>, allows compatibility with the C Standard {2}.

Many historical implementations (including Version 7) and the 1984 /usr/group Standard (B59) use long instead of time_t. POSIX.1 uses the latter type in order to agree with the C Standard {2}.

4.3BSD includes *time()* only as an interface to the more flexible *gettimeofday()* function.

2791 B.4.5.2 Get Process Times

The accuracy of the times reported is intentionally left unspecified to allow implementations flexibility in design, from uniprocessor to multiprocessor networks.

The inclusion of times of child processes is recursive, so that a parent process may collect the total times of all of its descendants. But the times of a child are only added to those of its parent when its parent successfully waits on the child. Thus, it is not guaranteed that a parent process will always be able to see the total times of all its descendants.

2799 (See also the discussion of the term "real time" in B.3.4.1.)

If the type *clock_t* is defined to be a signed 32-bit integer, it will overflow in somewhat more than a year if there are 60 clock ticks per second, or less than a year if | there are 100. There are individual systems that run continuously for longer than that. POSIX.1 permits an implementation to make the reference point for the returned value be the startup time of the process, rather than system startup time.

The term "charge" in this context has nothing to do with billing for services. The operating system accounts for time used in this way. That information must be correct, regardless of how that information is used.

2809 **B.4.6 Environment Variables**

2810 B.4.6.1 Environment Access

Additional functions *putenv()* and *clearenv()* were considered but rejected because they were considered to be more oriented towards system administration than ordinary application programs. This is being reconsidered for an amendment to POSIX.1 because uses from within an application have been identified since the decision was made.

It was proposed that this function is properly part of Section 8. It is an extension to a function in the C Standard {2}. Because this function should be available from any language, not just C, it appears here, to separate it from the material in Section 8, which is specific to the C binding. (The localization extensions to C are not, at this time, appropriate for other languages.)

I

2821 B.4.7 Terminal Identification

The difference between ctermid() and ttyname() is that ttyname() must be passed a file descriptor and returns the pathname of the terminal associated with that file descriptor, while ctermid() returns a string (such as /dev/tty) that will refer to the controlling terminal if used as a pathname. Thus ttyname() is useful only if the process already has at least one file open to a terminal.

The historical value of *ctermid()* is /dev/tty; this is acceptable. The *ctermid()* function should not be used to determine if a process actually has a controlling terminal, but merely the name that would be used.

2830 B.4.7.1 Generate Terminal Pathname

L_ctermid must be defined appropriately for a given implementation and must be greater than zero so that array declarations using it are accepted by the compiler. The value includes the terminating null byte.

2834 B.4.7.2 Determine Terminal Device Name

The term "terminal" is used instead of the historical term "terminal device" in order to avoid a reference to an undefined term.

2837 B.4.8 Configurable System Variables

This subclause was added in response to requirements of application developers
and of system vendors who deal with many international system configurations.
It is closely related to B.5.7 as well.

Although a portable application can run on all systems by never demanding more resources than the minimum values published in POSIX.1, it is useful for that application to be able to use the actual value for the quantity of a resource available on any given system. To do this, the application will make use of the value of a symbolic constant in <limits.h> or <unistd.h>.

- However, once compiled, the application must still be able to cope if the amount of resource available is increased. To that end, an application may need a means of determining the quantity of a resource, or the presence of an option, at execution time.
- 2850 Two examples are offered:
- (1) Applications may wish to act differently on systems with or without job control. Applications vendors who wish to distribute only a single binary package to all instances of a computer architecture would be forced to assume job control is never available if it were to rely solely on the
 2854
 2855
- (2) International applications vendors occasionally require knowledge of the number of clock ticks per second. Without the facilities of this subclause, they would be required to either distribute their applications partially in source form or to have 50 Hz and 60 Hz versions for the various countries in which they operate.

l

ł

It is the knowledge that many applications are actually distributed widely in executable form that lead to this facility. If limited to the most restrictive values in the headers, such applications would have to be prepared to accept the most limited environments offered by the smallest microcomputers. Although this is entirely portable, there was a consensus that they should be able to take advantage of the facilities offered by large systems, without the restrictions associated with source and object distributions.

During the discussions of this feature, it was pointed out that it is almost always possible for an application to discern what a value might be at run-time by suitably testing the various interfaces themselves. And, in any event, it could always be written to adequately deal with error returns from the various functions. In the end, it was felt that this imposed an unreasonable level of complication and sophistication on the application writer.

This run-time facility is not meant to provide ever-changing values that applications will have to check multiple times. The values are seen as changing no more frequently than once per system initialization, such as by a system administrator or operator with an automatic configuration program. POSIX.1 specifies that they shall not change within the lifetime of the process.

2879 Some values apply to the system overall and others vary at the file system or 2880 directory level. These latter are described in B.5.7.

2881 B.4.8.1 Get Configurable System Variables

Note that all values returned must be expressible as integers. String values were
considered, but the additional flexibility of this approach was rejected due to its
added complexity of implementation and use.

Some values, such as {PATH_MAX}, are sometimes so large that they must not be used to, say, allocate arrays. The *sysconf()* function will return a negative value to show that this symbolic constant is not even defined in this case.

2888 B.4.8.1.1 Special Symbol {CLK_TCK}

2889 {CLK_TCK} appears in POSIX.1 for backwards compatibility with IEEE Std 2890 1003.1-1988. Its use is obsolescent.

2891 B.4.8.2 Get Password From User

The *getpass*() function was explicitly excluded from POSIX.1 because it was found that the name was misleading, and it provided no functionality that the user could not easily implement within POSIX.1. The implication of some form of security, which was not actually provided, exceeded the small gain in convenience.

2896 **B.5 Files and Directories**

2897 See pathname resolution.

The wording regarding the group of a newly created regular file, directory, or 2898 FIFO in open(), mkdir(), mkfifo(), respectively, defines the two acceptable 2899 behaviors in order to permit both the System V (and Version 7) behavior (in which 2900 the group of the new object is set to the effective group ID of the creating process) 2901 and the 4.3BSD behavior (in which the new object has the group of its parent 2902 directory). An application that needs a file to be created specifically in one or the 2903 other of the possible groups should use *chown()* to ensure the new group regard-2904 less of the style of groups the interface implements. Most applications will not 2905 and should not be concerned with the group ID of the file. 2906

2907 **B.5.1 Directories**

Historical implementations prior to 4.2BSD had no special functions, types, or 2908 headers for directory access. Instead, directories were read with read() and each 2909 program that did so had code to understand the internal format of directory files. 2910 Many such programs did not correctly handle the case of a maximum-length (his-2911 torically fourteen character) filename and would neglect to add a null character 2912 string terminator when doing comparisons. The access methods in POSIX.1 elim-2913 inate that bug, as well as hiding differences in implementations of directories or 2914 file systems. 2915

The directory access functions originally selected for POSIX.1 were derived from 2916 4.2BSD, were adopted in System V Release 3, and are in SVID (B39) Volume 3, 2917 with the exception of a type difference for the *d* ino field. That field represents 2918 implementation-dependent or even file system-dependent information (the i-node 2919 number in most implementations). Since the directory access mechanism is 2920 intended to be implementation-independent, and since only system programs, not 2921 ordinary applications, need to know about the i-node number (or file serial 2922 number) in this context, the d_ino field does not appear in POSIX.1. Also, pro-2923 2924 grams that want this information can get it with stat().

2925 **B.5.1.1 Format of Directory Entries**

Information similar to that in the header <dirent.h> is contained in a file 2926 <sys/dir.h> in 4.2BSD and 4.3BSD. The equivalent in these implementations 2927 of struct direct from POSIX.1 is struct direct. The filename was changed because 2928 the name <sys/dir.h> was also used in earlier implementations to refer to 2929 definitions related to the older access method; this produced name conflicts. The 2930 name of the structure was changed because POSIX.1 does not completely define 2931 what is in the structure, so it could be different on some implementations from 2932 struct direct. 2933

- The name of an array of *char* of an unspecified size should not be used as an *lvalue*. Use of
- 2936 sizeof (d_name)
- 2937 is incorrect; use

2938 strlen (d_name)

2939 instead.

The array of char d_name is not a fixed size. Implementations may need to declare struct dirent with an array size for d_name of 1, but the actual number of characters provided matches (or only slightly exceeds) the length of the file name.

2943 Currently, implementations are excluded if they have *d_name* with type *char* *. 2944 Lacking experience of such implementations, the developers of POSIX.1 declined 2945 to try to describe in standards language what to do if either type were permitted.

2946 **B.5.1.2 Directory Operations**

Based on historical implementations, the rules about file descriptors apply to directory streams as well. However, POSIX.1 does not mandate that the directory stream be implemented using file descriptors. The description of *opendir()* clarifies that if a file descriptor is used for the directory stream it is mandatory that *closedir()* deallocate the file descriptor. When a file descriptor is used to implement the directory stream, it behaves as if the FD_CLOEXEC had been set for the file descriptor.

- The returned value of *readdir()* merely *represents* a directory entry. No equivalence should be inferred.
- 2956 The directory entries for dot and dot-dot are optional. POSIX.1 does not provide a way to test a priori for their existence because an application that is portable 2957 must be written to look for (and usually ignore) those entries. Writing code that 2958 presumes that they are the first two entries does not always work, as many imple-2959 mentations permit them to be other than the first two entries, with a "normal" 2960 entry preceding them. There is negligible value in providing a way to determine 2961 what the implementation does because the code to deal with dot and dot-dot must 2962 be written in any case and because such a flag would add to the list of those flags 2963 (which has proven in itself to be objectionable) and might be abused. 2964
- Since the structure and buffer allocation, if any, for directory operations are defined by the implementation, POSIX.1 imposes no portability requirements for erroneous program constructs, erroneous data, or the use of indeterminate values such as the use or referencing of a *dirp* value or a *dirent* structure value after a directory stream has been closed or after a *fork()* or one of the *exec* function calls.
- Historical implementations of readdir() obtain multiple directory entries on a single read operation, which permits subsequent readdir() operations to operate from the buffered information. Any wording that required each successful readdir() operation to mark the directory st_atime field for update would militate against the historical performance-oriented implementations.
- 2975 Since *readdir()* returns **NULL** both:
- 2976 (1) When it detects an error, and
- 2977 (2) When the end of the directory is encountered

an application that needs to tell the difference must set *errno* to zero before the call and check it if **NULL** is returned. Because the function must not change *errno* in case (2) and must set it to a nonzero value in case (1), a zero *errno* after a

- 2981 call returning NULL indicates end of directory, otherwise an error.
- 2982 Routines to deal with this problem more directly were proposed:

2983 2984	<pre>int derror (dirp DIR *dirp;</pre>))
2985 2986	<pre>void clearderr DIR *dirp;</pre>	(dirp)

The first would indicate whether an error had occurred, and the second would clear the error indication. The simpler method involving *errno* was adopted instead by requiring that *readdir()* not change *errno* when end-of-directory is encountered.

2991 Historical implementations include two more functions:

2992 2993	<pre>long telldir DIR *dirp;</pre>	(dirp)	
2994	void seekdir	(dirp,	loc)
2995	DIR <i>*dirp;</i>		
2996	long loc;		

The *telldir()* function returns the current location associated with the named directory stream.

The *seekdir()* function sets the position of the next *readdir()* operation on the directory stream. The new position reverts to the one associated with the directory stream when the *telldir()* operation was performed.

These functions have restrictions on their use related to implementation details. 3002 Their capability can usually be accomplished by saving a filename found by read-3003 dir() and later using rewinddir() and a loop on readdir() to relocate the position 3004 from which the filename was saved. Though this method is probably slower than 3005 using *seekdir()* and *telldir()*, there are few applications in which the capability is 3006 needed. Furthermore, directory systems that are implemented using technology 3007 such as balanced trees, where the order of presentation may vary from access to 3008 access, do not lend themselves well to any concept along these lines. For these 3009 reasons, *seekdir()* and *telldir()* are not included in POSIX.1. 3010

3011 An error or signal indicating that a directory has changed while open was con-3012 sidered but rejected.

3013 B.5.1.3 Set Position of Directory Stream

The *seekdir()* and *telldir()* functions were proposed for inclusion in POSIX.1, but were excluded because they are inherently unreliable when all the possible conforming implementations of the rest of POSIX.1 were considered. The problem is that returning to a given point in a directory is quite difficult to describe formally, in spite of its intuitive appeal, when systems that used B-trees, hashing functions, or other similar mechanisms for directory search are considered.

Even the simple goal of attempting to visit each directory entry that is unmodified
between the opendir() and closedir() calls exactly once is difficult to implement
reliably in the face of directory compaction and reorganization.

Since the primary need for *seekdir()* and *telldir()* is to implement file tree walks, and since such a function is likely to be included in a future revision of POSIX.1, and since in that more constrained context it appears that at least the goal of visiting unmodified nodes exactly once can be achieved, it was felt that waiting for the development of that function best served all the constituencies.

3028 B.5.2 Working Directory

3029 B.5.2.1 Change Current Working Directory

3030 The *chdir()* function only affects the working directory of the current process.

The result if a NULL argument is passed to *chdir()* is left implementation defined because some implementations dynamically allocate space in that case.

3033 B.5.2.2 Working Directory Pathname

Since the maximum pathname length is arbitrary unless {PATH_MAX} is defined, an application generally cannot supply a *buf* with *size* {{PATH_MAX} + 1}.

- Having getcwd() take no arguments and instead use the C function malloc() to 3036 produce space for the returned argument was considered. The advantage is that 3037 getcwd() knows how big the working directory pathname is and can allocate an 3038 appropriate amount of space. But the programmer would have to use the C func-3039 tion free() to free the resulting object, or each use of getcwd() would further 3040 reduce the available memory. Also, malloc() and free() are used nowhere else in 3041 POSIX.1. Finally, getcwd() is taken from the SVID (B39), where it has the two 3042 arguments used in POSIX.1. 3043
- The older function *getwd()* was rejected for use in this context because it had only a buffer argument and no size argument, and thus had no way to prevent overwriting the buffer, except to depend on the programmer to provide a large enough buffer.
- The result if a NULL argument is passed to *getcwd()* is left implementation defined because some implementations dynamically allocate space in that case.

If a program is operating in a directory where some (grand)parent directory does 3050 not permit reading, getcwd() may fail, as in most implementations it must read 3051 the directory to determine the name of the file. This can occur if search, but not 3052 read, permission is granted in an intermediate directory, or if the program is 3053 placed in that directory by some more privileged process (e.g., login). Including 3054 this error, [EACCES], makes the reporting of the error consistent and warns the 3055 application writer that getcwd() can fail for reasons beyond the control of the 3056 application writer or user. Some implementations can avoid this occurrence [e.g., 3057 by implementing getcwd() using pwd, where pwd is a set-user-root process], thus 3058 the error was made optional. 3059

Because POSIX.1 permits the addition of other errors, this would be a common addition and yet one that applications could not be expected to deal with without this addition. Part 1: SYSTEM API [C LANGUAGE]

Some current implementations use {PATH_MAX}+2 bytes. These will have to be changed. Many of those same implementations also may not diagnose the [ERANGE] error properly or deal with a common bug having to do with newline in a directory name (the fix to which is essentially the same as the fix for using +1 bytes), so this is not a severe hardship.

3068 B.5.2.3 Change Process's Root Directory

The *chroot()* function was excluded from POSIX.1 on the basis that it was not useful to portable applications. In particular, creating an environment in which an application could run after executing a *chroot()* call is well beyond the current scope of POSIX.1.

3073 **B.5.3 General File Creation**

Because there is no portable way to specify a value for the argument indicating the file mode bits (except zero), <sys/stat.h> is included with the functions that reference mode bits.

3077 **B.5.3.1 Open a File**

- 3078 Except as specified in POSIX.1, the flags allowed in *oflag* are not mutually 3079 exclusive and any number of them may be used simultaneously.
- Some implementations permit opening FIFOs with O_RDWR. Since FIFOs could be implemented in other ways, and since two file descriptors can be used to the same effect, this possibility is left as undefined.
- 3083 See B.4.2.3 about the group of a newly created file.
- The use of *open()* to create a regular file is preferable to the use of *creat()* because the latter is redundant and included only for historical reasons.
- The use of the O_TRUNC flag on FIFOs and directories [pipes cannot be *open()*-ed] must be permissible without unexpected side effects [e.g., *creat()* on a FIFO must not remove data]. Because terminal special files might have type-ahead data stored in the buffer, O_TRUNC should not affect their content, particularly if a program that normally opens a regular file should open the current controlling terminal instead. Other file types, particularly implementation-defined ones, are left implementation defined.
- 3093 Implementations may deny access and return [EACCES] for reasons other than 3094 just those listed in the [EACCES] definition.
- The O_NOCTTY flag was added to allow applications to avoid unintentionally acquiring a controlling terminal as a side effect of opening a terminal file. POSIX.1 does not specify how a controlling terminal is acquired, but it allows an implementation to provide this on *open()* if the O_NOCTTY flag is not set and other conditions specified in 7.1.1.3 are met. The O_NOCTTY flag is an effective no-op if the file being opened is not a terminal device.
- In historical implementations the value of O_RDONLY is zero. Because of that, it is not possible to detect the presence of O_RDONLY and another option. Future

- 3103 implementations should encode O_RDONLY and O_WRONLY as bit flags so that:
- $O_RDONLY | O_WRONLY == O_RDWR$
- See the rationale for the change from O_NDELAY to O_NONBLOCK in B.6.

3106 **B.5.3.2 Create a New File or Rewrite an Existing One**

The *creat()* function is redundant. Its services are also provided by the *open()* function. It has been included primarily for historical purposes since many existing applications depend on it. It is best considered a part of the C binding rather than a function that should be provided in other languages.

3111 B.5.3.3 Set File Creation Mask

Unsigned argument and return types for umask() were proposed. The return type and the argument were both changed to $mode_t$.

Historical implementations have made use of additional bits in *cmask* for their implementation-specific purposes. The addition of the text that the meaning of other bits of the field are implementation defined permits these implementations to conform to POSIX.1.

- 3118 **B.5.3.4 Link to a File**
- 3119 See B.2.2.2.

Linking to a directory is restricted to the super-user in most historical implementations because this capability may produce loops in the file hierarchy or otherwise corrupt the file system. POSIX.1 continues that philosophy by prohibiting link() and unlink() from doing this. Other functions could do it if the implementor designed such an extension.

Some historical implementations allow linking of files on different file systems. Wording was added to explicitly allow this optional behavior. Symbolic links are not discussed by POSIX.1. The exception for cross-file system links is intended to apply only to links that are programmatically indistinguishable from "hard" links.

3129 **B.5.4 Special File Creation**

- 3130 B.5.4.1 Make a Directory
- 3131 See B.2.5.

The *mkdir()* function originated in 4.2BSD and was added to System V in Release 3.0.

- 3134 4.3BSD detects [ENAMETOOLONG].
- S135 See B.4.2.3 about the group of a newly created directory.

L

Part 1: SYSTEM API [C LANGUAGE]

1

3136 **B.5.4.2 Make a FIFO Special File**

The syntax of this routine is intended to maintain compatibility with historical implementations of mknod(). The latter function was included in the 1984 /usr/group Standard {B59}, but only for use in creating FIFO special files. The mknod() function was excluded from POSIX.1 as implementation defined and replaced by mkdir() and mkfifo().

3142 See B.4.2.3 about the group of a newly created FIFO.

3143 B.5.5 File Removal

The *rmdir()* and *rename()* functions originated in 4.2BSD, and they used 3144 [ENOTEMPTY] for the condition when the directory to be removed does not exist 3145 or new already exists. When the 1984 /usr/group Standard (B59) was published, 3146 it contained [EEXIST] instead. When these functions were adopted into System V, 3147 the 1984 /usr/group Standard (B59) was used as a reference. Therefore, several 3148 existing applications and implementations support/use both forms, and no agree-3149 ment could be reached on either value. All implementations are required to sup-3150 ply both [EEXIST] and [ENOTEMPTY] in <errno.h> with distinct values so that 3151 applications can use both values in C language case statements. 3152

3153 **B.5.5.1 Remove Directory Entries**

Unlinking a directory is restricted to the super-user in many historical implementations for reasons given in B.5.3.4. But see B.5.5.3.

The meaning of [EBUSY] in historical implementations is "mount point busy." Since POSIX.1 does not cover the system administration concepts of mounting and unmounting, the description of the error was changed to "resource busy." (This meaning is used by some device drivers when a second process tries to open an exclusive use device.) The wording is also intended to allow implementations to refuse to remove a directory if it is the root or current working directory of any process.

- 3163 **B.5.5.2 Remove a Directory**
- 3164 See also B.5.5 and B.5.5.1.

3165 **B.5.5.3 Rename a File**

This *rename()* function is equivalent for regular files to that defined by the C Standard {2}. Its inclusion here expands that definition to include actions on directories and specifies behavior when the *new* parameter names a file that already exists. That specification requires that the action of the function be atomic.

One of the reasons for introducing this function was to have a means of renaming directories while permitting implementations to prohibit the use of link() and unlink() with directories, thus constraining links to directories to those made by mkdir().

1

ł

- The specification that if *old* and *new* refer to the same file describes existing, although undocumented, 4.3BSD behavior. It is intended to guarantee that:
- 3177 rename("x", "x");

does not remove the file.

Renaming *dot* or *dot-dot* is prohibited in order to prevent cyclical file system paths.

See also the descriptions of [ENOTEMPTY] and [ENAMETOOLONG] in B.5.5 and [EBUSY] in B.5.5.1. For a discussion of [EXDEV], see B.5.3.4.

3183 **B.5.6 File Characteristics**

The ustat() function, which appeared in the 1984 /usr/group Standard (B59) and is still in the SVID (B39), was excluded from POSIX.1 because it is:

- Not reliable. The amount of space available can change between the time the call is made and the time the calling process attempts to use it.
- ³¹⁸⁸ Not required. The only known program that uses it is the text editor ed.
- 3189 Not readily extensible to networked systems.

3190 **B.5.6.1 File Characteristics: Header and Data Structure**

3191 See B.2.5.

A conforming C language application must include $\langle sys/stat.h \rangle$ for functions that have arguments or return values of type $mode_t$, so that symbolic values for that type can be used. An alternative would be to require that these constants are also defined by including $\langle sys/types.h \rangle$.

- The S_ISUID and S_ISGID bits may be cleared on any write, not just on *open()*, as some historical implementations do it.
- System calls that update the time entry fields in the *stat* structure must be documented by the implementors. POSIX.1 conforming systems should not update the time entry fields for functions listed in POSIX.1 unless the standard requires that they do, except in the case of documented extensions to the standard.
- Note that *st_dev* must be unique within a Local Area Network (LAN) in a "system" made up of multiple computers' file systems connected by a LAN.
- Networked implementations of a POSIX.1 system must guarantee that all files visible within the file tree (including parts of the tree that may be remotely mounted from other machines on the network) on each individual processor are uniquely identified by the combination of the st_ino and st_dev fields.

3208 B.5.6.2 Get File Status

The intent of the paragraph describing "additional or alternate file access control mechanisms" is to allow a secure implementation where a process with a label that does not dominate the file's label cannot perform a *stat()* function. This is not related to read permission; a process with a label that dominates the file's

I

label will not need read permission. An implementation that supports write-up
operations could fail *fstat()* function calls even though it has a valid file descriptor
open for writing.

3216 B.5.6.3 File Accessibility

In early drafts of POSIX.1, some inadequacies in the *access()* function led to the creation of an *eaccess()* function because:

- (1) Historical implementations of access() do not test file access correctly
 when the process's real user ID is super-user. In particular, they always
 return zero when testing execute permissions without regard to whether
 the file is executable.
- 3223 (2) The super-user has complete access to all files on a system. As a conse 3224 quence, programs started by the super-user and switched to the effective
 3225 user ID with lesser privileges cannot use access() to test their file access
 3226 permissions.

However, the historical model of eaccess() does not resolve problem (1), so POSIX.1 now allows access() to behave in the desired way because several implementations have corrected the problem. It was also argued that problem (2) is more easily solved by using open(), chdir(), or one of the *exec* functions as appropriate and responding to the error, rather than creating a new function that would not be as reliable. Therefore, eaccess() was taken back out of POSIX.1.

Secure implementations will probably need an extended *access()*-like function, but there were not enough of the requirements to define it yet. This could be proposed as an extension for a future amendment to POSIX.1.

The sentence concerning appropriate privileges and execute permission bits reflects the two possibilities implemented by historical implementations when checking super-user access for X_OK.

3239 B.5.6.4 Change File Modes

POSIX.1 specifies that the S_ISGID bit is cleared by *chmod()* on a regular file under certain conditions. This is specified on the assumption that regular files may be executed, and the system should prevent users from making executable *setgid* files perform with privileges that the caller does not have. On implementations that support execution of other file types, the S_ISGID bit should be cleared for those file types under the same circumstances.

Implementations that use the S_ISUID bit to indicate some other function (for example, mandatory record locking) on nonexecutable files need not clear this bit on writing. They should clear the bit for executable files and any other cases where the bit grants special powers to processes that change the file contents. Similar comments apply to the S_ISGID bit.

3251 B.5.6.5 Change Owner and Group of File

System III and System V allow a user to give away files; that is, the owner of a file may change its user ID to anything. This is a serious problem for implementations that are intended to meet government security regulations. Version 7 and 4.3BSD permit only the super-user to change the user ID of a file. Some government agencies (usually not ones concerned directly with security) find this limitation too confining. POSIX.1 uses "may" to permit secure implementations while not disallowing System V.

System III and System V allow the owner of a file to change the group ID to anything. Version 7 permits only the super-user to change the group ID of a file. 4.3BSD permits the owner to change the group ID of a file to its effective group ID or to any of the groups in the list of supplementary group IDs, but to no others.

Although *chown*() can be used on some systems by the file owner to change the owner and group to any desired values, the only portable use of this function is to change the group of a file to the effective GID of the calling process or to a member of its group set.

The decision to require that, for nonprivileged processes, the S_ISUID and 3267 S_ISGID bits be cleared on regular files, but only may be cleared on nonregular 3268 files, was to allow plans for using these bits in implementation-specified manners 3269 on directories. Similar cases could be made for other file types, so POSIX.1 does 3270 not require that these bits be cleared except on regular files. As these cases arise, 3271 the system implementors will have to determine whether these features enable 3272 any security loopholes and specify appropriate restrictions. If the implementation 3273 supports executing any file types other than regular files, the S_ISUID and 3274 S_ISGID bits should be cleared for those file types in the same way as they are on 3275 regular files. 3276

3277 B.5.6.6 Set File Access and Modification Times

The actime structure member must be present so that an application may set it, even though an implementation may ignore it and not change the access time on the file. If an application intends to leave one of the times of a file unchanged while changing the other, it should use stat() to retrieve the file's st_atime and st_mtime parameters, set actime and modtime in the buffer, and change one of them before making the utime() call.

3284 **B.5.7 Configurable Pathname Variables**

When the run-time facility described in B.4.8 was designed, it was realized that some variables change depending on the file system. For example, it is quite feasible for a system to have two varieties of file systems mounted: a System V file system and a BSD "Fast File System."

If limited to strictly compile-time features, no application that was widely distributed in executable binary form could rely on more than 14 bytes in a pathname component, as that is the minimum published for {NAME_MAX} in POSIX.1. The *pathconf()* function allows the application to take advantage of the most liberal file system available at run-time. In many BSD-based systems, 255 bytes are allowed for pathname components.

l

These values are potentially changeable at the directory level, not just at the file system. And, unlike the overall system variables, there is no guarantee that these might not change during program execution.

3298 **B.5.7.1 Get Configurable Pathname Variables**

- The *pathconf()* function was proposed immediately after the *sysconf()* function when it was realized that some configurable values may differ across file system, directory, or device boundaries.
- For example, {NAME_MAX} frequently changes between System V and BSD-based file systems; System V uses a maximum of 14, BSD 255. On an implementation that provided both types of file systems, an application would be forced to limit all pathname components to 14 bytes, as this would be the value specified in imits.h> on such a system.
- Therefore, various useful values can be queried on any pathname or file descriptor, assuming that the appropriate permissions are in place.
- The value returned for the variable {PATH_MAX} indicates the longest relative pathname that could be given if the specified directory is the process's current working directory. A process may not always be able to generate a name that long and use it if a subdirectory in the pathname crosses into a more restrictive file system.
- The value returned for the variable {_POSIX_CHOWN_RESTRICTED} also applies to directories that do not have file systems mounted on them. The value may change when crossing a mount point, so applications that need to know should check for each directory. [An even easier check is to try the *chown*() function and look for an error in case it happens.]
- Unlike the values returned by *sysconf()*, the pathname-oriented variables are potentially more volatile and are not guaranteed to remain constant throughout the process's lifetime. For example, in between two calls to *pathconf()*, the file system in question may have been unmounted and remounted with different characteristics.
- Also note that most of the errors are optional. If one of the variables always has the same value on an implementation, the implementation need not look at *path* or *fildes* to return that value and is, therefore, not required to detect any of the errors except the meaning of [EINVAL] that indicates that the value of *name* is not valid for that variable.
- If the value of any of the limits described in 2.8.4 or 2.8.5 are indeterminate (logically infinite), they will not be defined in <limits.h> and the pathconf() and fpathconf() functions will return -1 without changing errno. This can be distinguished from the case of giving an unrecognized name argument because errno will be set to [EINVAL] in this case.
- Since -1 is a valid return value for the *pathconf()* and *fpathconf()* functions, applications should set *errno* to zero before calling them and check *errno* only if the return value is -1.

3337 **B.6 Input and Output Primitives**

System III and System V have included a flag, O_NDELAY, to mark file descriptors so that user processes would not block when doing I/O to them. If the flag is set, a *read()* or *write()* call that would otherwise need to block for data returns a value of zero instead. But a *read()* call also returns a value of zero on end-of-file, and applications have no way to distinguish between these two conditions.

BSD systems support a similar feature through a flag with the same name, but somewhat different semantics. The flag applies to all users of a file (or socket) rather than only to those sharing a file descriptor. The BSD interface provides a solution to the problem of distinguishing between a blocking condition and an end-of-file condition by returning an error, [EWOULDBLOCK], on a blocking condition.

The 1984 /usr/group Standard (B59) includes an interface with some features from both System III/V and BSD. The overall semantics are that it applies only to a file descriptor. However, the return indication for a blocking condition is an error, [EAGAIN]. This was the starting point for POSIX.1.

The problem with the 1984 /usr/group Standard {B59} is that it does not allow compatibility with existing applications. An implementation cannot both conform to that standard and support applications written for existing System V or BSD systems. Several changes have been considered address this issue. These include:

- (1) No change (from 1984 /usr/group Standard {B59})
- 3359 (2) Changing to System III/V semantics
- 3360 (3) Changing to BSD semantics
- (4) Broadening POSIX.1 to allow conforming implementation a choice among
 these semantics
- (5) Changing the name of the flag from O_NDELAY
- (6) Changing to System III/V semantics and providing a new call to distin guish between blocking and end-of-file conditions

Alternative (5) was the consensus choice. The new name is O_NONBLOCK. This 3366 alternative allows a conforming implementation to provide backward compatibil-3367 ity at the source and/or object level with either System III/V or BSD systems (but 3368 POSIX.1 does not require or even suggest that this be done). It also allows a Con-3369 forming POSIX.1 Application Using Extensions the functionality to distinguish 3370 between blocking and end-of-file conditions, and to do so in as simple a manner as 3371 any of the alternatives. The greatest shortcoming was that it forces all existing 3372 System III/V and BSD applications that use this facility to be modified in order to 3373 strictly conform to POSIX.1. This same shortcoming applies to (1) and (4) as well, 3374 and it applies to one group of applications for (2), (3), and (6). 3375

Systems may choose to implement both O_NDELAY and O_NONBLOCK, and there
 is no conflict as long as an application does not turn both flags on at the same
 time.

3379 See also the discussion of scope in B.6.5.1.

3380 **B.6.1 Pipes**

An implementation that fails write() operations on fildes[0] or read()s on fildes[1]is not required. Historical implementations (Version 7 and System V) return the error [EBADF] in such cases. This allows implementations to set up a second pipe for full duplex operation at the same time. A conforming application that uses the pipe() function as described in POSIX.1 will succeed whether this second pipe is present or not.

- 3387 B.6.1.1 Create an Inter-Process Channel
- The wording carefully avoids using the verb "to open" in order to avoid any implication of use of *open*().
- 3390 See also B.6.4.2.

3391 **B.6.2 File Descriptor Manipulation**

3392 B.6.2.1 Duplicate an Open File Descriptor

The dup() and dup2() functions are redundant. Their services are also provided by the *fcntl()* function. They have been included in POSIX.1 primarily for historical reasons, since many existing applications use them.

While the brief code segment shown is very similar in behavior to dup2(), a conforming implementation based on other functions defined by POSIX.1 is significantly more complex. Least obvious is the possible effect of a signalcatching function that could be invoked between steps and allocate or deallocate file descriptors. This could be avoided by blocking signals.

The dup2() function is not marked obsolescent because it presents a type-safe version of functionality provided in a type-unsafe version by fcntl(). It is used in the current draft of the Ada binding to POSIX.1.

The dup2() function is not intended for use in critical regions as a synchronization mechanism.

In the description of [EBADF], the case of *fildes* being out of range is covered by the given case of *fildes* not being valid. The descriptions for *fildes* and *fildes2* are different because the only kind of invalidity that is relevant for *fildes2* is whether it is out of range; that is, it does not matter whether *fildes2* refers to an open file when the dup2() call is made.

3411 If *fildes2* is a valid file descriptor, it shall be closed, regardless of whether the 3412 function returns an indication of success or failure, unless *fildes2* is equal to 3413 *fildes*.

3414 B.6.3 File Descriptor Deassignment

3415 **B.6.3.1 Close a File**

Once a file is closed, the file descriptor no longer exists, since the integer corresponding to it no longer refers to a file.

The use of interruptible device close routines should be discouraged to avoid problems with the implicit closes of file descriptors by *exec* and *exit()*. POSIX.1 only intends to permit such behavior by specifying the [EINTR] error case.

3421 B.6.4 Input and Output

The use of I/O with large byte counts has always presented problems. Ideas such 3422 as *lread()* and *lwrite()* (using and returning *longs*) were considered at one time. 3423 The current solution is to use abstract types on the C Standard (2) interface to 3424 read() and write() (and not to discuss common usage). The abstract types can be 3425 declared so that existing interfaces work, but can also be declared so that larger 3426 types can be represented in future implementations. It is presumed that what-3427 ever constraints limit the maximum range of *size_t* also limit portable I/O requests 3428 to the same range. POSIX.1 also limits the range further by requiring that the 3429 byte count be limited so that a signed return value remains meaningful. Since 3430 the return type is also a (signed) abstract type, the byte count can be defined by 3431 the implementation to be larger than an *int* can hold. 3432

POSIX.1 requires that no action be taken when *nbyte* is zero. This is not intended to take precedence over detection of errors (such as invalid buffer pointers or file descriptors). This is consistent with the rest of POSIX.1, but the phrasing here could be misread to require detection of the zero case before any other errors. A value of zero is to be considered a correct value, for which the semantics are a no-op.

There were recommendations to add format parameters to *read()* and *write()* in order to handle networked transfers among heterogeneous file system and base hardware types. Such a facility may be required for support by the OSI presentation of layer services. However, it was determined that this should correspond with similar C Language facilities, and that is beyond the scope of POSIX.1. The concept was suggested to the developers of the C Standard {2} for their consideration as a possible area for future work.

In 4.3BSD, a read() or write() that is interrupted by a signal before transferring 3446 any data does not by default return an [EINTR] error, but is restarted. In 4.2BSD, 3447 4.3BSD, and the Eighth Edition there is an additional function, select(), whose 3448 purpose is to pause until specified activity (data to read, space to write, etc.) is 3449 detected on specified file descriptors. It is common in applications written for 3450 those systems for select() to be used before read() in situations (such as keyboard 3451 input) where interruption of I/O due to a signal is desired. But this approach does 3452 not conform, because *select()* is not in POSIX.1. 4.3BSD semantics can be provided 3453 by extensions to POSIX.1. 3454

POSIX.1 permits *read()* and *write()* to return the number of bytes successfully transferred when interrupted by an error. This is not simply required because it
was not done by Version 7, System III, or System V, and because some hardware
may not be capable of returning information about partial transfers if a device
operation is interrupted. Unfortunately, this does make writing a Conforming
POSIX.1 Application more difficult in circumstances where this could occur.

Requiring this behavior does not address the situation of pipelined buffers, such as might be found in streaming tape drives or other devices that read ahead of the actual requests. The signal interruption will often indicate an exceptional condition and flush all buffers. Thus, the amount read from the device may be different from the amount transferred to the application.

- The issue of which files or file types are interruptible is considered an implementation design issue. This is often affected primarily by hardware and reliability issues.
- There are no references to actions taken following an "unrecoverable error." It is considered beyond the scope of POSIX.1 to describe what happens in the case of hardware errors.

3472 B.6.4.1 Read from a File

POSIX.1 does not specify the value of the file offset after an error is returned; there are too many cases. For programming errors, such as [EBADF], the concept is meaningless since no file is involved. For errors that are detected immediately, such as [EAGAIN], clearly the pointer should not change. After an interrupt or hardware error, however, an updated value would be very useful and is the behavior of many implementations.

Note that a *read()* of zero bytes does not modify *st_atime*. A *read()* that requests more than zero bytes, but returns zero, does modify *st_atime*.

3481 **B.6.4.2 Write to a File**

- 3482 An attempt to write to a pipe or FIFO has several major characteristics:
- 3483 Atomic/nonatomic

A write is atomic if the whole amount written in one operation is not interleaved with data from any other process. This is useful when there are multiple writers sending data to a single reader. Applications need to know how large a write request can be expected to be performed atomically. This maximum is called {PIPE_BUF}. POSIX.1 does not say whether write requests for more than {PIPE_BUF} bytes will be atomic, but requires that writes of {PIPE_BUF} or fewer bytes shall be atomic.

- 3491 Blocking/immediate
- Blocking is only possible with O_NONBLOCK clear. If there is enough space for all the data requested to be written immediately, the implementation should do so. Otherwise, the process may block; that is, pause until enough space is available for writing. The effective size of a pipe or FIFO (the maximum amount that can be written in one operation without blocking) may vary dynamically, depending on the implementation, so it is not possible to specify a fixed value for it.

ļ

1

3501int fildes; size_t nbyte; 35033502size_t ret; char *buf;3504char *buf;3505ret = write (fildes, buf, nbyte);3506may return3507complete: ret = nbyte3508partial: ret < nbyte3509This shall never happen if nbyte \leq (PIPE_BUF)3510guarantee atomicity, even if $ret \leq$ (PIPE_J3511guarantee atomicity, even if $ret \leq$ (PIPE_J3512atomicity is guaranteed according to the amound not the amount written.3514deferred: ret = -1, errno = [EAGAIN]3515This error indicates that a later request may does not indicate that it shall succeed, even (PIPE_BUF), because if no process reads from FIFO, the write will never succeed. An app usefully count the number of times [EAGAIN]3520particular value of nbyte > (PIPE_BUF) and later writes with a smaller value, on the as the effective size of the pipe may have decrear the effective size of the pipe may have decrear.3523Partial and deferred writes are only possible with O_NONB3524The relations of these properties are shown in the following tables.	3501 3502 3503 3504 3505 3506 3507 3508 3509 3510	<pre>int fi size_t ssize_ char * ret = may return complete:</pre>	ldes; nbyte; t ret; buf; write (fildes, bu	f, nbyte);				
3502size_t nbyte;3503ssize_t ret;3504char *buf;3505ret = write (fildes, buf, nbyte);3506may return3507complete: ret = nbyte3508partial: ret < nbyte	3502 3503 3504 3505 3506 3507 3508 3509 3510	size_t ssize_ char * ret = may return complete:	nbyte; t ret; buf; write (fildes, bu	f, nbyte);				
3503 $ssize_t ret;$ char *buf;3504 $char *buf;$ 3505 $ret = write (fildes, buf, nbyte);$ 3506may return3507complete: $ret = nbyte$ 3508partial: $ret < nbyte$ 3509This shall never happen if $nbyte \leq \{PIPE_BIF\}$ 3510happen (with $nbyte > \{PIPE_BUF\}$), POSI3511guarantee atomicity, even if $ret \leq \{PIPE_J\}$ 3512atomicity is guaranteed according to the amount written.3514deferred: $ret = -1$, $errno = [EAGAIN]$ 3515This error indicates that a later request may does not indicate that it shall succeed, even3518FIFO, the write will never succeed. An app usefully count the number of times [EAGAIN]3520particular value of $nbyte > \{PIPE_BUF\}$ and later writes with a smaller value, on the as the effective size of the pipe may have decreas the effective size of the pipe may have decreas3523Partial and deferred writes are only possible with O_NONB3524The relations of these properties are shown in the following tables.	3503 3504 3505 3506 3507 3508 3509 3510	ssize_ char * ret = may return complete:	t ret; buf; write (fildes, bu	f, nbyte);				
3504char *buf;3505ret = write (fildes, buf, nbyte);3506may return3507complete: ret = nbyte3508partial: ret < nbyte	3504 3505 3506 3507 3508 3509 3510	char * ret = may return complete:	buf; write (fildes, bu	f, nbyte);				
3505ret = write (fildes, buf, nbyte);3506may return3507complete: ret = $nbyte$ 3508partial: ret < $nbyte$ 3509This shall never happen if $nbyte \leq \{PIPE_BUF\}$, POSI3510happen (with $nbyte > \{PIPE_BUF\}$), POSI3511guarantee atomicity, even if $ret \leq \{PIPE_1\}$ 3512atomicity is guaranteed according to the amo3513not the amount written.3514deferred: ret = -1 , $errno = [EAGAIN]$ 3515This error indicates that a later request may3516does not indicate that it shall succeed, even3517(PIPE_BUF), because if no process reads from3518FIFO, the write will never succeed. An app3520particular value of $nbyte > \{PIPE_BUF\}$ and3521later writes with a smaller value, on the as3522Partial and deferred writes are only possible with O_NONB3524The relations of these properties are shown in the following tables.	3505 3506 3507 3508 3509 3510	ret = may return complete:	write (fildes, bu	f, nbyte);				
3506may return3507complete: ret = $nbyte$ 3508partial: ret < $nbyte$ 3509This shall never happen if $nbyte \leq \{PIPE_BI$ 3510happen (with $nbyte > \{PIPE_BUF\}$), POSI3511guarantee atomicity, even if $ret \leq \{PIPE_J\}$ 3512atomicity is guaranteed according to the amount written.3514deferred: ret = -1 , $errno = [EAGAIN]$ 3515This error indicates that a later request may does not indicate that it shall succeed, even (PIPE_BUF), because if no process reads from FIFO, the write will never succeed. An application usefully count the number of times [EAGAIN]3520particular value of $nbyte > \{PIPE_BUF\}$ and later writes with a smaller value, on the assister writes with a smaller value, on the assister writes are only possible with O_NONB3523Partial and deferred writes are only possible with O_NONB3524The relations of these properties are shown in the following tables.	3506 3507 3508 3509 3510	may return complete:							
3507complete: $ret = nbyte$ 3508partial: $ret < nbyte$ 3509This shall never happen if $nbyte \leq (PIPE_BI<)$ 3510happen (with $nbyte > (PIPE_BUF)$), POSI3511guarantee atomicity, even if $ret \leq (PIPE_)$ 3512atomicity is guaranteed according to the amount written.3514deferred: $ret = -1$, $errno = [EAGAIN]$ 3515This error indicates that a later request matches that it shall succeed, even (PIPE_BUF), because if no process reads from FIFO, the write will never succeed. An application usefully count the number of times [EAGAIN]3520particular value of $nbyte > \{PIPE_BUF\}$ and later writes with a smaller value, on the assist the effective size of the pipe may have decreased the effective size of the pipe have	3507 3508 3509 3510	complete:							
3508partial:ret < nbyte3509This shall never happen if $nbyte \leq \{PIPE_BI$ 3510happen (with $nbyte > \{PIPE_BUF\}$), POSI3511guarantee atomicity, even if $ret \leq \{PIPE_\}$ 3512atomicity is guaranteed according to the amonot the amount written.3513deferred:ret = -1, errno = [EAGAIN]3514deferred:ret = -1, errno = [EAGAIN]3515This error indicates that a later request matches does not indicate that it shall succeed, even (PIPE_BUF), because if no process reads from FIFO, the write will never succeed. An application usefully count the number of times [EAGAIN]3520particular value of $nbyte > \{PIPE_BUF\}$ and later writes with a smaller value, on the assist the effective size of the pipe may have decreaded the effective size of the p	3508 3509 3510		$ret = nb_{c}$	yte					
3509This shall never happen if $nbyte \leq \{PIPE_BI$ 3510happen (with $nbyte > \{PIPE_BUF\}$), POSI3511guarantee atomicity, even if $ret \leq \{PIPE_1\}$ 3512atomicity is guaranteed according to the amount is guaranteed according to the amount written.3513not the amount written.3514deferred: ret = -1, errno = [EAGAIN]3515This error indicates that a later request may does not indicate that it shall succeed, even (PIPE_BUF), because if no process reads from FIFO, the write will never succeed. An app usefully count the number of times [EAGAIN]3520particular value of $nbyte > \{PIPE_BUF\}$ and later writes with a smaller value, on the as the effective size of the pipe may have decreaded the pipe may have decreaded the pipe may have decread	3509 3510	partial:	ret < nb	yte					
3510happen (with $nbyte > \{PIPE_BUF\}$), POSI3511guarantee atomicity, even if $ret \leq \{PIPE_J\}$ 3512atomicity is guaranteed according to the amount is guaranteed according to the amount written.3513deferred: $ret = -1$, $errno = [EAGAIN]$ 3514deferred: $ret = -1$, $errno = [EAGAIN]$ 3515This error indicates that a later request matches a later request matches and a later request matches a later request matches and a later matches are a later request matches and a later write later a later write later a later write later a later write later a later request matches and a later request matches and a later write later a later request matches and a later write later write later a later write later writ	3510		This sha	all never ha	ppen if nbyt	$e \leq \{\text{PIPE}_BUB\}$	F}. If it does		
3511guarantee atomicity, even if $ret \leq (PIPE_)$ 3512atomicity is guaranteed according to the amount is guaranteed according to the amount written.3513deferred: ret = -1, errno = [EAGAIN]3514deferred: ret = -1, errno = [EAGAIN]3515This error indicates that a later request matches does not indicate that it shall succeed, even (PIPE_BUF), because if no process reads from FIFO, the write will never succeed. An approximation usefully count the number of times [EAGAIN]3520particular value of nbyte > (PIPE_BUF) and later writes with a smaller value, on the assist the effective size of the pipe may have decreased the effective size of the pipe may have decreased the relations of these properties are shown in the following tables.3525Interlations of these properties are shown in the following tables.	0511		happen	(with nby	te > {PIPE_	BUF}), POSIX	.1 does not		
3512atomicity is guaranteed according to the amount3513not the amount written.3514deferred: ret = -1, errno = [EAGAIN]3515This error indicates that a later request may does not indicate that it shall succeed, even3516Goes not indicate that it shall succeed, even3517(PIPE_BUF), because if no process reads from FIFO, the write will never succeed. An app usefully count the number of times [EAGAIN]3520particular value of nbyte > {PIPE_BUF} and later writes with a smaller value, on the as the effective size of the pipe may have decreas3523Partial and deferred writes are only possible with O_NONB3524The relations of these properties are shown in the following tables.	3511		guarant	ee atomicity	y, even if r	$et \leq \{\text{PIPE}_B\}$	JF}, because		
3513not the amount written.3514deferred: ret = -1, errno = [EAGAIN]3515This error indicates that a later request may does not indicate that it shall succeed, even3516(PIPE_BUF), because if no process reads from FIFO, the write will never succeed. An app usefully count the number of times [EAGAIN] particular value of nbyte > {PIPE_BUF} and later writes with a smaller value, on the as the effective size of the pipe may have decreas3523Partial and deferred writes are only possible with O_NONB The relations of these properties are shown in the following tables.	3512		atomicit	y is guarant	teed accordin	g to the amou	nt <i>requested</i> ,		
3514deferred:ret = -1, errno = [EAGAIN]3515This error indicates that a later request may does not indicate that it shall succeed, even (PIPE_BUF), because if no process reads from FIFO, the write will never succeed. An app usefully count the number of times [EAGAIN] particular value of nbyte > {PIPE_BUF} and later writes with a smaller value, on the as the effective size of the pipe may have decreased the effective size of the pipe may h	3513		not the	amount writ	ten.		-		
3515This error indicates that a later request ma3516does not indicate that it shall succeed, ev3517{PIPE_BUF}, because if no process reads fro3518FIFO, the write will never succeed. An app3519usefully count the number of times [EAGAIN]3520particular value of nbyte > {PIPE_BUF} an3521later writes with a smaller value, on the as3523Partial and deferred writes are only possible with O_NONB3524The relations of these properties are shown in the following tables.	3514	deferred:	ret = -1	, errno = [EA	GAIN]				
3516does not indicate that it shall succeed, ev3517{PIPE_BUF}, because if no process reads fro3518FIFO, the write will never succeed. An app3519usefully count the number of times [EAGAIN]3520particular value of nbyte > {PIPE_BUF} an3521later writes with a smaller value, on the as3522the effective size of the pipe may have decreas3523Partial and deferred writes are only possible with O_NONB3524The relations of these properties are shown in the following tables.	3515		This error indicates that a later request may succeed. It						
3517{PIPE_BUF}, because if no process reads fro3518FIFO, the write will never succeed. An app3519usefully count the number of times [EAGAIN]3520particular value of nbyte > {PIPE_BUF} an3521later writes with a smaller value, on the as3522the effective size of the pipe may have decreas3523Partial and deferred writes are only possible with O_NONB3524The relations of these properties are shown in the following tables.	3516		does no	t indicate t	that it shall	succeed, eve	n if <i>nbyte</i> \leq		
3518FIFO, the write will never succeed. An application3519usefully count the number of times [EAGAIN]3520particular value of nbyte > {PIPE_BUF} and3521later writes with a smaller value, on the as3522the effective size of the pipe may have decreas3523Partial and deferred writes are only possible with O_NONB3524The relations of these properties are shown in the following tables.	3517		{PIPE_B	UF}, becaus	e if no proce	ess reads from	the pipe or		
3519usefully count the number of times [EAGAIN]3520particular value of nbyte > {PIPE_BUF} and3521later writes with a smaller value, on the as3522the effective size of the pipe may have decread3523Partial and deferred writes are only possible with O_NONB3524The relations of these properties are shown in the following tables.	3518		FIFO, th	ne write wil	l never succ	eed. An appli	ication could		
3520 particular value of nbyte > {PIPE_BUF} and later writes with a smaller value, on the as the effective size of the pipe may have decreas 3521 Partial and deferred writes are only possible with O_NONB 3524 The relations of these properties are shown in the following tables.	3519		usefully	count the n	umber of tim	es [EAGAIN] is	s caused by a		
3521 later writes with a smaller value, on the as 3522 the effective size of the pipe may have decreas 3523 Partial and deferred writes are only possible with O_NONB 3524 The relations of these properties are shown in the following tables. 3525	3520		particul	ar value of	nbyte > {PI	PE_BUF} and	perhaps do		
3522 the effective size of the pipe may have decrea 3523 Partial and deferred writes are only possible with O_NONB 3524 The relations of these properties are shown in the following tables. 3525	3521		later wr	ites with a	smaller valu	e, on the assu	umption that		
 Partial and deferred writes are only possible with O_NONB The relations of these properties are shown in the following tables. 	3522		the effec	ctive size of t	the pipe may	have decrease	ed.		
The relations of these properties are shown in the following tables.	3523	Partial and de	eferred w	vrites are on	ly possible w	ith O_NONBLO	OCK set.		
3525	3524 Th	ne relations of these	properti	es are showr	n in the follow	ving tables.			
	3525								
3526 Write to a Pipe or FIFO with O_NONBLOCK clear	3526	Write to a Pipe or FIFO with O_NONBLOCK clear							
3527 Immediately	3527	Immed	iately		~				
3528 Writable: None Some nbyte	3528	Writa	ble:	None	Some	nbyte			
3529 Atomic Atomic Atomic	0500			Atomic	Atomic	Atomic			
$nbyte \leq blocking blocking immediate$	3529	nbyt	e≤	blocking	blocking	immediate			
3531 (PIPE_BUF) nbyte nbyte nbyte	3529 3530	{PIPE_	BUF}	nbyte	nbyte	nbyte			
3532 <i>nbyte</i> > Blocking Blocking	3529 3530 3531	nhvt	e >	Blocking	Blocking	Blocking			
3533 {PIPE BUF} nbyte nbyte	3529 3530 3531 3532		BIIE)	nbvte	nbyte	nbvte			

If the O_NONBLOCK flag is clear, a write request shall block if the amount writable immediately is less than that requested. If the flag is set [by *fcntl*()], a write request shall never block.

3537	Write to a Pip	Write to a Pipe or FIFO with O_NONBLOCK set							
3538 3539	Immediately Writable:	None	Some	nbyte					
3540 3541	$nbyte \leq \\ \{PIPE_BUF\}$	–1, [EAGAIN]	–1, [EAGAIN]	Atomic nbyte					
3542 3543 3544	nbyte > {PIPE_BUF}	–1, [EAGAIN]	< nbyte or –1, [EAGAIN]	≤ nbyte or −1, [EAGAIN]					

There is no exception regarding partial writes when O_NONBLOCK is set. With the exception of writing to an empty pipe, POSIX.1 does not specify exactly when a partial write will be performed since that would require specifying internal details of the implementation. Every application should be prepared to handle partial writes when O_NONBLOCK is set and the requested amount is greater than {PIPE_BUF}, just as every application should be prepared to handle partial writes on other kinds of file descriptors.

The intent of forcing writing at least one byte if any can be written is to assure that each write will make progress if there is any room in the pipe. If the pipe is empty, {PIPE_BUF} bytes must be written; if not, at least some progress must have been made.

- Where POSIX.1 requires -1 to be returned and *errno* set to [EAGAIN], most histori-3556 cal implementations return zero (with the O_NDELAY flag set—that flag is the 3557 historical predecessor of O_NONBLOCK, but is not itself in POSIX.1). The error 3558 indications in POSIX.1 were chosen so that an application can distinguish these 3559 cases from end-of-file. While write() cannot receive an indication of end-of-file, 3560 read() can, and the two functions have similar return values. Also, some existing 3561 systems (e.g., Eighth Edition) permit a write of zero bytes to mean that the reader 3562 should get an end-of-file indication; for those systems, a return value of zero from 3563 write() indicates a successful write of an end-of-file indication. 3564
- The concept of a {PIPE_MAX} limit (indicating the maximum number of bytes that can be written to a pipe in a single operation) was considered, but rejected, because this concept would unnecessarily limit application writing.
- 3568 See also the discussion of O_NONBLOCK in B.6.

Writes can be serialized with respect to other reads and writes. If a *read()* of file data can be proven (by any means) to occur after a *write()* of the data, it must reflect that *write()*, even if the calls are made by different processes. A similar requirement applies to multiple write operations to the same file position. This is needed to guarantee the propagation of data from *write()* calls to subsequent *read()* calls. This requirement is particularly significant for networked file systems, where some caching schemes violate these semantics.

Note that this is specified in terms of *read()* and *write()*. Additional calls such as the common *readv()* and *writev()* would want to obey these semantics. A new "high-performance" write analog that did not follow these serialization requirements would also be permitted by this wording. POSIX.1 is also silent about any effects of application-level caching (such as that done by *stdio*). POSIX.1 does not specify the value of the file offset after an error is returned; there are too many cases. For programming errors, such as [EBADF], the concept is meaningless since no file is involved. For errors that are detected immediately, such as [EAGAIN], clearly the pointer should not change. After an interrupt or hardware error, however, an updated value would be very useful and is the behavior of many implementations.

POSIX.1 does not specify behavior of concurrent writes to a file from multiple processes. Applications should use some form of concurrency control.

3589 **B.6.5 Control Operations on Files**

3590 **B.6.5.1 Data Definitions for File Control Operations**

The main distinction between the file descriptor flags and the file status flags is scope. The former apply to a single file descriptor only, while the latter apply to all file descriptors that share a common open file description [by inheritance through fork() or an F_DUPFD operation with fcntl()]. For O_NONBLOCK, this scoping is like that of O_NDELAY in System V rather than in 4.3BSD, where the scoping for O_NDELAY is different from all the other flags accessed via the same commands.

3598 For example:

3599	fd1	==	open	(pathname,	oflags);
3600	fd2	===	dup	(fd1);	
3601	fd3	-	open	(pathname,	oflags);

Boes an fcntl() call on fd1 also apply to fd2 or fd3 or to both? According to POSIX.1, F_SETFD applies only to fd1, while F_SETFL applies to fd1 and fd2 but not to fd3. This is in agreement with all common historical implementations except for BSD with the F_SETFL command and the O_NDELAY flag (which would apply to fd3 as well). Note that this does not force any incompatibilities in BSD implementations, because O_NDELAY is not in POSIX.1. See also B.6.

Historically, the file descriptor flags have had only the literal values 0 and 1. POSIX.1 defines the symbolic name FD_CLOEXEC to permit a more graceful extension of this functionality. Owners of existing applications should be aware of the need to change applications using the literal values, and implementors should be aware of the existence of this practice in existing applications.

3613 B.6.5.2 File Control

The ellipsis in the Synopsis is the syntax specified by the C Standard {2} for a variable number of arguments. It is used because System V uses pointers for the implementation of file locking functions.

The *arg* values to F_GETFD, F_SETFD, F_GETFL, and F_SETFL all represent flag values to allow for future growth. Applications using these functions should do a read-modify-write operation on them, rather than assuming that only the values defined by POSIX.1 are valid. It is a common error to forget this, particularly in the case of F_SETFD, because there is only one flag in POSIX.1.

POSIX.1 permits concurrent read and write access to file data using the *fcntl()* 3622 function; this is a change from the 1984 /usr/group Standard (B59) and early 3623 POSIX.1 drafts, which included a *lockf()* function. Without concurrency controls, 3624 this feature may not be fully utilized without occasional loss of data. Since other 3625 mechanisms for creating critical regions, such as semaphores, are not included, a 3626 file record locking mechanism was thought to be appropriate. The *fcntl()* mechan-3627 ism may be used to implement semaphores, although access is not first-in-first-3628 out without extra application development effort. 3629

- Data losses occur in several ways. One is that read and write operations are not 3630 atomic, and as such a reader may get segments of new and old data if con-3631 currently written by another process. Another occurs when several processes try 3632 to update the same record, without sequencing controls; several updates may 3633 occur in parallel and the last writer will "win." Another case is a b-tree or other 3634 internal list-based database that is undergoing reorganization. Without exclusive 3635 use to the tree segment by the updating process, other reading processes chance 3636 getting lost in the database when the index blocks are split, condensed, inserted, 3637 or deleted. While *fcntl()* is useful for many applications, it is not intended to be 3638 overly general and will not handle the b-tree example well. 3639
- This facility is only required for regular files because it is not appropriate for many devices such as terminals and network connections.
- Since fcntl() works with "any file descriptor associated with that file, however it is obtained," the file descriptor may have been inherited through a fork() or *exec* operation and thus may affect a file that another process also has open.
- The use of the open file description to identify what to lock requires extra calls and presents problems if several processes are sharing an open file description, but there are too many implementations of the existing mechanism for POSIX.1 to use different specifications.
- Another consequence of this model is that closing any file descriptor for a given file (whether or not it is the same open file description that created the lock) causes the locks on that file to be relinquished for that process. Equivalently, any close for any file/process pair relinquishes the locks owned on that file for that process. But note that while an open file description may be shared through *fork()*, locks are not inherited through *fork()*. Yet locks may be inherited through one of the *exec* functions.
- The identification of a machine in a network environment is outside of the scope of POSIX.1. Thus, an *l_sysid* member, such as found in System V, is not included in the locking structure.
- Since locking is performed with fcntl(), rather than lockf(), this specification prohibits use of advisory exclusive locking on a file that is not open for writing.
- Before successful return from a F_SETLK or F_SETLKW request, the previous lock type for each byte in the specified region shall be replaced by the new lock type. This can result in a previously locked region being split into smaller regions. If this would cause the number of regions being held by all processes in the system to exceed a system-imposed limit, the *fcntl()* function returns -1 with *errno* set to [ENOLCK].

Mandatory locking was a major feature of the 1984 /usr/group Standard (B59). 3667 For advisory file record locking to be effective, all processes that have access to a 3668 file must cooperate and use the advisory mechanism before doing I/O on the file. 3669 Enforcement-mode record locking is important when it cannot be assumed that all 3670 processes are cooperating. For example, if one user uses an editor to update a file 3671 at the same time that a second user executes another process that updates the 3672 same file and if only one of the two processes is using advisory locking, the 3673 processes are not cooperating. Enforcement-mode record locking would protect 3674 against accidental collisions. 3675

Secondly, advisory record locking requires a process using locking to bracket each I/O operation with lock (or test) and unlock operations. With enforcement-mode file and record locking, a process can lock the file once and unlock when all I/O operations have been completed. Enforcement-mode record locking provides a base that can be enhanced, for example, with sharable locks. That is, the mechanism could be enhanced to allow a process to lock a file so other processes could read it, but none of them could write it.

- 3683 Mandatory locks were omitted for several reasons:
- Mandatory lock setting was done by multiplexing the set-group-ID bit in
 most implementations; this was confusing, at best.
- 3686 (2) The relationship to file truncation as supported in 4.2BSD was not well 3687 specified.
- 3688 (3) Any publicly readable file could be locked by anyone. Many historical
 3689 implementations keep the password database in a publicly readable file.
 3690 A malicious user could thus prohibit logins. Another possibility would be
 3691 to hold open a long-distance telephone line.
- 3692 (4) Some demand-paged historical implementations offer memory mapped
 3693 files, and enforcement cannot be done on that type of file.

Since sleeping on a region is interrupted with any signal, *alarm()* may be used to provide a timeout facility in applications requiring it. This is useful in deadlock detection. Because implementation of full deadlock detection is not always feasible, the [EDEADLK] error was made optional.

3698

3699 B.6.5.3 Reposition Read/Write File Offset

The C Standard {2} includes the functions *fgetpos()* and *fsetpos()*, which work on very large files by use of a special positioning type.

Although *lseek()* may position the file offset beyond the end of the file, this function does not itself extend the size of the file. While the only function in POSIX.1 that may extend the size of the file is *write()*, several C Standard {2} functions, such as *fwrite()*, *fprintf()*, etc., may do so [by causing calls on *write()*].

An invalid file offset that would cause [EINVAL] to be returned may be both implementation defined and device dependent (for example, memory may have few invalid values). A negative file offset may be valid for some devices in some implementations.

1

1

L

3710 See B.6.5.2 for a explanation of the use of signed and unsigned offsets with 3711 *lseek()*.

3712 **B.7 Device- and Class-Specific Functions**

There were several sources of difficulties involved with using historical interfaces as the basis of this section:

- (1) The basic Version 7 *ioctl()* mechanism is difficult to specify adequately,
 due to its use of a third argument that varies in both size and type
 according to the second, command, argument.
- 3718 (2) System III introduced and System V continued *ioctl()* commands that are completely different from those of Version 7.
- (3) 4.2BSD and other BSD systems added to the basic Version 7 *ioctl()* command set; some of these were for features such as job control that POSIX.1 eventually adopted.
- 3723 (4) None of the basic historical implementations are adequate in an interna 3724 tional environment. This concern is not technically within the scope of
 3725 POSIX.1, but the goal of POSIX.1 was to mandate no unnecessary impedi 3726 ments to internationalization.

The 1984 /usr/group Standard (B59) attempted to specify a portable mechanism that application writers could use to get and set the modes of an asynchronous terminal. The intention of that committee was to provide an interface that was neither implementation specific nor hardware dependent. Initial proposals dealt with high-level routines similar to the *curses* library (available on most historical implementations). In such an implementation, the user interface would consist of calls similar to:

3734 setraw(); 3735 setcooked();

It was quickly pointed out that if such routines were standardized, the definition of "raw" and "cooked" would have to be provided. If these modes were not well defined in POSIX.1, application code could not be written in a portable way. However, the definition of the terms would force low-level concepts to be included in a supposedly high-level interface definition.

Focus was given to the necessary low-level attributes that were needed to support the necessary terminal characteristics (e.g., line speeds, raw mode, cooked mode, etc.). After considerable debate, a structure similar to, but more flexible than, the System III *termio* was accepted. The format of that structure, referred to as the *termios* structure, has formed the basis for the current section.

- A method was needed to communicate with the system about the *termios* information. Proposals included:
- 3748 (1) The *ioctl()* function as in System V. This had the same problems as men 3749 tioned previously for the Version 7 *ioctl()* function and was basically
 3750 identical to it. Another problem was that the direction of the command
 3751 (whether information is written from or read into the third argument)

was not specified—in historical implementations, only the device driver knows this information. This was a problem for networked implementations. It was also a problem that there was no size parameter to specify the variable size of the third argument, and there was a similar problem with its type.

- (2) An *iocntl()* function with additional arguments specifying direction, type, 3757 and size. But these new arguments did not help application writers, who 3758 would have no control over their values, which would have to match each 3759 command exactly. The new arguments did, however, solve the problems 3760 of networked implementations. And *iocntl()* would have been implement-3761 able in terms of *ioctl()* on historical implementations (without need for 3762 modifying existing code), although it would have been easy to update 3763 existing code to use the arguments directly. 3764
- (3) A termcntl() function with the same arguments as proposed for the *iocntl*() function. The difference was that termcntl() would be limited to terminal interface functions; there would be other interface functions, such as a tapecntl() function for tape interfaces, rather than a single general device interface routine.
- (4) Unspecified functions. The issue of what the interface function(s) should be called was avoided for many of the early drafts while details of the information to be handled was of prime concern. The resulting specification resembled the information in System V, but attempted to avoid problems of case, speed, networks, and internationalization.
- Specific tc*() functions³⁾ to replace each ioctl() function were finally incorporated into POSIX.1, instead of any of the previously mentioned proposals.
- 3777 The issue of modem control was excluded from POSIX.1 on the grounds that
- 3778 It was concerned with setting and control of hardware timers.
- The appropriate timers and settings vary widely internationally.
- Feedback from European computer manufacturers indicated that this facility was not consistent with European needs and that specification of such a facility was not a requirement for portability.

3783 B.7.1 General Terminal Interface

If the implementation does not support this interface on any device types, it should behave as if it were being used on a device that is not a terminal device (in most cases *errno* will be set to [ENOTTY]) on return from functions defined by this interface. This is based on the fact that many applications are written to run both interactively and in some noninteractive mode, and they adapt themselves at run time. Requiring that they all be modified to test an environment variable to

I

^{3790 3)} The notation $tc^*()$ is reminiscent of shell pattern matching notation and is an abbreviated way of 3791 referring to all functions beginning with the letters "tc."

determine if they should try to adapt is unnecessary. On a system that provides no Section 7 interface, providing all the entry points as stubs that return [ENOTTY] (or an equivalent, as appropriate) has the same effect and requires no changes to the application.

Although the needs of both interface implementors and application developers 3796 were addressed throughout POSIX.1, this section pays more attention to the needs 3797 of the latter. This is because, while many aspects of the programming interface 3798 can be hidden from the user by the application developer, the terminal interface is 3799 usually a large part of the user interface. Although to some extent the application 3800 developer can build missing features or work around inappropriate ones, the 3801 difficulties of doing that are greater in the terminal interface than elsewhere. For 3802 example, efficiency prohibits the average program from interpreting every charac-3803 ter passing through it in order to simulate character erase, line kill, etc. These 3804 functions should usually be done by the operating system, possibly at the inter-3805 rupt level. 3806

The tc*() functions were introduced as a way of avoiding the problems inherent in the traditional ioctl() function and in variants of it that were proposed. For example, tcsetattr() is specified in place of the use of the TCSETA ioctl() command function. This allows specification of all the arguments in a manner consistent with the C Standard {2}, unlike the varying third argument of ioctl(), which is sometimes a pointer (to any of many different types) and sometimes an int.

- 3813 The advantages of this new method include:
- 3814 It allows strict type checking.
- 3815 The direction of transfer of control data is explicit.
- 3816 Portable capabilities are clearly identified.
- 3817 The need for a general interface routine is avoided.
- 3818 Size of the argument is well-defined (there is only one type).
- 3819 The disadvantages include:
- 3820 No historical implementation uses the new method.
- 3821 There are many small routines instead of one general-purpose one.
- 3822 The historical parallel with *fcntl()* is broken.
- 3823 **B.7.1.1 Interface Characteristics**
- 3824 B.7.1.1.1 Opening a Terminal Device File
- ³⁸²⁵ Further implications of the effects of CLOCAL are discussed in 7.1.2.4.
- 3826 B.7.1.1.2 Process Groups

There is a potential race when the members of the foreground process group on a terminal leave that process group, either by exit or by changing process groups. After the last process exits the process group, but before the foreground process group ID of the terminal is changed (usually by a job-control shell), it would be possible for a new process to be created with its process ID equal to the terminal's foreground process group ID. That process might then become the process group leader and accidentally be placed into the foreground on a terminal that was not necessarily its controlling terminal. As a result of this problem, the controlling terminal is defined to not have a foreground process group during this time.

The cases where a controlling terminal has no foreground process group occur 3836 when all processes in the foreground process group either terminate and are 3837 waited for or join other process groups via setpgid() or setsid(). If the process 3838 group leader terminates, this is the first case described; if it leaves the process 3839 group via setpgid(), this is the second case described [a process group leader can-3840 not successfully call *setsid()*]. When one of those cases causes a controlling termi-3841 nal to have no foreground process group, it has two visible effects on applications. 3842 The first is the value returned by tcgetpgrp(), as discussed in 7.2.3 and B.7.2.3. 3843 The second (which occurs only in the case where the process group leader ter-3844 minates) is the sending of signals in response to special input characters. The 3845 intent of POSIX.1 is that no process group be wrongly identified as the foreground 3846 process group by *tcgetpgrp()* or unintentionally receive signals because of place-3847 ment into the foreground. 3848

In 4.3BSD, the old process group ID continues to be used to identify the fore-3849 ground process group and is returned by the function equivalent to tcgetpgrp(). 3850 In that implementation it is possible for a newly created process to be assigned 3851 the same value as a process ID and then form a new process group with the same 3852 value as a process group ID. The result is that the new process group would 3853 receive signals from this terminal for no apparent reason, and POSIX.1 precludes 3854 this by forbidding a process group from entering the foreground in this way. It 3855 would be more direct to place part of the requirement made by the last sentence 3856 under 3.1.1, but there is no convenient way for that subclause to refer to the value 3857 that tcgetpgrp() returns, since in this case there is no process group and thus no 3858 process group ID. 3859

One possibility for a conforming implementation is to behave similarly to 4.3BSD, but to prevent this reuse of the ID, probably in the implementation of fork(), as long as it is in use by the terminal.

Another possibility is to recognize when the last process stops using the terminal's foreground process group ID, which is when the process group lifetime ends, and to change the terminal's foreground process group ID to a reserved value that is never used as a process ID or process group ID. (See the definition of *process group lifetime* in 2.2.2.) The process ID can then be reserved until the terminal has another foreground process group.

The 4.3BSD implementation permits the leader (and only member) of the foreground process group to leave the process group by calling the equivalent of *setpgid()* and to later return, expecting to return to the foreground. There are no known application needs for this behavior, and POSIX.1 neither requires nor forbids it (except that it is forbidden for session leaders) by leaving it unspecified.

3874 B.7.1.1.3 The Controlling Terminal

POSIX.1 does not specify a mechanism by which to allocate a controlling terminal.
This is normally done by a system utility (such as getty) and is considered an
administrative feature outside the scope of POSIX.1.

Historical implementations allocate controlling terminals on certain open() calls. Since open() is part of POSIX.1, its behavior had to be dealt with. The traditional behavior is not required because it is not very straightforward or flexible for either implementations or applications. However, because of its prevalence, it was not practical to disallow this behavior either. Thus, a mechanism was standardized to ensure portable, predictable behavior in open().

Some historical implementations deallocate a controlling terminal on its last systemwide close. This behavior in neither required nor prohibited. Even on implementations that do provide this behavior, applications generally cannot depend on it due to its systemwide nature.

3888 B.7.1.1.4 Terminal Access Control

The access controls described in this subclause apply only to a process that is accessing its controlling terminal. A process accessing a terminal that is not its controlling terminal is effectively treated the same as a member of the foreground process group. While this may seem unintuitive, note that these controls are for the purpose of job control, not security, and job control relates only to a process's controlling terminal. Normal file access permissions handle security.

If the process calling *read()* or *write()* is in a background process group that is orphaned, it is not desirable to stop the process group, as it is no longer under the control of a job-control shell that could put it into foreground again. Accordingly, calls to *read()* or *write()* functions by such processes receive an immediate error return. This is different than in 4.2BSD, which kills orphaned processes that receive terminal stop signals.

The foreground/background/orphaned process group check performed by the ter-3901 minal driver must be repeatedly performed until the calling process moves into 3902 the foreground or until the process group of the calling process becomes orphaned. 3903 That is, when the terminal driver determines that the calling process is in the 3904 background and should receive a job-control signal, it sends the appropriate sig-3905 nal (SIGTTIN or SIGTTOU) to every process in the process group of the calling pro-3906 cess and then it allows the calling process to immediately receive the signal. The 3907 latter is typically performed by blocking the process so that the signal is immedi-3908 ately noticed. Note, however, that after the process finishes receiving the signal 3909 and control is returned to the driver, the terminal driver must reexecute the fore-3910 ground/background/orphaned process group check. The process may still be in 3911 the background, either because it was continued in the background by a job-3912 control shell, or because it caught the signal and did nothing. 3913

The terminal driver repeatedly performs the foreground/background/orphaned 3914 process group checks whenever a process is about to access the terminal. In the 3915 case of write() or the control functions in 7.2, the check is performed at the entry 3916 of the function. In the case of *read()*, the check is performed not only at the entry 3917 of the function, but also after blocking the process to wait for input characters (if 3918 necessary). That is, once the driver has determined that the process calling the 3919 read() function is in the foreground, it attempts to retrieve characters from the 3920 input queue. If the queue is empty, it blocks the process waiting for characters. 3921 When characters are available and control is returned to the driver, the terminal 3922 driver must return to the repeated foreground/background/orphaned process 3923 group check again. The process may have moved from the foreground to the 3924

3925 background while it was blocked waiting for input characters.

3926 B.7.1.1.5 Input Processing and Reading Data

³⁹²⁷ There is no additional rationale provided for this subclause.

3928 B.7.1.1.6 Canonical Mode Input Processing

The term "character" is intended here. ERASE should erase the last character, not the last byte. In the case of multibyte characters, these two may be different.

4.3BSD has a WERASE character that erases the last "word" typed (but not any 3931 preceding blanks or tabs). A word is defined as a sequence of nonblank charac-3932 ters, with tabs counted as blanks. Like ERASE, WERASE does not erase beyond 3933 the beginning of the line. This WERASE feature has not been specified in POSIX.1 3934 because it is difficult to define in the international environment. It is only useful 3935 for languages where words are delimited by blanks. In some ideographic 3936 languages, such as Japanese and Chinese, words are not delimited at all. The 3937 WERASE character should presumably take one back to the beginning of a sen-3938 tence in those cases; practically, this means it would not get much use for those 3939 languages. 3940

It should be noted that there is a possible inherent deadlock if the application and implementation conflict on the value of MAX_CANON. With ICANON set (if IXOFF is enabled) and more than MAX_CANON characters transmitted without a linefeed, transmission will be stopped, the linefeed (or carriage return when ICRLF is set) will never arrive, and the *read()* will never be satisfied.

An application should not set IXOFF if it is using canonical mode unless it knows that (even in the face of a transmission error) the conditions described previously cannot be met or unless it is prepared to deal with the possible deadlock in some other way, such as timeouts.

It should also be noted that this can be made to happen in noncanonical mode if the trigger value for sending IXOFF is less than VMIN and VTIME is zero.

3952 B.7.1.1.7 Noncanonical Mode Input Processing

- 3953 Some points to note about MIN and TIME:
- 3954(1)The interactions of MIN and TIME are not symmetric. For example, when3955MIN > 0 and TIME = 0, TIME has no effect. However, in the opposite case3956where MIN = 0 and TIME > 0, both MIN and TIME play a role in that MIN3957is satisfied with the receipt of a single character.
- Also note that in case A (MIN > 0, TIME > 0), TIME represents an inter character timer while in case C (MIN = 0, TIME > 0) TIME represents a
 read timer.

These two points highlight the dual purpose of the MIN/TIME feature. Cases A and B, where MIN > 0, exist to handle burst-mode activity (e.g., file transfer programs) where a program would like to process at least MIN characters at a time. In case A, the intercharacter timer is activated by a user as a safety measure; in case B, it is turned off. Cases C and D exist to handle single-character timed transfers. These cases are readily adaptable to screen-based applications that need to know if a character is present in the input queue before refreshing the screen. In case C the read is timed; in case D, it is not.

Another important note is that MIN is always just a minimum. It does not denote a record length. That is, if a program does a read of 20 bytes, MIN is 10, and 25 characters are present, 20 characters shall be returned to the user. In the special case of MIN=0, this still applies: if more than one character is available, they all will be returned immediately.

3975 B.7.1.1.8 Writing Data and Output Processing

3976 There is no additional rationale provided for this subclause.

3977 B.7.1.1.9 Special Characters

3978 There is no additional rationale provided for this subclause.

3979 B.7.1.1.10 Modem Disconnect

3980 There is no additional rationale provided for this subclause.

3981 B.7.1.1.11 Closing a Terminal Device File

POSIX.1 is silent on whether a *close()* will block on waiting for transmission to
drain, or even if a *close()* might cause a flush of pending output. If the application
is concerned about this, it should call the appropriate function, such as *tcdrain()*,
to ensure the desired behavior.

- 3986 B.7.1.2 Parameters That Can Be Set
- 3987 B.7.1.2.1 termios Structure
- This structure is part of an interface that, in general, retains the historic grouping of flags. Although a more optimal structure for implementations may be possible, the degree of change to applications would be significantly larger.

3991 **B.7.1.2.2 Input Modes**

Some historical implementations treated a long break as multiple events, as many as one per character time. The wording in POSIX.1 explicitly prohibits this.

Although the ISTRIP flag is normally superfluous with today's terminal hardware and software, it is historically supported. Therefore, applications may be using ISTRIP, and there is no technical problem with supporting this flag. Also, applications may wish to receive only 7-bit input bytes and may not be connected directly to the hardware terminal device (for example, when a connection traverses a network).

Also, there is no requirement in general that the terminal device ensures that
high-order bits beyond the specified character size are cleared. ISTRIP provides
this function for 7-bit characters, which are common.

In dealing with multibyte characters, the consequences of a parity error in such a character, or in an escape sequence affecting the current character set, are beyond the scope of POSIX.1 and are best dealt with by the application processing the multibyte characters.

4007 **B.7.1.2.3 Output Modes**

POSIX.1 does not describe postprocessing of output to a terminal or detailed control of that from a portable application. (That is, translation of newline to carriage return followed by linefeed or tab processing.) There is nothing that a portable application should do to its output for a terminal because that would require knowledge of the operation of the terminal. It is the responsibility of the operating system to provide postprocessing appropriate to the output device, whether it is a terminal or some other type of device.

Extensions to POSIX.1 to control the type of postprocessing already exist and are expected to continue into the future. The control of these features is primarily to adjust the interface between the system and the terminal device so the output appears on the display correctly. This should be set up before use by any application.

In general, both the input and output modes should not be set absolutely, but
 rather modified from the inherited state.

4022 **B.7.1.2.4 Control Modes**

This subclause could be misread that the symbol "CSIZE" is a title in Table 7-3. Although it does serve that function, it is also a required symbol, as a literal reading of POSIX.1 (and the caveats about typography) would indicate.

4026 **B.7.1.2.5 Local Modes**

- 4027 Noncanonical mode is provided to allow fast bursts of input to be read efficiently
 4028 while still allowing single-character input.
- The ECHONL function historically has been in many implementations. Since
 there seems to be no technical problem with supporting ECHONL, it is included in
 POSIX.1 to increase consensus.
- The alternate behavior possible when ECHOK or ECHOE are specified with ICANON is permitted as a compromise depending on what the actual terminal hardware can do. Erasing characters and lines is preferred, but is not always possible.

4036 B.7.1.2.6 Special Control Characters

Permitting VMIN and VTIME to overlap with VEOF and VEOL was a compromise
for historical implementations. Only when backwards compatibility of object code
is a serious concern to an implementor should an implementation continue this
practice. Correct applications that work with the overlap (at the source level)
should also work if it is not present, but not the reverse.

ł

1

1

1

4042 B.7.1.2.7 Baud Rate Values

4043 There is no additional rationale provided for this subclause.

4044 B.7.1.3 Baud Rate Functions

The term *baud* is used historically here, but is not technically correct. This is properly "bits per second," which may not be the same as "baud." However, the term is used because of the historical usage and understanding.

4048 These functions do not take numbers as arguments, but rather symbolic names.
4049 There are two reasons for this:

- Historically, numbers were not used because of the way the rate was stored
 in the data structure. This is retained even though an interface function is
 now used.
- 4053 More importantly, only a limited set of possible rates is at all portable, and 4054 this constrains the application to that set.

There is nothing to prevent an implementation to accept, as an extension, a number (such as 126) if it wished, and because the encoding of the Bxxx symbols is not specified, this can be done so no ambiguity is introduced.

Setting the input baud rate to zero was a mechanism to allow for split baud rates. 4058 Clarifications to this version of POSIX.1 have made it possible to determine if split 4059 rates are supported and to support them without having to treat zero as a special 4060 case. Since this functionality is also confusing, it has been declared obsolescent. 4061 The 0 argument referred to is the literal constant 0, not the symbolic constant B0. 4062 POSIX.1 does not preclude B0 from being defined as the value 0; in fact, imple-4063 mentations will likely benefit from the two being equivalent. POSIX.1 does not 4064 fully specify whether the previous *cfsetispeed()* value is retained after a *tcgetattr()* 4065 as the actual value or as zero. Therefore, portable applications should always set 4066 both the input speed and output speed when setting either. 4067

In historical implementations, the baud rate information is traditionally kept in c_cflag . Applications should be written to presume that this might be the case (and thus not blindly copy c_cflag) but not to rely on it, in case it is in some other field of the structure. Setting the c_cflag field absolutely after setting a baud rate is a nonportable action because of this. In general, the unused parts of the flag fields might be used by the implementation and should not be blindly copied from the descriptions of one terminal device to another.

4075 **B.7.2 General Terminal Interface Control Functions**

The restrictions described in this subclause on access from processes in background process groups controls apply only to a process that is accessing its controlling terminal. (See B.7.1.1.4).

4079 Care must be taken when changing the terminal attributes. Applications should
4080 always do a tcgetattr(), save the termios structure values returned, and then do a
4081 tcsetattr() changing only the necessary fields. The application should use the
4082 values saved from the tcgetattr() to reset the terminal state whenever it is done

with the terminal. This is necessary because terminal attributes apply to the
underlying port and not to each individual open instance; that is, all processes
that have used the terminal see the latest attribute changes.

4086 A program that uses these functions should be written to catch all signals and 4087 take other appropriate actions to assure that when the program terminates, 4088 whether planned or not, the terminal device's state is restored to its original 4089 state. See also B.7.1.

4090 Existing practice dealing with error returns when only part of a request can be 4091 honored is based on calls to the *ioctl()* function. In historical BSD and System V 4092 implementations, the corresponding *ioctl()* returns zero if the requested actions 4093 were semantically correct, even if some of the requested changes could not be 4094 made. Many existing applications assume this behavior and would no longer 4095 work correctly if the return value were changed from zero to -1 in this case.

4096 Note that either specification has a problem. When zero is returned, it implies
4097 everything succeeded even if some of the changes were not made. When -1 is
4098 returned, it implies everything failed even though some of the changes were
4099 made.

4100 Applications that need all of the requested changes made to work properly should 4101 follow *tcsetattr()* with a call to *tcgetattr()* and compare the appropriate field 4102 values.

- 4103 B.7.2.1 Get and Set State
- 4104 The *tcsetattr()* function can be interrupted in the following situations:
- 4105 It is interrupted while waiting for output to drain.
- 4106 It is called from a process in a background process group and SIGTTOU is 4107 caught.
- 4108 B.7.2.2 Line Control Functions
- 4109 There is no additional rationale provided for this subclause.

4110 B.7.2.3 Get Foreground Process Group ID

The tcgetpgrp() function has identical functionality to the 4.2BSD ioctl() function
TIOCGPGRP except for the additional security restriction that the referenced terminal must be the controlling terminal for the calling process.

In the case where there is no foreground process group, returning an error rather 4114 than a positive value was considered. This was rejected because existing applica-4115 tions based on either IEEE Std 1003.1-1988 or 4.3BSD are likely to consider errors 4116 from this call or the BSD equivalent to be catastrophic and respond inappropri-4117 ately. Such applications implicitly assume that this case does not exist, and the 4118 positive return value is the only solution that permits them to behave properly 4119 even when they do encounter it. No application has been identified that can 4120benefit from distinguishing between this case and the case of a valid foreground 4121process group other than its own. Therefore, requiring or permitting any other 4122

solution would cause more application portability problems with no corresponding
benefit to applications. The value must be positive, not zero, because applications
may use the negation as the *pid* argument to *kill()*. In addition, the value 1 must
not be used so that an attempt to send a signal to this (nonexistent) process group
does not result in broadcasting a signal unintentionally. See also B.7.1.1.2.

4128 B.7.2.4 Set Foreground Process Group ID

The tcsetpgrp() function has identical functionality to the 4.2BSD *ioctl*() function TIOCSPGRP except for the additional security restrictions that the referenced terminal must be the controlling terminal for the calling process and the specified new process group must be currently in use in the caller's session.

4133 **B.8 Language-Specific Services for the C Programming Language**

4134 See the discussion of C functions in B.1.1.1.

Common usage may be defined by historical publications such as *The C Programming Language* {B46}.

4137 The null set of supported languages is allowed.

The list of functions comprises the list of "common-usage" functions and also 4138 includes those that are not in common usage that are addressed by POSIX.1. The 4139 rules for common-usage conformance to POSIX.1 address whether the functions 4140 that are not generally considered in common usage are implemented. There are a 4141 large number of functions found in various systems that, although frequently 4142 found, are not broadly enough available to be considered in common usage. The 4143 signal() function (although in common usage) is omitted because applications con-4144 forming to POSIX.1 should use the more reliable *sigaction()* interface instead. 4145

4146 **B.8.1 Referenced C Language Routines**

4147 **B.8.1.1 Extensions to Time Functions**

System V uses the TZ environment variable to set some information about time.
It has the form (spaces inserted for clarity):

4150 *std offset dst*

where the first three characters (std) are the name of the standard time zone, the 4151 digits that follow (offset) represent the time added to the local time zone to arrive 4152 at Coordinated Universal Time, and the next three characters (dst) are the name 4153 of the summer time zone. The meaning of offset implies that most sites west of 4154 the Prime Meridian will have a positive offset (preceded by an optional plus sign, 4155 "+"), while most sites east of the Prime Meridian will have a negative offset (pre-4156 ceded by a minus sign, "-"). Both std and offset are required; if dst is missing, 4157 summer time does not apply. 4158

Currently, the UNIX system *localtime()* function translates a number of seconds
 since the Epoch into a detailed breakdown of that time. This breakdown includes

- 4161 (1) Time of day: Hours, minutes, and seconds.
- (2) Day of the month, month of the year, and the year.
- (3) Day of the week and day of the year (Julian day).
- (4) Whether or not summer (daylight saving) time is in effect.

It is the first and last items that present a problem: the time of the day depends
on whether or not summer time is in effect. Whether or not summer time is in
effect depends on the locale and date.

4168 Most historical systems had time-zone rules compiled into the C library. These 4169 rules usually represented United States rules for 1970 to 1986. This did not 4170 accommodate the changes of 1987, nor other world variations (½-hour time, dou-4171 ble daylight time, and solar time being common, but not complete, examples). 4172 Some recent systems addressed these problems in various ways.

Having the rules compiled into the program made binary distributions that
accommodated all the variations (including sudden changes to the law), and perprocess rule changes, difficult at best.

POSIX.1 includes a way of specifying the time zone in the TZ string, but only permits one time-zone pattern at a time, thus not dealing with different patterns in previous years and with such issues as solar time. Methods exist to deal with all the problems above. The method in POSIX.1 appears to be simpler to implement and may be faster in execution when it is adequate. POSIX.1 also permits an implementation-defined rule set that begins with a colon. (The previous format cannot begin with a colon.)

Rules of the form AAAn or AAAnBBB (the style used in many historical implemen-4183 tations) do not carry with them any statement about the start and end of daylight 4184 time (neither the date nor the time of day; the default to 02:00 not applying if no 4185 rule is present at all), thus implying that the implementation must provide the 4186 appropriate *rules*. An implementation may provide those rules in any way it sees 4187 fit, as long as the constraints implied by the TZ string as provided by the user are 4188 met. Specifically, the implementation may use the string as an index into a table, 4189 which may reside either on disk or in memory. Such tables could contain rules 4190that are sensitive to the year to which they are applied, again since the user did 4191 not specify the exact rule. (Although impractical, every possible TZ string could 4192 be represented in a table, as a detail of implementation; the less specific the user 4193 is about the TZ string, the more freedom the implementation has to interpret it.) 4194

There is at least one public-domain time-zone implementation (the Olson/Harris method) that uses nonspecific TZ strings and a table, as described previously, and handles all the general time-zone problems mentioned above. This implementation also appears in a late release of 4.3BSD. If this implementation honors all the specifications provided in the TZ string, it would conform to POSIX.1. Nothing precludes the implementation from adding information beyond that given by the user in the TZ string.

The fully specified TZ environment variable extends the historical meaning to also include a rule for when to use standard time and when to use summer time. 4204 Southern hemisphere time zones are supported by allowing the first *rule date* 4205 (change to summer time) to be later in the year than the second *rule date* (change 4206 to standard time).

This mechanism accommodates the "floating day" rules (for example "last Sunday
in October") used in the United States and Canada (and the European Economic
Community for the last several years). In theory, TZ only has to be set once and
then never touched again unless the law is changed.

Julian dates are proposed with two syntaxes, one zero-based, the other one-based. They are here for historical reasons. The one-based counting (J) is used more commonly in Europe (and on calendars people may use for reference). The zerobased counting (n) is used currently in some implementations and should be kept for historical reasons as well as being the only way to specify Leap Day.

- It is expected that the leading colon format will allow systems to implement an even broader range of specifications for the time zone without having to resort to a file or permit naming an explicit file containing the appropriate rules.
- The specification in POSIX.1 for TZ assumes that very few programs need to be historically accurate as long as the relative timing of two events is preserved.
- Summer time is governed by both locale and date. This proposal only handles the
 locale dependency. Using an implementation-defined file format for either the
 entire TZ variable or to specify the *rules* for a particular time zone is allowed as a
 means by which both the locale and date dependency can be handled.
- Since historical implementations do not examine TZ beyond the assumed end of dst, it is possible literally to extend TZ and break very little existing software. Since much historical software does not function outside the US time zones, minor changes to TZ (such as extending offset to be hh:mm—as long as the colon and minutes, :mm, are optional) should have little effect.
- 4230 POSIX.1 is intentionally silent about values of TZ that do not fit either of the 4231 specified forms. It simply requires that TZ values that follow those forms be 4232 interpreted as specified.
- 4233 B.8.1.2 Extensions to setlocale() Function
- 4234 The C Standard (2) defines a collection of interfaces to support internationaliza-4235 tion. One of the most significant aspects of these interfaces is a facility to set and 4236 query the *international environment*. The international environment is a reposi-4237 tory of information that affects the behavior of certain functionality, namely
- 4238 (1) Character Handling
- 4239 (2) String Handling (i.e., collating)
- 4240 (3) Date/Time Formatting
- 4241 (4) Numeric Editing

The *setlocale()* function provides the application developer with the ability to set all or portions, called *categories*, of the international environment. These categories correspond to the areas of functionality, mentioned above. The syntax for *setlocale()* is the following:

1

4246 char *setlocale(int category, const char *locale);

4247 where *category* is the name of one of five categories, namely

- 4248LC_CTYPE4249LC_COLLATE4250LC_TIME4251LC_MONETARY4252LC_NUMERIC
- ⁴²⁵³ In addition, a special value, called LC_ALL, directs *setlocale()* to set all categories.
- The *locale* argument is a character string that points to a specific setting for the international environment, or locale. There are three preset values for the locale argument, namely
- 4257"C"Specifies the minimal environment for C translation. If setlocale()4258is not invoked, the "C" locale is the default.
- POSIX" Specifies a locale that is the same as "C" for the attributes defined by the C Standard {2} and POSIX.1, but may contain extensions.
 The wording permits extensions by standards, specifically that of ISO/IEC 9945-2 (B36), which is expected to use the same symbol, and by future versions of POSIX.1.
- 4264 "" Specifies an implementation-defined native environment.
- 4265 **NULL** Used to direct *setlocale()* to query the current international 4266 environment and return the name of the locale.
- This subclause describes the behavior of an implementation of *setlocale()* and its use of environment variables in controlling this behavior on POSIX.1-based systems. There are two primary uses of *setlocale()*:
- (1) Querying the international environment to find out what it is set to;
- 4271 (2) Setting the international environment, or *locale*, to a specific value.
- The following subclauses describe the behavior of *setlocale()* in these two areas. Since it is difficult to describe the behavior in words, examples will be used to illustrate the behavior of specific uses.
- To query the international environment, *setlocale()* is invoked with a specific category and the NULL pointer as the locale. The NULL pointer is a special directive to *setlocale()* that tells it to query rather than set the international environment. The following syntax is used to query the name of the international environment:

The *setlocale()* function returns the string corresponding to the current international environment. This value may be used by a subsequent call to *setlocale()* to

Part 1: SYSTEM API [C LANGUAGE]

reset the international environment to this value. However, it should be noted that the return value from *setlocale()* is a pointer to a static area within the function and is not guaranteed to remain unchanged [i.e., it may be modified by a subsequent call to *setlocale()*]. Therefore, if the purpose of calling *setlocale()* is to save the value of the current international environment so it can be changed and reset later, the return value should be copied to an array of *char* in the calling program.

4289 There are three ways to set the international environment with *setlocale()*:

4290 setlocale (*category*, *string*)

This usage will set a specific *category* in the international environment to a specific value corresponding to the value of the *string*. A specific example is provided below:

4294 setlocale(LC_ALL, "Fr_FR.8859");

- 4295In this example, all categories of the international environment4296will be set to the locale corresponding to the string4297"Fr_FR.8859", or to the French language as spoken in France4298using the ISO 8859-1 code set.
- 4299If the string does not correspond to a valid locale, setlocale()4300will return a NULL pointer and the international environment4301is not changed. Otherwise, setlocale() will return the name of4302the locale just set.

4303 setlocale(*category*, "C")

4304The C Standard {2} states that one locale must exist on all con-4305forming implementations. The name of the locale is "C" and4306corresponds to a minimal international environment needed to4307support the C programming language.

4308 setlocale (category, "")

- 4309This will set a specific category to an implementation-defined4310default. For POSIX.1-based systems, this corresponds to the4311value of the environment variables.
- 4312 B.8.2 C Language Input/Output Functions

4313 **B.8.2.1 Map a Stream Pointer to a File Descriptor**

4314 Without some specification of which file descriptors are associated with these 4315 streams, it is impossible for an application to set up the streams for another appli-4316 cation it starts with fork() and *exec*. In particular, it would not be possible to 4317 write a portable version of the sh command interpreter (although there may be 4318 other constraints that would prevent that portability).

4319 **B.8.2.2 Open a Stream on a File Descriptor**

4320 The file descriptor may have been obtained from open(), creat(), pipe(), dup(), or 4321 fcntl(); inherited through fork() or *exec*; or perhaps obtained by implementation-4322 dependent means, such as the 4.3BSD socket() call.

I

I

The meanings of the *type* arguments of *fdopen()* and *fopen()* differ. With *fdopen()*, open for write ("w" or "w+") does not truncate, and append ("a" or "a+") cannot create for writing. There is no need for "b" in the format due to the equivalence of binary and text files in POSIX.1. See B.1.1.1. Although not explicitly required by POSIX.1, a good implementation of append ("a") mode would cause the O_APPEND flag to be set.

4329 **B.8.2.3 Interactions of Other FILE-Type C Functions**

Note that the existence of open streams on a file implies open file descriptors and thus affects the timestamps of the file. The intent is that using *stdio* routines to read a file must eventually update the access time, and using them to write a file must eventually update the modify and change times. However, the exact timing of marking the *st_atime*, *st_ctime*, and *st_mtime* fields cannot be specified, as that would imply a particular buffering strategy.

The purpose of the rules about handles is to allow the writing of a program that uses *stdio* and does some shell-like things; in particular, creating an open file for a child process to use, where both the parent and child wish to use *stdio*, with the consequences of buffering. In most cases, this cannot happen in the C Standard {2} (because there is no way to create a second handle), but the *system()* function can cause this to occur, at least in most historical implementations.

Presently, POSIX.1 deals mostly with output streams; input is implementation defined. It should be possible to make input on seekable devices work for seekable files without affecting buffering strategies significantly. However, the details have not been worked out fully and will be addressed in a future revision of POSIX.1. The requirements on applications are unlikely to change [basically, serving notice to the implementation that the use of a particular handle is (temporarily) completed] and are symmetric to those for output.

There are some implied rules about interprocess synchronization, but no mechan-4349 ism is given, intentionally. In the simplest case, if the parent meets the require-4350 ments on all its files and then performs a fork() and a wait() before further 4351 activity on them [and a *fflush()* on input files after that], the desired synchroniza-4352 tion will be achieved. Synchronization could in theory be done with signals, but 4353 the only likely case is the one just described. The terms handle and active handle 4354 were required to make the text readable and are not intended for use outside this 4355 discussion. 4356

4357 Note that since *exit()* implies *_exit()*, a file descriptor is also closed by *exit()*.

Because a handle is either freshly opened, or if not must have handed off control of the open file description as specified, the new handle is always ready to be used (except for seeks) with no initialization. [A freshly opened stream has not yet done any reads, as required by the C Standard (2), at least implicitly by the rules associated with *setvbuf*().]

In requiring the seek to an appropriate location for the new handle, the application is required to know what it is doing if it is passing streams with seeks involved. If the required seek is not done, the results are undefined (and in fact the program probably will not work on many common implementations). A naive program used as a utility can be reasonably expected to work properly
when the constraints are met by the calling program because it will not hand off
file descriptors except with closes.

The exec functions are treated specially because the application should always fflush() everything before performing one of the exec functions. If stdout is available on the same open file description after the exec, it is a different stream, at least because any unflushed data will be discarded during the exec (similarly for stdin). Process termination is also special because a process terminating due to a signal or _exit() will not have the buffers flushed.

The fork() function also must be specially treated because it clones a number of 4376 file descriptors simultaneously. Thus, all of them should be prepared for handoff 4377 before the fork(). In effect, fork() creates a pair of handles that are improperly 4378 dealt with unless, before the fork(), the first part of a handoff occurred. Note that 4379 fflush (NULL) in the C Standard {2} is an appropriate way to do this for output. A 4380 subsequent exec call [that does not succeed in calling exit() in some way] will 4381 reduce the number of handles back to the original value (allowing for files that 4382 are not close-on-exec), and, thus, preparations for exec need not necessarily do the 4383 flush. However, because exit() closes all streams, if the exec fails, the application 4384 must be careful to terminate with exit(). 4385

POSIX.1 does not specify asynchronous I/O, and when dealing with asynchronous
I/O the problem of coordinating access to streams will be more difficult. If asynchronous I/O is provided as an extension, the problems it introduces in this area
should be addressed as part of that extension.

It may be that functions such as system() and popen(), currently being considered
for ISO/IEC 9945-2 (B36), will have to perform some of these operations.

The introduction of underlying functions allows generic reference to *errno* values returned by those functions and also to other side effects (as required in the *handles* discussion above). It is not intended to specify implementation, although many implementations may in fact use those functions. The C Standard {2} says very little about *errno* in the context of *stdio*. In the more restricted POSIX.1 environment, providing a reasonable set of *errno* values become possible.

- 4398
- 4399 **B.8.2.3.1** fopen()

4400 There is no additional rationale provided for this subclause.

4401 **B.8.2.3.2** *fclose()*

The *fclose()* function is required to synchronize the buffer pointer with the file pointer (unless it already is, which would be the case at EOF). Functionality equivalent to

4405 fseek(stream, ftell(stream), SEEK_SET)

does this nicely. The exception for devices incapable of seeking is an obvious requirement, but the implication is that there is no way to reliably read a buffered pipe and hand off handles. This is the situation in historical implementations and is inherent in any "read-ahead" buffering scheme. This limitation is also reflected in the handle hand-off rules.

B.8 Language-Specific Services for the C Programming Language

I

- 4411 Note that the last byte read from a stream and the last byte read from an open 4412 file description are not necessarily the same; in most cases the open file 4413 description's pointer will be past that of the stream because of the stream's read-4414 ahead.
- 4415 **B.8.2.3.3** freopen()
- 4416 There is no additional rationale provided for this subclause.
- 4417 **B.8.2.3.4** *fflush()*
- 4418 There is no additional rationale provided for this subclause.
- 4419 B.8.2.3.5 fgetc(), fgets(), fread(), getc(), getchar(), gets(), scanf(), fscanf()
- 4420 There is no additional rationale provided for this subclause.
- 4421 B.8.2.3.6 fputc(), fputs(), fwrite(), putc(), putchar(), puts(), printf(), 4422 vprintf(), vfprintf()
- 4423 There is no additional rationale provided for this subclause.
- 4424 B.8.2.3.7 fseek(), rewind()
- The fseek() function must operate as specified to make the case where seeking is being done work. The key requirement is to avoid an optimization such that an fseek() would not result in an lseek() if the fseek() pointed within the current buffer. This optimization is valuable in general, so it is only required after an fflush().
- 4430 **B.8.2.3.8** *perror*()
- 4431 There is no additional rationale provided for this subclause.
- 4432 **B.8.2.3.9** *tmpfile()*
- 4433 There is no additional rationale provided for this subclause.
- 4434 **B.8.2.3.10** *ftell()*

In append mode, a *fflush()* will change the seek pointer because of possible writes by other processes on the same file. An *fseek()* reflects the underlying file's file offset, which is not necessarily the end of the file. Implementors should be aware that the operating system itself (not some in-memory approximation) of the file offset should be queried when in append mode.

4440 B.8.2.3.11 Error Reporting

POSIX.1 intentionally does not require that all errors detected by the underlying functions be detected by the functions listed here. There are many reasonable cases where this might not occur; for example, many of the functions with *write()* as an underlying function might not detect a number of error conditions in cases where they simply buffer output for a subsequent flush.

1

1

[ENOMEM] was considered for addition as an explicit possible error because most 4446 implementations use *malloc()*. This was not done because the scope does not 4447 include "out of resource" errors. Nevertheless this is the most likely error to be 4448 added to the possible error conditions. Other implementation-defined errors, par-4449 ticularly in the f*open() family, are to be expected, and the generic rules about 4450 adding (or deleting) possible errors apply, except that it is expected that 4451 implementation-defined changes in the error set returned by open() would also 4452 apply to fopen() [unless the condition cannot possibly happen in fopen(), which 4453 may be possible, but appears unlikely]. 4454

4455 **B.8.2.3.12** exit(), abort()

POSIX.1 intends that processing related to the *abort()* function will occur unless
"the signal SIGABRT is being caught, and the signal handler does not return," as
defined by the C Standard {2}. This processing includes at least the effect of *fclose()* on all open streams, and the default actions defined for SIGABRT.

The *abort()* function will override blocking or ignoring the SIGABRT signal.
Catching the signal is intended to provide the application writer with a portable
means to abort processing, free from possible interference from any
implementation-provided library functions.

- 4464 Note that the term "program termination" in the C Standard {2} is equivalent to 4465 "process termination" in POSIX.1.
- 4466 **B.8.2.4** Operations on Files the *remove*() Function
- 4467 There is no additional rationale provided for this subclause.
- 4468 **B.8.3 Other C Language Functions**

4469 **B.8.3.1 Nonlocal Jumps**

- 4470 The C Standard $\{2\}$ specifies various restrictions on the usage of the setjmp()4471 macro in order to permit implementors to recognize the name in the compiler and 4472 not implement an actual function. These same restrictions apply to the sig-4473 setjmp() macro.
- There are processors that cannot easily support these calls, but this was not considered a sufficient reason to exclude them.
- The distinction between setjmp()/longjmp() and sigsetjmp()/siglongjmp() is only
 significant for programs that use the sigaction(), sigprocmask(), or sigsuspend()
 functions. Since earlier implementations did not have signal masks, only a single
 pair was provided.

4480
4.2BSD and 4.3BSD systems provide functions named _setjmp() and _longjmp()
4481
4481
4482
4482
4482
4483
4483
4484
4484
4484
4484
4484
4485
4485
4485
4486
4486
4486
4486
4487
4486
4486
4488
4486
4486
4489
4486
4486
4480
4486
4486
4486
4487
4486
4486
4486
4486
4486
4487
4486
4486
4486
4486
4487
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486
4486</l

I

I

existing practice in both cases, the relation of setjmp() and longjmp() to signal masks is not specified, and a new set of functions is defined instead.

4488 **B.8.3.2 Set Time Zone**

4489 There is no additional rationale provided for this subclause.

4490 B.8.3.3 Omitted Memory Management

The brk() and sbrk() functions frequently were proposed for inclusion in POSIX.1, 4491 but they were excluded deliberately. See also B.1.1. The rationale for including 4492 them is usually addressed to the argument that it is the sbrk() primitive that 4493 makes it possible to implement some more general heap management system, 4494 such as that provided for C by malloc(). The need for such functionality is fully 4495 understood, but specifying it as a part of a standard would have the effect of limit-4496 ing the number of architectures that could support POSIX.1. It might also con-4497 strain languages whose memory-management model was not served by the *sbrk()* 4498 model. 4499

Memory management is not excluded from POSIX.1: POSIX.1 relies on the 4500 language to provide it, and in the C binding (as reflected in Section 8) it is pro-4501 vided by malloc(). It would be provided by new() in Pascal. In a language like 4502 FORTRAN, which does not supply memory management to the user, it would be 4503 undesirable to force the language binding to attempt to include such a function. 4504 It is reasonable to imagine a language that required a more powerful primitive 4505 than sbrk() to be implemented, and standardizing sbrk() would only constrain 4506 such future languages. 4507

POSIX.1 is silent about mixed languages. Mixing languages that provide incompatible memory-management mechanisms can yield unpredictable results. Future
standards that address mixing of languages should consider this issue.

4511 Architectures that could not support sbrk() are also a limiting factor. In particu-4512 lar, architectures that do not present a model of a single linear address space 4513 would be severely constrained by sbrk(), but are not so constrained by malloc() or 4514 new().

Each language should specify the memory-management primitives best suited to 4515 that language. Whether the implementor chooses to use a more primitive 4516 mechanism to implement that, or the implementor chooses to directly implement 4517 the language function in the kernel, is not a proper concern of the developers of 4518 POSIX.1, nor should it be for any portable application. An application that 4519 presumes the sbrk() model of memory management will not port to all architec-4520 tures in any case, for the same reasons that sbrk() itself does not work on those 4521 architectures. No true gain in application portability would be achieved by man-4522 dating such an interface. This implies that an implementor of software that 4523 wishes to port to multiple platforms and that attempts to implement its own 4524 memory management rather than relying on language-supplied functions must be 4525 prepared to deal with multiple platform-supplied primitives and, because it is 4526 4527 doing its own memory management inherently, cannot be considered, or be made to be, portable in that regard. 4528

1

4529 **B.9 System Databases**

At one time, this section was entitled Passwords, but this title was changed as all references to a "password file" were changed to refer to a "user database."

4532 **B.9.1 System Databases**

There are no references in POSIX.1 to a *passwd file* or a *group file*, and there is no requirement that the *group* or *passwd* databases be kept in files containing editable text. Many large timesharing systems use *passwd* databases that are hashed for speed. Certain security classifications prohibit certain information in the *passwd* database from being publicly readable.

- The term "encoded" is used instead of "encrypted" in order to avoid the implementation connotations (such as reversibility or use of a particular algorithm) of the latter term.
- The getgrent(), setgrent(), endgrent(), getpwent(), setpwent(), and endpwent() functions are not included in POSIX.1 because they provide a linear database search capability that is not generally useful [the getpwuid(), getpwnam(), getgrgid(), and getgrnam() functions are provided for keyed lookup] and because in certain distributed systems, especially those with different authentication domains, it may not be possible or desirable to provide an application with the ability to browse the system databases indiscriminately.
- A change from historical implementations is that the structures used by these functions have fields of the types gid_t and uid_t , which are required to be defined in the header <sys/types.h>. POSIX.1 has not changed the synopses of these functions to require the inclusion of this header, since that would invalidate a large number of existing applications. Implementations must ensure that these types are defined by the inclusion of <grp.h> and <pwd.h>, respectively, without imposing any namespace pollution or errors from redefinition of types.

POSIX.1 is silent about the content of the strings containing user or group names.
These could be digit strings. POSIX.1 is also silent as to whether such digit strings bear any relationship to the corresponding (numeric) user or group ID.

4558 **B.9.2 Database Access**

4559 **B.9.2.1 Group Database Access**

- 4560 There is no additional rationale provided for this subclause.
- 4561 **B.9.2.2 User Database Access**
- 4562 There is no additional rationale provided for this subclause.

4563 **B.10 Data Interchange Format**

4564 **B.10.1 Archive/Interchange File Format**

- 4565 There are three areas of interest associated with file interchange:
- 4566 (1) Media. There are other existing standards that define the media used for
 4567 data interchange.
- 4568 (2) User Interface. This rightfully should be in the shell and utilities stan-4569 dard, under development as ISO/IEC 9945-2 {B36}.
- (3) Format of the Data. None of the groups currently developing POSIX standards address topics that match this area. The groups felt that this area is closest to the types of things that should be in the POSIX.1 document, as the level of that document most closely matches the level of data required.
- There are two programs in wide use today: tar and cpio. There are many supporters for each program. Four options were considered for POSIX.1:
- 4577 (1) Make both formats optional. This was considered unacceptable because
 4578 it does not allow any portable method for data interchange.
- 4579 (2) Require one format.
- 4580 (3) Require one format with the other optional.
- 4581 (4) Require both formats.
- 4582 Both the Extended cpio and the Extended tar Formats are required by POSIX.1.

There are a number of concerns about defining extensions that are known to be required by historical implementations. Failure to specify a consistent method to implement these extensions will limit portability of the data and, more importantly, will create confusion if these extensions are later standardized.

- Two of these extensions that should be documented are symbolic links, which were defined by 4.2BSD and 4.3BSD systems, and high-performance (or contiguous) files, which exist in a number of implementations and are now being considered for future amendments to POSIX.1.
- By defining these extensions, implementors are able to recognize these features and take appropriate implementation-defined actions for these files. For example, a high-performance file could be converted to a regular file if the system did not support high-performance files; symbolic links might be replaced by normal hard links.
- The policy of not defining user interfaces to utilities preempted any description of a tar or cpio command. The behavior of the former command was described in some detail in previous drafts.
- 4599 The possibilities for transportable media include, but are not limited to
- 4600 (1) 12,7 mm (0,5 in) magnetic tape, 9 track, 63 bpmm (1600 bpi)
- 4601 (2) 12,7 mm (0,5 in) magnetic tape, 9 track, 246 cpmm (6250 cpi)

1

I

- 4602 (3) QIC-11, 6,30 mm (0,25 in) streamer tape
- 4603 (4) QIC-24, 6,30 mm (0,25 in) streamer tape
- 4604 (5) 130 mm (5,25 in) diskettes, 9 512-byte sectors/track, 3,8 tpmm (96 tpi)
- (6) 130 mm (5,25 in) diskettes, 9 512-byte sectors/track, 1,9 tpmm (48 tpi)

4606 When selecting media, issues such as character frame size also need to be 4607 addressed. The easiest environment for interchange occurs when 8-bit frames are 4608 used.

The utilities are not restricted to work only with *transportable* media: existing related utilities are often used to transport data from one place to another in the file hierarchy.

The formats are included to provide implementation-independent ways to move 4612 files from one system to another and also to provide ways for a user to save data 4613 on a transportable medium to be restored at a later date. Unfortunately, these 4614 two goals can contradict each other, as system security problems are easy to find 4615 in tape systems if they are not protected. Thus, there are strict requirements 4616 about how the mechanism to copy files shall react when operated by both 4617 privileged and nonprivileged users. The general concept is that a privileged (his-4618 torically using the ISUID bit in the file's mode with the owner UID of the file set to 4619 super-user) version of the utility (or one operated by a privileged user) can be 4620 used as a save/restore scheme, but a nonprivileged version is used to interpret 4621 media from a different system without compromising system security. 4622

Regardless of the archive format used, guidelines should be observed when writing tapes to be read on other systems. Assuming the target system conforms to POSIX.1, archives created should only use definitions found in POSIX.1 (e.g., file types, minimum values as found in Section 2) and should only use relative pathnames (i.e., no leading slash).

Both tar and cpio formats have traditionally been used for both exchange of 4628 information and archiving. These formats have a number of features that facili-4629 tate archiving, for example, the ability to store information about a file that is a 4630 device. POSIX.1 does not assume this kind of data is portable. It is intended that 4631 these formats provide for the portable exchange of source information between 4632 dissimilar systems. This requires specification of the character set to be used 4633 (ISO/IEC 646 {1}) when these formats are used to write source information. The 4634 1990 version of ISO/IEC 646 {1} IRV was selected as the international character set 4635 that corresponds most directly to the ASCII set used in many historical implemen-4636 tations. The 1990 version was chosen over the 1983 version because it defines 4637 '\$' as the currency symbol in the IRV, as opposed to the starburst-like generic 4638 currency symbol. Note that ISO/IEC 646 {1} is a safe lowest-common-denominator 4639 character set and that interchange of larger character sets is permitted by mutual 4640 agreement. Using any other character set (such as ISO 8859-1 {B34} or some mul-4641 tibyte character set) reduces the number of machines to which interchange is 4642 guaranteed. 4643

All data written by format-creating utilities and read by format-reading utilities is an ordered stream of bytes. The first byte of the stream should be first on the medium, the second byte second, etc. On systems where the hardware swaps bytes or otherwise rearranges the byte stream on output or input, the

I

implementor of these utilities must compensate for this so that the data on thestorage device retains its ordered nature.

POSIX.1 describes two different formats for data archiving and interchange. 4650 Strong support for both formats was evident through the balloting process. This 4651 is a clear indication of the need for both formats due to existing practice. The bal-4652 loting process also defined a number of deficiencies of each format. The strong 4653 support indicates that these deficiencies are not sufficient to remove either format 4654 from POSIX.1, but will need to be addressed in future amendments to POSIX.1. It 4655 was not practical to remedy these deficiencies during the balloting process. Con-4656 siderable thought and review must occur before making any changes to these for-4657 mats. It was felt that the best solution is to advise implementors and application 4658 writers of these deficiencies by documenting them in the rationale and to include 4659 both formats in POSIX.1. 4660

The developers of POSIX.1 recognize the desirability for migration toward one common format and have been made aware of some strong inputs to consider both formats in light of existing practice, current technology trends, and the POSIX standards activities such as security and high-performance systems, and to develop one format that is technically superior. This format will be incorporated into a future amendment to POSIX.1 when it is developed.

- The deficiencies that have been identified in the existing formats are as follows. 4667 The size of a file link is limited to 100 characters in tar. A number of fields in 4668 the cpio header (c_filesize, c_dev, c_ino, c_mode, and c_rdev) are too short to 4669 support values that POSIX.1 allows these fields to contain. Some existing imple-4670 mentations and current trends in development will require the ability to 4671 represent even larger values in these fields. The cpio format does not provide a 4672 mechanism to represent the user and group IDs symbolically, and a range of 4673 implementation-defined file types have not been reserved for the user. The cpio 4674 format specification does not reserve any formats for implementation-defined 4675 usage. The extensions that have been made to cpio for POSIX.1 are compatible 4676 with existing versions of cpio. Correction of some of these deficiencies would 4677 make existing versions of cpio behave unpredictably. When these changes are 4678 made the cpio magic number will have to be changed. 4679
- This clause uses the term *file name*; note that *filename* and *file name* are not synonyms; the latter is a synonym for *pathname*, in that it includes the slashes between filenames.
- In earlier drafts, the word "local" was used in the context of "file system" and was
 taken (incorrectly) to be related to "remotely mounted file system." This was not
 intended. The term "(local) file system" refers to the file hierarchy as seen by the
 utilities, and "local" was removed because of this confusion.

4687 B.10.1.1 Extended tar Format

The original model for this facility is the 4.3BSD or Version 7 tar program and
format, but the format given here is an extension of the traditional tar format.
The name USTAR was adopted to reflect this.

4691 This description reflects numerous enhancements over previous versions. The 4692 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. Archives written using this new format may be read by applications designed to use the old format as long as the functional enhancements provided here are not used. This means the user is limited to archiving only regular type files and nonsymbolic links to such files.
- 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 file name 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 4703 requirements for {PATH_MAX}. If a pathname will fit within the name field, it is 4704 recommended that the pathname be stored there without the use of the prefix 4705 field. Although the name field is known to be too small to contain {PATH_MAX} 4706 characters, the value was not changed in this version of the archive file format to 4707 retain backward compatibility, and instead the prefix was introduced. Also, 4708 because of the earlier version of the format, there is no way to remove the restric-4709 tion on the *linkname* field being limited in size to just that of the *name* field. 4710
- 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.
- 4714 It is suggested that if device special files need to be represented that cannot be 4715 represented in the standard format that one of the extension types ('A'-'Z') be 4716 used, and that the additional information for the special file be represented as 4717 data and be reflected in the size field.
- Attempting to restore a special file type, where it is converted to ordinary data 4718 and conflicts with an existing file name, need not be specially detected by the util-4719 ity. If run as an ordinary user, a format-reading utility should not be able to 4720 overwrite the entries in, for example, /dev in any case (whether the file is con-4721 verted to another type or not). If run as a privileged user, it should be able to do 4722 so, and it would be considered a bug if it did not. The same is true of ordinary 4723 data files and similarly named special files; it is impossible to anticipate the 4724 user's needs (who could really intend to overwrite the file), so the behavior should 4725 be predictable (and thus regular) and rely on the protection system as required. 4726
- The values '2' and '7' in the typeflag field are intended to define how symbolic links and contiguous files can be stored in a tar archive. POSIX.1 does not require the symbolic link or contiguous file extensions, but does define a standard way of archiving these 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 format-interpreting utility should do the best it can with the file and go on to the next.
- 4734 B.10.1.2 Extended cpio Format

The model for this format is the existing System V cpio -c data interchange format. This model documents the portable version of cpio format and not the binary version. It has the flexibility to transfer data of any type described within the POSIX.1 standard, yet is extensible to transfer data types specific to extensions beyond POSIX.1 (e.g., symbolic links or contiguous files). Because it describes existing practice, there is no question of maintaining upward compatibility.

This subclause does not standardize behavior for the utility when the file type is not understood or supported. It is useful for the utility to report to the user whatever action is taken in this case, though POSIX.1 neither requires nor recommends this.

4746 **B.10.1.2.1** 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 i-node 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 i-node number of the file in the location from which that file is extracted.

4752 **B.10.1.2.2 cpio File Name**

For most historical implementations of the cpio utility, {PATH_MAX} bytes can be used to describe the pathname without the addition of any other header fields (the null byte would be included in this count). {PATH_MAX} is the minimum value for pathname size, documented as 256 bytes in Section 2. However, an implementation may use $c_namesize$ to determine the exact length of the pathname. With the current description of the cpio header, this pathname size can be as large as a number that is described in six octal digits.

- 4760 **B.10.1.2.3** cpio File Data
- 4761 There is no additional rationale provided for this subclause.
- 4762 B.10.1.2.4 cpio Special Entries
- 4763 These are provided to maintain backward compatibility.
- 4764 **B.10.1.2.5** cpio Values
- Three values are documented under the c_mode field values to provide for extensibility for known file types:

4767	0110000	Reserved for contiguous files. The implementation may treat
4768		the rest of the information for this archive like a regular file. If
4769		this file type is undefined, the implementation may create the
4770		file as a regular file.
4771	0120000	Reserved for files with symbolic links. The implementation
4772		may store the link name within the data portion of the file. If
4773		this type is undefined, the implementation may not know how
4774		to link this file or be able to understand the data section. The
4775		implementation may decide to ignore this file type and output a
4776		warning message.

47770140000Reserved for sockets. If this type is undefined on the target4778system, the implementation may decide to ignore this file type4779and output a warning message.

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 format-interpreting utility should do the best it can with the file and go on to the next.

4785 **B.10.1.3 Multiple Volumes**

Multivolume archives have been introduced in a manner that has become a de 4786 facto standard in many implementations. Though it is not required by POSIX.1. 4787 classical implementations of the format-reading and -creating utility, upon read-4788 ing logical end-of-file, check to see if an error channel is open to a controlling ter-4789 minal. The utility then produces a message requesting a new medium to be made 4790 available. The utility waits for a new medium to be made available by attempting 4791 to read a message to restart from the controlling terminal. In all cases, the com-4792 munication with the controlling terminal is in an implementation-defined 4793 manner. 4794

This subclause (10.1.3) is intended to handle the issue of multiple volume archives. Since the end-of-medium and transition between media are not properly part of POSIX.1, the transition is described in terms of files; the word "file" is used in a very broad, but correct, sense—a tape drive is a file. The intent is that files will be read serially until the end-of-archive indication is encountered and that file or media change will be handled by the utilities in an implementation-defined manner.

Note that there was an issue with the representation of this on magnetic tape, 4802 and POSIX.1 is intended to be interpreted such that each byte of the format is 4803 represented on the media exactly once. In some current implementations, it is 4804 not deterministic whether encountering the end-of-medium reflector foil on mag-4805 netic tape during a write will yield an error during a subsequent read() of that 4806 record, or if that record is actually recorded on the tape. It is also possible that 4807 read() will encounter the end-of-medium when end-of-medium was not encoun-4808 tered when the data was written. This has to do with conditions where the end of 4809 [magnetic] record is in such a position that the reflector foil is on the verge of 4810 being detected by the sensor and is detected during one operation and not on a 4811 later one, or vice versa. 4812

An implementation of the format-creating utility must assure when it writes a record that the data appears on the tape exactly once. This implies that the program and the tape driver work in concert. An implementation of the formatreading utility must assure that an error in a boundary condition described above will not cause loss of data.

The general consensus was that the following would be considered as correct operation of a tape driver when end-of-medium is detected:

4820 (1) During writing, either

- 4821(a) The record where the reflector spot was detected is backspaced over
by the driver so that the trailing tape mark that will be written on
close() will overwrite. Writing the tape mark should not yield an
end-of-medium condition, or
- (b) The condition is reported as an error on the *write()* following the one where the end-of-medium is detected (the one where the endof-medium is actually detected completing successfully). No data will be actually transferred on the *write()* reporting the error. The subsequent *close()* would write a tape mark following the last record actually written. Writing the tape mark, and writing any subsequent records, should not yield any end-of-medium conditions.
- 4832[The latter behavior permits the implementation of ANSI standard labels4833because several records (the trailer records) can be written after the end-4834of-medium indications. It also permits dealing with, for example, COBOL4835"ON" statements.]
- 4836 (2) During reading, the end-of-medium indicator is simply ignored, presuming that a tape mark (end-of-file) will be recorded on the magnetic medium and that the reflector foil was advisory only to the *write()*.

4839 Systems where these conditions are not met by the tape driver should assure that 4840 the format-creating and -reading utilities assure proper representation and 4841 interpretations of the files on the media in a way consistent with the above recom-4842 mendations.

- 4843 The typical failures on systems that do not meet the above conditions are either
- (1) To leave the record written when the end-of-medium is encountered on the tape, but to report that it was not written. The format-creating utility would then rewrite it, and then the format-reading utility could see the record twice if the end-of-medium is not sensed during the read operations, or
- (2) The *write()* occurs uneventfully, but the *read()* senses the error and does not actually see the data, causing a record to be omitted.

Nothing in POSIX.1 requires that end-of-medium be determined by anything on the medium itself (for example, a predetermined maximum size would be an acceptable solution for the format-creating utility). The format-reading utility must be able to *read()* tapes written by machines that do use the whole medium, however.

4856 On media where end-of-medium and end-of-file are reliably coincident, such as 4857 disks, end-of-medium and end-of-file can be treated as synonyms.

Note that partial physical records [corresponding to a single *write()*] can be written on some media, but that only full physical records will actually be written to magnetic tape, given the manner in which the tape operates.

Annex C (informative)

Header Contents Samples

1 The material in this informative annex serves as an index to which symbols 2 should appear in which headers in a system that conforms to POSIX.1 with C 3 Standard Language-Dependent System support.

This is only an index, and any conflicts with the actual body of any relevant standard shall be resolved in favor of that standard. The actual body of the declaration was omitted in part because this is an index and in part to avoid any possible conflict with the standards.

8 Where it is known that a symbol or header is not required for Common Usage C 9 Language-Dependent System support, the name is followed by an asterisk (*). 10 Omission of an asterisk does not imply that the symbol is required for Common-11 Usage C. For Common-Usage C, although the location of symbols is typical, it is 12 not to be taken as a requirement: POSIX.1 is quite explicit that there is no 13 requirement except that differences from the C Standard (2) be documented.

Generally, where it is stated that functions are defined in a header, macros are permitted as acceptable alternatives by both standards. See the bodies of the standards for details.

- 17 <assert.h>
- 18 The header defines the macro
- 19 *assert()*
- 20 and makes reference to the macro
- 21 NDEBUG
- 22 <ctype.h>
- 23 The header declares the functions

24	isalnum()	isdigit()	islower()	ispunct()	isupper()	tolower()
25	isalpha()	isgraph()	isprint()	isspace()	isxdigit()	toupper()
26	iscntrl()					

|

27	<dirent.h></dirent.h>									
28	The header defines the typedef									
29	DIR									
30	and declares the structure									
31	dirent									
32	with structure element member									
33	d_name									
34	and declares functions									
35	closedir()	opendir() read	dir()	rewindd	lir()				
36	<errno.h></errno.h>									
37	The header defines	the macros	5							
	Dopra				-	-				
38	EZBIG	EEXIST	EMLINK		ENOM.	EM	EPERM			
39	EACCES	EFAULT	ENAMETO	JOLONG	ENOSE	νC	EPIPE			
40	EAGAIN	EFBIG	ENFILE		ENOSY		ERANGE			
41	EBADF	EINTR	ENODEV		ENOTI	DIR	EROFS			
42	EBUSY	EINVAL	ENOENT		ENOTE	EMPTY	ESPIPE			
43	ECHILD	EIO	ENOEXEC	3	ENOTI	Y	ESRCH			
44	EDEADLK	EISDIR	ENOLCK		ENXIO		EXDEV			
45	EDOM	EMFILE								
46	and declares the ex	xternal vari	able							
47	errno									
48	<fcntl.h></fcntl.h>									
49	The header defines	the macros	5							
50	FD_CLOEXE	C F_SE	FFD	O_ACCM	ODE	O_NONI	BLOCK			
51	F_DUPFD	F_SE	FFL	O_APPEN	٧D	O_RDON	1LY			
52	FGETFD	FSE	FLK	O CREAT	Г	O RDWI	R			
53	FGETFL	FSE	FLKW	OEXCL		O TRUN	IC			
54	F GETLK	F UN	LCK	O NOCT	ΓY	O WROI	NLY			
55	F_RDLCK	F_WR	LCK	0_11001		0				
56	and declares the st	ructure								
57	flock									
58	with structure eler	nents								
59	l_len l_pid l_start l_type l_whence									
--	--									
60	and the functions									
61	creat() fcntl() open()									
62	and may contain the macros									
63 64 65 66	SEEK_CUR S_IROTH S_IRWXU S_ISFIFO S_IWGRP S_IXGRP SEEK_END S_IRUSR S_ISBLK S_ISGID S_IWOTH S_IXOTH SEEK_SET S_IRWXG S_ISCHR S_ISREG S_IWUSR S_IXUSR S_IRGRP S_IRWXO S_ISDIR S_ISUID									
67	<float.h>*</float.h>									
68	The header defines the macros									
69 70 71 72 73 74 75 76 77 78	DBL_DIG*FLT_EPSILON*LDBL_DIG*DBL_EPSILON*FLT_MANT_DIG*LDBL_EPSILON*DBL_MANT_DIG*FLT_MAX*LDBL_MANT_DIG*DBL_MAX*FLT_MAX_10_EXP*LDBL_MAX*DBL_MAX_10_EXP*FLT_MAX_EXP*LDBL_MAX_10_EXP*DBL_MAX_EXP*FLT_MIN*LDBL_MAX_EXP*DBL_MIN*FLT_MIN_10_EXP*LDBL_MIN*DBL_MIN_10_EXP*FLT_MIN_EXP*LDBL_MIN_10_EXP*FLT_DIG*FLT_RADIX*LDBL_MIN_EXP*									
79	<grp.h></grp.h>									
80	The header declares the structure									
81	group									
82	with structure elements									
83	gr_gid gr_mem gr_name									
84	and the functions									
85	getgrgid() getgrnam()									

I

86 <limits.h>

87 The header defines the macros

88	ARG_MAX•	NGROUPS_MAX	USHRT_MAX
89	CHAR_BIT	OPEN_MAX•	_POSIX_ARG_MAX
90	CHAR_MAX	PATH_MAX•	_POSIX_CHILD_MAX
91	CHAR_MIN	PIPE_BUF•	_POSIX_LINK_MAX
92	CHILD_MAX•	SCHAR_MAX	_POSIX_MAX_CANON
93	INT_MAX	SCHAR_MIN	_POSIX_MAX_INPUT
94	IN T_MI N	SHRT_MAX	_POSIX_NAME_MAX
95	LINK_MAX•	SHRT_MIN	_POSIX_NGROUPS_MAX
96	LONG_MAX	SSIZE_MAX	_POSIX_OPEN_MAX
97	LONG_MIN	STREAM_MAX□	_POSIX_PATH_MAX
98	MAX_CANON•	TZNAME_MAX	_POSIX_PIPE_BUF
99	MAX_INPUT•	UCHAR_MAX	_POSIX_SSIZE_MAX
100	MB_LEN_MAX	UINT_MAX	_POSIX_STREAM_MAX
101	NAME_MAX•	ULONG_MAX	_POSIX_TZNAME_MAX

The macros marked with \Box shall be omitted from <limits.h> on specific implementations where the corresponding value is greater than or equal to the stated minimum, but is indeterminate. The macros marked with • shall be omitted from <limits.h> on specific implementations where the corresponding value is greater than or equal to the stated minimum, but where the value can vary depending on the file to which it is applied.

- 108 <locale.h>
- 109 The header defines the macros

110	LC_ALL*	LC_CTYPE*	LC_NUMERIC*	NULL*
111	LC_COLLATE*	LC_MONETARY*	LC_TIME*	

112 and declares the structure

113 *lconv**

114 with structure elements

115	currency_symbol*	mon_decimal_point*	negative_sign*
116	decimal_point*	mon_grouping*	$p_cs_precedes*$
117	frac_digits*	mon_thousands_sep*	p_sep_by_space*
118	grouping*	$n_cs_precedes*$	p_sign_posn*
119	int_curr_symbol*	n_sep_by_space*	positive_sign*
120	int_frac_digits*	n_sign_posn*	thousands_sep*

121 and the functions

122 *localeconv()* setlocale()**

	Part 1: SYSTEM API [C LANGUAGE]					IEEE Std 1003.1-1990		
123	<math.h></math.h>							
124	The header defines the macro							
125	HUGE_VAL							
126	and declares the func	tions						
127	acos() ceil() $exp()$	fmod()	log10()	pow()	sart()		
128	asin() cos() $fabs()$	frexp()	log()	sin()	tan()		
129	atan2() cosh	() floor()	ldexp()	modf()	sinh()	tanh()		
130	atan()	() 1 ()						
131	<pwd.h></pwd.h>							
132	The header defines th	e structure						
133	passwd							
134	with structure element	nts						
135	pw_dir pu	v_gid pw_	name pw_	_shell pw	_uid			
136	and declares the func	tions						
137	getpwnam()	getpwuid()						
138	<setjmp.h></setjmp.h>							
139	The header defines th	e types						
140	jmp_buf s	igjmp_buf						
141	and declares the func	tions						
142	longjmp()	setjmp()	siglongjm	p() sigsetj	imp()			
143	Note that the C Star	dard [2] and	this part of	f 190/1FC 99	15 hoth	normit those		
140	functions to be define	d sololy as ma	cros	1 150/160 55	40 0000	permit these		
144	<pre><signal.h></signal.h></pre>	u solely as ma	ci 05.					
146	The header defines th	e macros						
147	SA NOCLDSTO	P SIGHUP	SIGOUIT	SIGTTIN	SIG	DFL		
148	SIGARRT	SIGUL	SIGSEGV	SIGTTOU	SIG	ERR*		
149	SIGALRM	SIGINT	SIGSTOP	SIGUSRI	SIG	IGN		
150	SIGCHLD	SIGKILL	SIGTERM	SIGUSR2	SIG	SETMASK		
151	SIGCONT	SIGPIPE	SIGTSTP	SIG BLOO	CK SIG	UNBLOCK		
152	SIGFPE		NIGIDII	510_5100	···			

153 and the types

ISO/IEC 9945-1: 1990

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

I

I

I

154	sig_atomic_t* sig	set_t			
155	and declares the structure				
156	sigaction				
157	with structure elements				
158	sa_flags sa_ha	undler sa_m	ask		
159	and the functions				
160	kill() sig	addset()	sigfillset()	signending	J()
161	raise()* sie	delset()	sigismember()	sigprocma	sk()
162	sigaction() sig	gemptyset()	signal()*	sigsuspend	<i>d</i> ()
163	<stdarg.h>*</stdarg.h>				
164	The header defines the ma	acros			
165	va_arg* va_end	* va_list*	va_start*		
166	<stddef.h>*</stddef.h>				
167	The header defines the ma	acros			
168	NULL* offsetof*				
169	and the types				
170	ptrdiff_t* size_t*	wchar_t	*		
171	<stdio.h></stdio.h>				
172	The header defines the ma	acros			
173	BUFSIZ	L tmpnam*	STREAM	MAX*	stdout
174	EOF	NULL	TMP MA	X	IOFBF*
175	FILENAME MAX*	SEEK CUR	stderr		IOLBF*
176	L_ctermid	SEEK_END	stdin		_IONBF*
177	L_cuserid	SEEK_SET			
178	NOTE: The L_cuserid symbol is	s permitted in thi	s header, but need no	t be supplied.	See 2.7.2.
179	and the types				
180	fpos_t* size_t				

181 and declares the type

FILE

182

and the functions 183 clearerr() putchar() fileno() fsetpos()* sprintf() 184 fclose() fopen() ftell() 185 puts() sscanf() fdopen() fprintf() *fwrite()* remove() *tmpfile()* 186 feof() fputc() getc() rename() 187 tmpnam() *ferror()* fputs() getchar() rewind() ungetc() 188 fflush() fread() gets() scanf() vfprintf()* 189 190 fgetc() freopen() perror() setbuf() vprintf()* fscanf() fgetpos()* printf() setvbuf()* vsprintf()* 191 fgets() fseek() putc() 192 <stdlib.h> 193 The header defines the macros 194 EXIT_FAILURE MB_CUR_MAX* RAND_MAX 195 NULL EXIT_SUCCESS 196 and the types 197 div t* ldiv t* wchar_t* 198 size t and declares the functions 199 bsearch() abort() labs()* qsort() strtol()* 200 abs() calloc() ldiv()* rand() strtoul()* 201 malloc() realloc() atexit()* div()*system()* 202 mblen()* srand() wcstombs()* atof() exit() 203 strtod()* wctomb()* atoi() free() mbstowcs()* 204 mbtowc()* atol() getenv() 205 <string.h> 206 The header defines the macro 207 NULL 208 and the type 209

210 *size_t*

ł

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

I

I

1

I

1

211	and declares the functions				
212 213	memchr()* strc. memcmp()* strc.	at() stro hr() stre	spn() rror()*	<pre>strncmp() strncpy()</pre>	strspn() strstr()
214	memcpy()* strc	mp() strl	en()	strpbrk()	strtok()
215	memmove()* strc	oll()* strr	ncat()	strrchr()	strxfrm()*
216	memset()* strc	by()			
217	<sys stat.h=""></sys>				
218	The header defines the ma	cros			
219	S IRGRP S IRW	KO S ISCHR	S ISGID	S IWGRP	S IXGRP
220	S IROTH S IRW	KU S ISDIR	S ISREG	SIWOTH	SIXOTH
221	S IRUSR S ISBL	K S ISFIFO	S ISUID	S IWUSR	S IXUSR
222	S_IRWXG		210012	~	2
	_				
223	and declares the structure				
224	stat				
225	with structure elements				
22 6	st atime st dev	st ino	st mtime	st size	
227	st_ctime st_gid	st_mode	st_nlink	st_uid	
228	and the functions				
229	chmod() mkdir()	stat()			
230	fstat() mkfifo()	umask()			
231	<sys times.h=""></sys>				
232	The header defines the typ	e			
233	$clock_t$				
234	and declares the structure				
235	tms				
236	with structure elements				

237 tms_cstime tms_cutime tms_stime tms_utime

and the function

239 times()

240	<sys th="" types.h<=""><th>1></th><th></th><th></th><th></th></sys>	1>			
241	The header defin	es the types			
	-				
242	dev_t	ino_t nlink_t	pid_t ssi	ze_t	
243	gid_t	mode_t off_t	sıze_t uıa	t_t	
244	<sys td="" utsname<=""><td>a.h></td><td></td><td></td><td></td></sys>	a.h>			
245	The header decla	res the structure			
246	utsname				
	with structure of				
247	with structure el	ements			
248	machine	nodename r	elease svsn	ame version	
249	and the function				
250	uname()				
951	cove/wait h				
201	The second				
252	The header defin	es the macros			
959	WEXITSTA	TUS WIFSIGN		NG WTFR	MSIG
253	WIFEXITE	D WIFSTOP	PED WSTOPS	NG WINT	RACED
	,,				
255	and declares the	functions			
	••••	• • • • •			
256	wait()	waitpid()			
257	<termios.h></termios.h>				
258	The header defin	es the macros			
259	BO	B75	ECHONL	NCCS	TCSAFLUSH
26 0	B110	B9600	HUPCL	NOFLSH	TCSANOW
261	B1200	BRKINT	ICANON	OPOST	TOSTOP
262	B134	CLOCAL	ICRNL	PARENB	VEOF
263	B150	CREAD	IEXTEN	PARMRK	VEOL
264	B1800	CS5	IGNBRK	PARODD	VERASE
265	B19200	CS6	IGNCR	TCIFLUSH	VINTR
266	B200	CS7	IGNPAR	TCIOFF	VKILL
267	B2400	CS8	INLCR	TCIOFLUSH	VMIN
268	B300	CSIZE	INPCK	TCION	VQUIT
269	B38400	CSTOPB	ISIG	TCOFLUSH	VSTART
270	B4800	ECHO	ISTRIP	TCOOFF	VSTOP
271	B50	ECHOE	IXOFF	TCOON	VSUSP
272	B600	ECHOK	IXON	TCSADRAIN	VTIME

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

INFORMATION TECHNOLOGY—POSIX

273	and the types
274	cc_t speed_t $tcflag_t$
275	and declares the structure
276	termios
277	with structure elements
278	c_cc c_cflag c_iflag c_lflag c_oflag
279	and the functions
280 281 282	cfgetispeed()cfsetospeed()tcflush()tcsendbreak()cfgetospeed()tcdrain()tcgetattr()tcsetattr()cfsetispeed()tcflow()tcflow()tcsetattr()
283	<time.h></time.h>
284	The header defines the macros
285	CLK_TCK CLOCKS_PER_SEC NULL
286	the types
287	clock_t size_t time_t
288	and declares the structure
289	tm
290	with structure elements
291 292	tm_hour tm_mday tm_mon tm_wday tm_year tm_isdst tm_min tm_sec tm_yday
293	and the functions
294 295	asctime()ctime()gmtime()mktime()time()clock()*difftime()*localtime()strftime()tzset()
296	and declares the external variable
297	tzname

1

298 <unistd.h>

299 The header defines the macros

300	F_OK	_PC_PIPE_BUF
301	NULL	_PC_VDISABLE
302	R_OK	_POSIX_CHOWN_RESTRICTED
303	SEEK_CUR	_POSIX_JOB_CONTROL
304	SEEK_END	_POSIX_NO_TRUNC
305	SEEK_SET	_POSIX_SAVED_IDS
306	STDERR_FILENO	_POSIX_VDISABLE
307	STDIN_FILENO	_POSIX_VERSION
308	STDOUT_FILENO	_SC_ARG_MAX
309	W_OK	_SC_CHILD_MAX
310	X_OK	_SC_CLK_TCK
311	_PC_CHOWN_RESTRICTED	_SC_JOB_CONTROL
312	_PC_LINK_MAX	_SC_NGROUPS_MAX
313	_PC_MAX_CANON	_SC_OPEN_MAX
314	_PC_MAX_INPUT	_SC_SAVED_IDS
315	_PC_NAME_MAX	_SC_STREAM_MAX
316	_PC_NO_TRUNC	_SC_TZNAME_MAX
317	_PC_PATH_MAX	_SC_VERSION

318 and defines the types

319 size_t* ssize_t*

320 and declares the functions

321	$_exit()$	execl()	geteuid()	link()	setsid()
322	access()	execle()	getgid()	lseek()	setuid()
323	alarm()	execlp()	getgroups()	pathconf()	sleep()
324	chdir()	execv()	getlogin()	pause()	sysconf()
325	chown()	execve()	getpgrp()	pipe()	tcgetpgrp()
326	close()	execvp()	getpid()	read()	tcsetpgrp()
327	ctermid()	fork()	getppid()	rmdir()	ttyname()
328	cuserid()	fpathconf()	getuid()	setgid()	unlink()
329	dup2()	getcwd()	isatty()	setpgid()	write()
330	dup()	getegid()			

331 NOTE: The *cuserid()* symbol is permitted in this header, but need not be supplied. See 2.7.2.

- 332 <utime.h>
- 333 The header declares the structure
- 334 utimbuf
- 335 with structure elements

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

336 actime modtime

337 and the function

338 *utime()*

1

1

1

Annex D (informative)

Profiles

This standard contains a number of options and variables that reflect the range of 1 systems and environments that might be encountered. In general, it will be use-2 ful for applications to take the full range of these possibilities into account and 3 either accommodate them or exclude them. However, there are significant com-4 munities of interest that may have common needs that warrant focusing on a 5 specific suite of these options and parameters. This annex discusses the concept 6 of profiles (also known as functional standards) and how they address this 7 problem. 8

9 This annex reflects current thinking. It is clear that a concept such as this will 10 help significantly in clarifying the intended use of these standards. It is to be 11 expected that some of the details of this material will be changed before it is fully 12 stabilized.

As background: the OSI model has over 170 standards (and consequent combinations thereof) that fit within it. Only a fraction of those are actually useful for any given application environment. The concept of *profiles* was developed to address this issue and appears also to apply to the area of application portability. The ISO/IEC term for such profiles is ISP, or "International Standardized Profile."

18 **D.1 Definitions**

The following definitions are proposed for use in the area covered by this part of ISO/IEC 9945.

D.1.1 Applications Environment Profile (AEP) [profile]: The specification of a complete and coherent subset of an Open System Environment, together with the options and parameters necessary to support a class of applications for interoperability or applications portability, including consistency of data access and human interfaces. Where there are several AEPs for the same OSE, they are harmonized.

AEPs are the basis for procurement and conformance testing and are the target environment for software development.

D.1.2 Application Specific Environment (ASE): A complete and coherent
 subset of an Applications Environment Profile, together with interfaces, services,
 or supporting formats outside of the profile, that are required by a particular

32 application for its installation and execution.

D.1.3 Application Specific Environment Description (ASED): The specification of an Application Specific Environment, together with the specific options or parameters required; interfaces, services, or supporting formats outside of the profile; and resource requirements necessary for the satisfactory operation of the application. (For example, storage and performance requirements.)

- (This term is intended for use in Applications Conformance clauses found in
 profiles.)
- 40 **D.1.4 coherent:** The parts are logically connected. (For example, if both FOR-41 TRAN and COBOL are specified, whether they can share files is specified.)

42 **D.1.5 complete:** Having all the necessary parts. (For example, if COBOL and 43 SQL are both specified, then there is a COBOL binding to SQL, or at least an expla-44 nation of why not.)

45 **D.1.6 comprehensive:** A sufficiently broad range of functionality is covered 46 that the needs of most Applications Environment Profiles are met.

D.1.7 consistent: The parts of the Open System Environment do not inherently conflict with each other. This does not preclude options that conflict, as long as an Applications Environment Profile can select a set that does not conflict.

50 **D.1.8 harmonized:** Where same functionality is needed in several profiles, it 51 appears identically in all of them.

D.1.9 Open System Environment (OSE): A comprehensive and consistent set of international information technology standards and functional standards (profiles) that specify interfaces, services, and supporting formats to accomplish interoperability and portability of applications, data, and people. These are based on International Standards (ISO, IEC, CCITT, ...)

D.1.10 POSIX Open System Environment: A comprehensive and consistent set of ISO/IEC, regional, and national information technology standards and functional standards (profiles) that specify interfaces, services, and supporting formats for interoperability and portability of applications, data, and people that are in accord with ISO/IEC 9945 (POSIX).

No single component of the OSE, including ISO/IEC 9945, is expected to be required in all such profiles.

64 **D.2 Options in This Part of ISO/IEC 9945**

In terms of this part of ISO/IEC 9945, there are a number of features that could be specified in a profile. This list includes:

- The options listed in 1.3.1.3.
- The limits in 2.8. Regarding the the C Language Limits for the type char,
 care should be taken that those limits are not for the POSIX.1 definition of
 character, but for the one in the C language. For the POSIX.1 definition of

- *character*, the following limits from the C Standard (2) could be specified as
 well: {MB_LEN_MAX} and {MB_CUR_MAX}.
- The flags in 2.9.4.
- Instances of the word "may" throughout the document. (Note that not all instances of "may" constitute behavior that could or should be considered appropriate for specification in a profile. Some reflect implementation variants that should not matter to applications.)
- Features that are specified in a generic way for broad portability of the standard, that might reasonably be constrained in the more limited context of a profile. For such features, the constraint shall be documented in a profile so that the effects of the constraint on the standard would be clarified.

83 D.3 Related Standards

The other POSIX standards (ISO/IEC 9945-2 {B36}, in particular) are expected to form a major part of the POSIX OSE. Other formal standards, such as those listed in Annex A, are also expected to be part of such a document (in particular, the C Standard {2}.)

Standards such as other languages, SQL, graphics standards such as GKS, and
 networking standards are also probable candidates for inclusion in the POSIX
 OSE.

91 D.4 Related Activities

In many ways, the work of NIST (in terms of FIPS), OSF, UNIX International, and X/Open often act like early (but sophisticated) profiles or metaprofiles, specifying a range of standards from which true profiles could select. They collect together many standards, specify options, and specify the relationship between the parts. These activities go well beyond profiles, as they add specifications that are not formal standards to the suite as well. Often these additional specifications point to areas where formal standards are required.

99 **D.5 Relationship to IEEE Draft Project 1003.0**

The IEEE P1003.0 working group is writing a *Guide to Open Systems*. In many ways, this is the specification of the POSIX Open Systems Environment. It could be the specification of the collection of standards that might be used to specify many Open Systems Environments, depending on how exactly that work proceeds.



Annex E (informative)

Sample National Profile

One class of "community of interest" for which profiles (as discussed in Annex D) are useful is specific countries, where the general characteristics warrant specific focus to serve the needs of users in those countries. Such needs lead to a number of implications concerning the options available within this part of ISO/IEC 9945 and may warrant specification of complementary standards as well.

The following is an example of a country's needs with respect to this part of ISO/IEC 9945 and how those needs relate to other international standards as well as national standards. The example provided is included here for informative purposes and is not a formal standard in the country in question. It is provided by the Danish Standards Association and is as accurate as possible with regards to Danish needs.¹⁾ This example national profile is worded as if it were a national standard.

A subclass of conforming implementations can be identified that meet the requirements of a specific profile. By documenting these either in national standards, in a document similar to an ISO/IEC ISP (an International Standardized Profile), or in an informative annex (such as this), these can be referenced in a consistent manner.

Further information may be obtained from the Danish Standards Association, Attn: S142u22A11
 POSIX WG, Box 77, DK-2900 Hellerup, Denmark; FAX: +45 31 62 30 77; Email: posix@itc.dk

²⁰ The data is also available electronically by anonymous FTAM or FTP at the site dkuug.dk in the 21 directory isp, where some other example national profiles, locales, and *charmaps* may also be 22 found. They are also available by an archive server reached at archive@dkuug.dk; use 23 "Subject: help" for further information.

²⁴ More complete examples of profiles are expected to be available in future revisions of this part of 25 ISO/IEC 9945 and in other POSIX standards.

E.1 (Example) Profile for Denmark

NOTE: This profile is chosen both for its instructive value by being a European profile and the generality in the provisions it makes, addressing most of the relevant points. It does claim to be correct for Denmark, and the style is what would be expected in a real ISP. A collection of real ISPs would be as useful, and work is underway collecting these.

This is the definition of the Danish Standards Association POSIX.1 profile. Information on the actual data for the locale and coded character set mapping definitions are under development as part of an informative annex in ISO/IEC 9945-2 (B36).²⁾

The subset of conforming implementations that provide the required characteristics below is referred to as conforming to the "Danish Standards Association (DS) Environment Profile" for this part of ISO/IEC 9945.

The profile specifies no options according to POSIX.1 section 2.9.3. For section 2.8.4 in the <limits.h> specification, the {_POSIX_TZNAME_MAX} value shall be 7.

41 E.1.1 Character Encoding

- Any character encoding with the required repertoire of the POSIX profile plus the following repertoire shall be allowed.
- A "character set description file," as described in ISO/IEC 9945-2, {B36} shall use the symbolic character names of the ISO_10646 *charmap* file described in the ISO/IEC 9945-2 {B36} sample profile annex for the characters encoded in the character set.

For the Danish and Greenlandic languages, the following characters shall be present in addition to the repertoire required by the POSIX profile: <ae>, <o/>, <ae>, <AE>, <O/>, and <AA>. For the Faroese language, the following characters shall be present in addition to the required POSIX locale characters and Danish repertoire: <a'>, <i'>, <o'>, <u'>, <y'>, <d->, <A'>, <I'>, <O'>, <U'>, <Y'>, and <D->.

Recommended character sets are ISO 8859-1 (B34) or ISO 10646 (B37). The CHARSET environment variable shall be used to specify the processing character set; for instance, ISO_8859-1 or ISO_10646. This shall be used to select the encoded character-set-specific versions of the locale definitions. If no CHARSET variable is present, ISO_8859-1 shall be assumed.

59 60

²⁾ The 9945-2 document, "POSIX.2," is currently in the state of a Committee Document (CD), to be approved as a Draft International Standard.

61 E.1.2 Character Encoding and Display

For terminal equipment not capable of generating or showing the processing char-62 acter set, the character names defined in the current charmap file shall be used: 63 characters in the charmap file having two-character names shall be specified by 64 the two-character name preceded by the <intro> character, and characters hav-65 ing *charmap* names longer than two characters shall be specified by the character 66 name preceded by the <intro> character and an <underline> and followed by 67 an <underline>. In names longer than two characters, an <intro> character 68 and an <underline> character in sequence shall signify a literal <underline> 69 character part of the character name. Two <intro> characters in sequence shall 70 signify one <intro> character, both in names and in the general stream. 71

For input, if the character name is undefined in the current *charmap* file, the data shall be left untouched (including the <intro> character) and the behavior is implementation defined.

75 E.1.3 Locale Definitions

The following guideline is used for specifying the locale identification string:³⁾

77 78

94

"%2.2s_%2.2s.%s,%s", <language>, <territory>, <character-set>, <version>

where *<language>* shall be taken from ISO 639 (B32) and *<territory>* shall be the 79 two-letter country code of ISO 3166 (B33), if possible. The <language> shall be 80 specified with lowercase letters only, and the *<territory>* shall be specified in 81 uppercase letters only. An optional *<character-set>* specification may follow after 82 a <period> for the name of the character set; if just a numeric specification is 83 present, this shall represent the number of the international standard describing 84 the character set. If the <character-set> specification is not present, the encoded 85 character set specific locale shall be determined by the CHARSET environment 86 variable, and if this is unset or null, the encoding of ISO 8859-1 (B34) shall be 87 assumed. A parameter specifying a *<version>* of the profile may be placed after 88 the optional <character-set> specification, delimited by <comma>. This may be 89 used to discriminate between different cultural needs; for instance, dictionary 90 order versus a more systems-oriented collating order. 91

Following the above guidelines for locale names, the national Danish locale stringshall be

da_DK

In the following, the TZ variable shall be specified according to the current official
 daylight-saving-time rules in Denmark. Since Daylight Saving Time is politically
 decided and thus changeable, this is only a recommendation.

3) The guideline was inspired by the X/Open Portability Guide (B61). It is presented in the file
 format notation used by ISO/IEC 9945-2 (B36).

100	The locale definition for Denmark shall be as follows:	
101	01 LANG da_DK	
102	02 TZ CET-1CET DST, M3.5.0/M9.5.0	
103	⁰³ The locale definition for the Faroe Islands shall be as follows:	
104	04 LANG fo_DK	
105	05 TZ UTCOUTC DST, M3.5.0/M9.5.0	
106	⁰⁶ The locale definition for Western Greenland shall be as follows:	
107	07 LANG kl_DK	
108	08 TZ UTZ+3VTZ,M3.5.0/M9.5.0	
109	⁰⁹ The locale definition for Eastern Greenland shall be as follows:	
110	10 LANG kl DK	
111	11 TZ VTZ+2WTZ, M3.5.0/M9.5.0	
112	¹² For the Faroe Islands and Greenland, only the LC_TIME and LC_M	ESSAGES

data are different from the Danish language specifications.

Identifier Index

abort()	Referenced C Language Routines {8.1}	151
abs()	Referenced C Language Routines {8.1}	151
access()	Check File Accessibility {5.6.3}	104
acos()	Referenced C Language Routines {8.1}	151
alarm()	Schedule Alarm {3.4.1}	63
asctime()	Extensions to Time Functions {8.1.1}	152
asin()	Referenced C Language Routines {8.1}	151
assert()	Referenced C Language Routines {8.1}	151
atan()	Referenced C Language Routines {8.1}	151
atan2()	Referenced C Language Routines (8.1)	151
atof()	Referenced C Language Routines (8.1)	151
atoi()	Referenced C Language Routines {8.1}	151
atol()	Referenced C Language Routines {8.1}	151
bsearch()	Referenced C Language Routines {8.1}	151
calloc()	Referenced C Language Routines {8.1}	151
ceil()	Referenced C Language Routines {8.1}	151
cfgetispeed()	Baud Rate Functions {7.1.3}	141
cfgetospeed()	Baud Rate Functions {7.1.3}	141
cfsetispeed()	Baud Rate Functions (7.1.3)	141
cfsetospeed()	Baud Rate Functions (7.1.3)	141
chdir()	Change Current Working Directory {5.2.1}	86
chmod()	Change File Modes (5.6.4)	106
chown()	Change Owner and Group of a File (5.6.5)	107
clearerr()	Referenced C Language Routines (8.1)	151
close()	Close a File {6.3.1}	115
closedir()	Directory Operations {5.1.2}	83
cos()	Referenced C Language Routines [8.1]	151
cosh()	Referenced C Language Routines (8.1)	151
cpio	Extended cpio Format (10.1.2)	173
creat()	Create a New File or Rewrite an Existing One {5.3.2}	91
ctermid()	Generate Terminal Pathname {4.7.1}	78
ctime()	Referenced C Language Routines [8.1]	152
dev_t	Primitive System Data Types {2.5}	27
directory	Directory Operations (5.1.2)	83
<pre><dirent.h></dirent.h></pre>	Format of Directory Entries {5.1.1}	83
dup()	Duplicate an Open File Descriptor {6.2.1}	114
dup2()	Duplicate an Open File Descriptor {6.2.1}	114
environ	Execute a File (3.1.2)	42
errno	Error Numbers {2.4}	23
<errno.h></errno.h>	Error Numbers {2.4}	23
exec	Execute a File {3.1.2}	42
execl()	Execute a File {3.1.2}	42
execle()	Execute a File (3.1.2)	42
execlp()	Execute a File (3.1.2)	42
execv()	Execute a File {3.1.2}	42
execve()	Execute a File (3.1.2)	42

execup()	Execute a File {3.1.2}	42
exit()	Referenced C Language Routines [8.1]	151
exit()	Terminate a Process {3.2.2}	49
exp()	Referenced C Language Routines [8.1]	151
fabs()	Referenced C Language Routines [8.1]	151
fclose()	Referenced C Language Routines [8,1]	151
fcntl()	File Control {6.5.2}	121
<fcntl.h></fcntl.h>	Data Definitions for File Control Operations (6.5.1)	121
fdonen()	Open a Stream on a File Descriptor {8.2.2}	157
feof()	Referenced C Language Routines (8.1)	151
ferror()	Referenced C Language Routines (8.1)	151
fflush()	Referenced C Language Routines (8.1)	151
facto()	Referenced C Language Routines (8.1)	151
facto()	Referenced C Language Routines (8.1)	151
flong()	Mon a Stream Dointon to a File Deceminton (9.2.1)	156
fliend()	Paferenced C Language Deutinos (9.1)	150
f(00r())	Referenced C Language Routines (0.1)	151
(moa()	Referenced C Language Routines (0.1)	101
fopen()	Referenced C Language Routines (8.1)	191
forr()	Process Creation {3.1.1}	41
fpathconf()	Get Configurable Pathname Variables (5.7.1)	110
fprintf()	Referenced C Language Routines {8.1}	151
fputc()	Referenced C Language Routines {8.1}	151
fputs()	Referenced C Language Routines [8.1]	151
fread()	Referenced C Language Routines (8.1)	151
free()	Referenced C Language Routines (8.1)	151
freopen()	Referenced C Language Routines (8.1)	151
frexp()	Referenced C Language Routines (8.1)	151
fscanf()	Referenced C Language Routines (8.1)	151
fseek()	Referenced C Language Routines (8.1)	151
fstat()	Get File Status {5.6.2}	103
ftell()	Referenced C Language Routines (8.1)	151
fwrite()	Referenced C Language Routines (8.1)	151
getc()	Referenced C Language Routines (8.1)	151
getchar()	Referenced C Language Routines (8.1)	151
getcwd()	Get Working Directory Pathname {5.2.2}	87
getegid()	Get Real User, Effective User, Real Group, and Effective	
0 0 0	Group IDs {4.2.1}	68
getenv()	Environment Access {4.6.1}	77
geteuid()	Get Real User, Effective User, Real Group, and Effective	
0	Group IDs {4.2.1}	68
getgid()	Get Real User, Effective User, Real Group, and Effective	
0000000	Group IDs {4 2 1}	68
getgraid()	Group Database Access (9 2 1)	166
gotgram()	Group Database Access (9.2.1)	166
getgroune()	Get Supplementary Group IDc [1 23]	70
getgroups()	Get User Name (4.2.4)	71
getnarn()	Cot Process Crown ID (4.2.1)	71 79
geipgip()	Cot Process and Porcest Process IDs (4.1.1)	67
getpia()	Cet Process and Parent Process IDs (4.1.1)	01
getpunger ()	User Detabase Asses (0.0.0)	167
geipwnam()	User Database Access (9.2.2)	107
getpwuid()	User Database Access (9.2.2)	101

gets()	Referenced C Language Routines {8.1}	151
getuid()	Get Real User, Effective User, Real Group, and Effective	
	Group IDs {4.2.1}	68
gid_t	Primitive System Data Types {2.5}	27
gmtime()	Referenced C Language Routines {8.1}	152
<grp.h></grp.h>	Group Database Access (9.2.1)	166
ino_t	Primitive System Data Types (2.5)	27
isalnum()	Referenced C Language Routines [8.1]	151
isalpha()	Referenced C Language Routines (8.1)	151
isatty()	Determine Terminal Device Name {4.7.2}	79
iscntrl()	Referenced C Language Routines [8.1]	151
isdigit()	Referenced C Language Routines (8.1)	151
isgraph()	Referenced C Language Routines {8.1}	151
islower()	Referenced C Language Routines [8.1]	151
isprint()	Referenced C Language Routines [8.1]	151
ispunct()	Referenced C Language Routines [8.1]	151
isspace()	Referenced C Language Routines [8.1]	151
isupper()	Referenced C Language Routines [8.1]	151
isxdigit()	Referenced C Language Routines {8.1}	151
kill()	Send a Signal to a Process {3.3.2}	56
ldexn()	Referenced C Language Routines {8.1}	151
imits.h>	Numerical Limits (2.8)	34
link()	Link to a File $\{5,3,4\}$	92
localtime()	Referenced C Language Routines (8.1)	152
log()	Referenced C Language Routines (81)	151
log()	Referenced C Language Routines (8.1)	151
longimn()	Referenced C Language Routines (8.1)	151
least()	Reposition Read/Write File Offset (6.5.3)	101
main()	Description (3.1.2.2)	43
malloc()	Referenced C Language Boutines (8.1)	151
mkdir()	Make a Directory (5.4.1)	94
mbfifo()	Make a FIFO Special File (5.4.2)	95
mbtime()	Referenced C Language Routines (8.1)	152
mode t	Primitivo System Data Tymos (2.5)	97
mode_i	Referenced C Language Routines (8.1)	151
nlinh t	Drimitivo System Dato Trmas (2.5)	97
off t	Drimitive System Data Types (2.5)	21
onen()	Open a Fila (5.2.1)	00
open()	Directory Operations (5.1.9)	00
openuir ()	Cot Configurable Dathgame Variables (5.7.1)	110
paincon()	Get Configurable Patnname variables (5.7.1)	110
pause()	Defense l O Le process Execution (3.4.2)	151
perror()	Referenced C Language Routines {8.1}	101
pia_t	Primitive System Data Types {2.5}	27
pipe()	Create an Inter-Process Channel (6.1.1)	113
pow()	Referenced C Language Routines (8.1)	151
printf()	Referenced C Language Routines {8.1}	151
putc()	Reterenced C Language Routines (8.1)	151
putchar()	Referenced C Language Routines {8.1}	151
puts()	Referenced C Language Routines {8.1}	151
<pwd.h></pwd.h>	User Database Access {9.2.2}	167
qsort()	Referenced C Language Routines {8.1}	151

rand()	Referenced C Language Routines (8.1)	151
read()	Read from a File {6.4.1}	116
readdir()	Directory Operations {5.1.2}	83
realloc()	Referenced C Language Routines (8.1)	151
remove()	Referenced C Language Routines {8.1}	151
rename()	Rename a File {5.5.3}	99
rewind()	Referenced C Language Routines {8.1}	151
rewinddir()	Directory Operations {5.1.2}	83
rmdir()	Remove a Directory {5.5.2}	98
scanf()	Referenced C Language Routines (8.1)	151
setbuf()	Referenced C Language Routines (8.1)	151
setgid()	Set User and Group IDs {4.2.2}	68
setjmp()	Referenced C Language Routines (8.1)	151
setlocale()	Extensions to <i>setlocale()</i> Function {8.1.2}	154
setpgid()	Set Process Group ID for Job Control {4.3.3}	73
setsid()	Create Session and Set Process Group ID {4.3.2}	72
setuid()	Set User and Group IDs {4.2.2}	68
sigaction()	Examine and Change Signal Action (3.3.4)	58
sigaddset()	Manipulate Signal Sets (3.3.3)	57
sigdelset()	Manipulate Signal Sets {3.3.3}	57
sigemptyset()	Manipulate Signal Sets {3.3.3}	57
sigfillset()	Manipulate Signal Sets {3.3.3}	57
sigismember()	Manipulate Signal Sets {3.3.3}	57
siglongjmp()	Nonlocal Jumps (8.3.1)	162
<signal.h></signal.h>	Signal Concepts (3.3.1)	51
sigpending()	Examine Pending Signals (3.3.6)	62
sigprocmask()	Examine and Change Blocked Signals (3.3.5)	60
sigsetjmp()	Nonlocal Jumps {8.3.1}	162
sigsetops	Manipulate Signal Sets {3.3.3}	57
sigsuspend()	Wait for a Signal {3.3.7}	62
sin()	Referenced C Language Routines {8.1}	151
sinh()	Referenced C Language Routines (8.1)	151
size_t	Primitive System Data Types {2.5}	27
sleep()	Delay Process Execution (3.4.3)	65
sprintf()	Referenced C Language Routines [8.1]	151
sqrt()	Referenced C Language Routines (8.1)	151
srand()	Referenced C Language Routines (8.1)	151
sscanf()	Referenced C Language Routines (8.1)	151
ssize_t	Primitive System Data Types {2.5}	27
stat()	Get File Status (5.6.2)	103
strcat()	Referenced C Language Routines [8.1]	151
strchr()	Referenced C Language Routines {8.1}	151
strcmp()	Referenced C Language Routines [8,1]	151
strcpy()	Referenced C Language Routines {8.1}	151
strcspn()	Referenced C Language Routines {8.1}	151
strftime()	Referenced C Language Routines {8.1}	152
strlen()	Referenced C Language Routines {8.1}	151
strncat()	Referenced C Language Routines (8.1)	151
strncmp()	Referenced C Language Routines {8.1}	151
strncpy()	Referenced C Language Routines {8.1}	151
strpbrk()	Referenced C Language Routines {8.1}	151

Part 1: SYSTEM API [C LANGUAGE]

strrchr()	Referenced C Language Routines (8.1)	151
strspn()	Referenced C Language Routines (8.1)	151
strstr()	Referenced C Language Routines {8.1}	151
strtok()	Referenced C Language Routines (8.1)	151
sysconf()	Get Configurable System Variables (4.8.1)	80
<sys stat.h=""></sys>	File Characteristics: Header and Data Structure {5.6.1}	101
<sys times.h=""></sys>	Get Process Times {4.5.2}	76
<sys types.h=""></sys>	Primitive System Data Types (2.5)	27
<sys utsname.h=""></sys>	Get System Name {4.4.1}	74
<sys wait.h=""></sys>	Wait for Process Termination {3.2.1}	47
tan()	Referenced C Language Routines {8.1}	151
tanh()	Referenced C Language Routines {8.1}	151
tar	Extended tar Format {10.1.1}	169
tcdrain()	Line Control Functions {7.2.2}	145
tcflow()	Line Control Functions {7.2.2}	145
tcflush()	Line Control Functions {7.2.2}	145
tcgetattr()	Get and Set State {7.2.1}	143
tcgetpgrp()	Get Foreground Process Group ID (7.2.3)	147
tcsendbreak()	Line Control Functions (7.2.2)	145
tcsetattr()	Get and Set State (7.2.1)	143
tcsetpgrp()	Set Foreground Process Group ID {7.2.4}	148
termios	General Terminal Interface {7.1}	129
<termios.h></termios.h>	Parameters That Can Be Set {7.1.2}	135
time()	Get System Time {4.5.1}	75
times()	Get Process Times {4.5.2}	76
time_t	Primitive System Data Types {2.5}	27
tmpfile()	Referenced C Language Routines (8.1)	151
tmpnam()	Referenced C Language Routines (8.1)	151
tolower()	Referenced C Language Routines (8.1)	151
toupper()	Referenced C Language Routines (8.1)	151
ttyname()	Determine Terminal Device Name {4.7.2}	79
tzset()	Set Time Zone {8.3.2}	162
uid_t	Primitive System Data Types {2.5}	27
umask()	Set File Creation Mask {5.3.3}	91
uname()	Get System Name {4.4.1}	74
ungetc()	Referenced C Language Routines (8.1)	151
<unistd.h></unistd.h>	Symbolic Constants {2.9}	37
unlink()	Remove Directory Entries (5.5.1)	96
utime()	Set File Access and Modification Times (5.6.6)	108
wait	Wait for Process Termination {3.2.1}	47
waitpid()	Wait for Process Termination (3.2.1)	47
write()	Write to a File {6.4.2}	118



Alphabetic Topical Index

A

Abbreviations ... 20, 208 abort() ... 46, 158, 160-161, 291, 307 definition of ... 151 rationale ... 291 abs() ... 307 definition of ... 151 absolute pathname definition of ... 11 access control lists ... 209 access mode definition of ... 11 access() ... 37, 104-105, 208, 261, 311 definition of ... 104 rationale ... 261 symbolic constants ... 37 acos() ... 305 definition of ... 151 actime ... 109, 312 address space definition of ... 11 AEP (see Applications Environment Profile) alarm schedule ... 63 alarm() ... 42, 46, 63-66, 241, 244-246, 272, 311 definition of ... 63 rationale ... 244 ANSI ... 300 Application Conformance ... 4, 192 Application Specific Environment (ASE) definition of ... 313 Application Specific Environment Description (ASED) definition of ... 314 Applications Environment Profile (AEP) [profile] definition of ... 313 appropriate privileges ... 17, 21, 26, 39, 56, 69-70, 92-93, 97, 105-111, 169, 172, 176-177, 199, 208, 246-247, 261 definition of ... 11 rationale ... 199 Archive/Interchange File Format ... 169, 294 (ARG_MAX) ... 24, 29, 44-45, 80, 230, 304 definition of ... 36

array of char definition of ... 29 ASCII ... 207 asctime() ... 310 definition of ... 152 ASED (see Application Specific Environment Description) asin() ... 305 definition of ... 151 asm builtin atoi() ... 195 assert() ... 301 definition of ... 151 <assert.h> ... 301 asynchronous I/O ... 289 atan2() ... 305 definition of ... 151 atan() ... 305 definition of ... 151 atexit() ... 189, 232, 307 atof() ... 307 definition of ... 151 atoi() ... 194-195, 307 definition of ... 151 atol() ... 307 definition of ... 151

Β

background process group definition of ... 12 background process definition of ... 12 background ... 12, 15, 51, 118, 121, 130-131, 139, 143, 147, 202-204, 206, 248, 277-278, 281-282 Base by POSIX.1, Additions by the C Standard ... 188 Base by the C Standard Additions by POSIX.1 ... 188 Baud Rate Functions ... 141, 281 Baud Rate Values ... 141, 281 Bibliography ... 179 binary stream ... 152 block special file definition of ... 12

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

blocked signals ... 60 brk() ... 188, 292 rationale for omission ... 292 BRKINT ... 136 BSD ... 199, 202, 205-207, 209, 213-214, 216, 222, 231-234, 236-239, 241-244, 247-250, 253, 258-260, 262-264, 266, 270, 272-273, 276-278, 282-284, 287, 291, 294, 296 signals ... 234 bsearch() ... 188, 307 definition of ... 151 By Neither POSIX.1 Nor the C Standard ... 188 byte definition of ... 29 Byte-Oriented cpio Archive Entry ... 174

С

C Language Definitions ... 29, 217 C Language Input/Output Functions ... 155, 287 C Language Limits ... 34, 223 C language FILE functions ... 158 input/output functions ... 155 language-dependent services ... 5 related specifications ... 6 C locale ... 154-155, 286 C Shell ... 202-204, 216 C Standard Language-Dependent Support ... 32 C Standard Language-Dependent System Support ... 5, 194 C Standard ... 2, 5-7, 12, 20, 24, 26-27, 29-30, 32-35, 46, 51, 54-55, 60-61, 103, 134, 151-152, 154-157, 159, 161-162, 186-190, 192-195, 197, 199-200, 212-214, 216-218, 220-223, 228-229, 232-235, 239, 243, 245-246, 250, 259, 266, 270, 272, 275, 285-289, 291, 301, 305, 315 abbreviation ... 20 definition of ... 20 symbols ... 29 calloc() ... 307 definition of ... 151 Canonical Mode Input Processing ... 132, 278 case folding ... 209-210 c_cc ... 132, 140-141, 310 c_cflag ... 135, 138, 310 cc_t ... 310 definition of ... 135

ceil() ... 305 definition of ... 151 cfgetispeed() ... 33, 141-142, 310 definition of ... 141 rationale ... 281 cfgetospeed() ... 33, 141-142, 310 definition of ... 141 rationale ... 281 cfsetispeed() ... 33, 141-142, 144, 281, 310 definition of ... 141 rationale ... 281 cfsetospeed() ... 33, 141-142, 310 definition of ... 141 rationale ... 281 Change Current Working Directory ... 86, 256Change File Modes ... 106, 261 Change Owner and Group of a File ... 107 Change Owner and Group of File ... 262 Change Process's Root Directory ... 257 Character Encoding and Display ... 319 Character Encoding ... 318 character set portable filename ... 207 character special file definition of ... 12 character definition of ... 12 {CHAR_BIT} ... 34, 304 {CHAR_MAX} ... 34, 304 {CHAR_MIN} ... 34, 223, 304 CHARSET variable ... 318-319 chdir() ... 86-87, 176, 256, 261, 311 definition of ... 86 rationale ... 256 Check File Accessibility ... 104 child process definition of ... 12 {CHILD_MAX} ... 80, 215, 224, 304 definition of ... 36 chmod() ... 21, 33, 46, 92, 95-96, 103, 105-106, 108, 177, 261, 308 definition of ... 106 rationale ... 261 chown() ... 39, 103, 107-108, 253, 262-263, 311 definition of ... 107 rationale ... 262 chroot() ... 207, 257 rationale for omission ... 257

Part 1: SYSTEM API [C LANGUAGE]

c iflag ... 131, 136, 310 clearenv() ... 250 clearerr() ... 307 definition of ... 151 c_lflag ... 131, 139, 310 {CLK_TCK} ... 81, 252 definition of ... 81 CLOCAL ... 129, 135, 138 clock tick ... 12, 29, 76-77, 80, 199, 217, 244, 249-251 definition of ... 12 rationale ... 199 clock() ... 199, 310 CLOCKS_PER_SEC ... 310 clock_t ... 12, 29, 76-77, 249-250, 308, 310 definition of ... 29 Close a File ... 115, 266 close() ... 45, 50, 91, 97, 115-116, 127, 158, 160, 279, 300, 311 definition of ... 115 rationale ... 266 closedir() ... 33, 83-85, 254-255, 302 definition of ... 83 rationale ... 254 Closing a Terminal Device File ... 135, 279 cmd Values for fcntl() ... 122 c_oflag ... 133, 137, 310 coherent definition of ... 314 command interpreter portable ... 231 Common-Usage C Language-Dependent System Support ... 6, 194 Common-Usage-Dependent Support ... 32 Compile-Time Symbolic Constants for Portability Specifications ... 38, 225 Compile-Time Symbolic Constants ... 38 complete definition of ... 314 compliance ... 189 comprehensive definition of ... 314 Configurable Pathname Variables ... 110-111,262 Configurable System Variables ... 80, 251 configurable variables ... 80 conformance document definition of ... 10 rationale ... 197 conformance ... 2-7, 9-11, 23, 33, 59, 109, 129, 151, 187-190, 192, 194-197, 199, 211, 213, 218, 222, 283, 313-314 application ... 4

implementation ... 2 conforming application ... 3-5, 11, 188, 190-192, 216, 235-236, 246, 264-265, 267 strictly ... 2-4, 10, 28, 34, 38, 144, 190, 192, 196, 228, 239, 244, 264 Conforming Implementation Options ... 4, 191 Conforming POSIX.1 Application Using Extensions ... 5, 192 Conforming POSIX.1 Application ... 4, 192 consistent definition of ... 314 constants compile-time ... 38 execution-time ... 38 symbolic ... 37 contiguous file ... 294, 297-298 Control Modes ... 138, 280 Control Operations on Files ... 121, 270 controlling process definition of ... 12 Controlling Terminal ... 130, 276 controlling terminal definition of ... 12 rationale ... 199 Conventions ... 9, 196 cooked mode ... 273 core ... 209, 230-232 cos() ... 305 definition of ... 151 cosh() ... 305 definition of ... 151 covert channel ... 211, 242 cpio ... 173-174, 176, 294-299 Archive Entry ... 174 File Data ... 176, 298 File Name ... 175, 298 format ... 173 Header ... 174, 298 Special Entries ... 176, 298 Values ... 176, 298 CREAD ... 138 creat() ... 21, 33, 91-92, 103-104, 115-116, 118, 121, 128, 155, 169, 193, 257-258, 287, 303 definition of ... 91 rationale ... 258 Create a New File or Rewrite an Existing One ... 91, 258 Create an Inter-Process Channel ... 113, 265 Create Session and Set Process Group ID ... 72, 248

Cross-References ... 42, 46, 49-50, 57-58, 60, 62-68, 70-74, 77-79, 85-87, 91-93, 95-97, 99, 101, 103-105, 107-108, 110, 114-116, 118, 121, 127-128, 142, 145, 147-148, 156-157, 162, 166-167, 173, 177 CR ... 134, 137 CRT ... 134 CSIZE ... 138 CSTOPB ... 138 ctermid() ... 33, 78-79, 188, 251, 311 definition of ... 78 rationale ... 251 ctime() ... 152-153, 162, 310 definition of ... 152 <ctype.h> ... 301 currency symbols ... 295 current working directory change ... 86 definition of ... 12, 20 cuserid() ... 247, 311 rationale ... 247

D

Data Definitions for File Control Operations ... 121, 270 Data Interchange Format ... 169, 294 Database Access ... 166, 293 Database Standards ... 181 database group ... 165 user ... 165 **Definitions and General Requirements** ... 196 Definitions ... 10, 197, 313 Delay Process Execution ... 65, 245 Denmark example profile ... 318 Determine Terminal Device Name ... 79, 251 /dev ... 297 Device- and Class-Specific Functions ... 129, 273 device number ... 199 device definition of ... 13 logical ... 206 dev_t ... 101, 214, 309 definition of ... 27 /dev/tty ... 206,251 difftime() ... 310 d ino ... 253

direct ... 253 Directories ... 83, 253 directory entry definition of ... 13 format ... 83 rationale ... 199 remove ... 96 Directory Operations ... 83, 254 directory ... 199 change current working ... 86 current working ... 12 definition of ... 13 empty ... 13 make ... 94 parent ... 16 remove ... 98 root ... 19, 207 working ... 20 working pathname ... 87 dirent ... 83, 85, 253-254, 302 <dirent.h> ... 83-85, 253, 302 definition of ... 83 DIR ... 83-85, 255, 302 definition of ... 84 div() ... 307 div_t ... 307 d_name ... 83, 254, 302 document conformance ... 10 Documentation ... 3, 190 documentation system ... 11 dot ... 13-14, 23, 26, 84, 98, 100, 192, 199, 216, 254, 260 definition of ... 13 dot-dot ... 13-14, 16, 23, 26, 84, 98, 100, 192, 199, 207, 211, 254, 260 definition of ... 13 $dup2() \dots 114-115, 265, 311$ definition of ... 114 rationale ... 265 $dup() \dots 91, 104, 114-116, 118, 121, 128, 158,$ 265, 287, 311 definition of ... 114 rationale ... 265 Duplicate an Open File Descriptor ... 114, 265

Ε

[E2BIG] ... 24, 45, 230, 302 definition of ... 24

Part 1: SYSTEM API [C LANGUAGE]

[EACCES] ... 24, 45, 74, 85-87, 90, 93-94, 96-98, 100, 104-106, 108-109, 112, 249, 256-257, 302 definition of ... 24 eaccess() ... 261 [EAGAIN] ... 24, 42, 117-118, 120, 131, 227, 264, 267-270, 302 definition of ... 24 [EBADF] ... 24, 85, 104, 112, 114-116, 118, 120, 126, 128, 145-149, 265, 267, 270, 302 definition of ... 24 [EBUSY] ... 24, 97-98, 100, 213, 259-260, 302 definition of ... 24 [ECHILD] ... 24, 49, 302 definition of ... 24 ЕСНОЕ ... 139 ЕСНОК ... 139 ECHONL ... 139 ЕСНО ... 139 ed ... 260 [EDEADLK] ... 24, 125, 127, 272, 302 definition of ... 24 [EDOM] ... 24, 213, 302 definition of ... 24 [EEXIST] ... 24, 90, 93-94, 96, 98, 100, 259, 302 definition of ... 24 [EFAULT] ... 24, 212, 230, 302 definition of ... 24 [EFBIG] ... 24, 121, 302 definition of ... 24 effective group ID ... 14-15, 19-20, 39, 44, 68-69, 88, 94-95, 102, 105-107, 229, 246-247, 253, 262 definition of ... 13 effective user ID ... 14, 19-20, 44, 56, 68-69, 88, 94-95, 102, 105-110, 208, 229, 241, 246, 261 definition of ... 13 Eighth Edition UNIX ... 198, 266, 269 [EINTR] ... 24, 49, 63, 65, 90, 115-119, 121, 124, 126, 145-146, 213, 227, 230, 234, 240-241, 266, 302 definition of ... 24 [EINVAL] ... 25, 49, 57-58, 60-61, 69-70, 74, 81, 87, 100, 105, 108, 112, 126, 128, 144-145, 147, 149, 243, 263, 272, 302 definition of ... 25 [EIO] ... 25, 118, 121, 130-131, 135, 302 definition of ... 25 [EISDIR] ... 25, 90, 100, 302 definition of ... 25

[EMFILE] ... 25, 85, 90, 114-115, 126, 302 definition of ... 25 [EMLINK] ... 25, 93-94, 100, 302 definition of ... 25 empty directory definition of ... 13 empty string definition of ... 13 [ENAMETOOLONG] ... 25, 45, 85-86, 90, 93, 95-98, 100, 104-106, 108-109, 112, 230, 258, 260, 302 definition of ... 25 endgrent() ... 293 rationale for omission ... 293 endpwent() ... 293 [ENFILE] ... 25, 85, 90, 114, 302 definition of ... 25 [ENODEV] ... 25, 302 definition of ... 25 [ENOENT] ... 25, 46, 85-86, 90, 93, 95-97, 99-100, 104-105, 107-108, 110, 112, 302 definition of ... 25 [ENOEXEC] ... 25, 46, 228, 230, 302 definition of ... 25 [ENOLCK] ... 25, 126, 271, 302 definition of ... 25 [ENOMEM] ... 25, 42, 46, 213, 227, 291, 302 definition of ... 25 [ENOSPC] ... 26, 90, 93, 95-96, 100, 121, 302 definition of ... 26 [ENOSYS] ... 26, 74, 148-149, 197, 302 definition of ... 26 [ENOTDIR] ... 26, 46, 85-86, 90, 93, 95-97, 99, 101, 104-106, 108, 110, 112, 302 definition of ... 26 [ENOTEMPTY] ... 26, 98, 100, 259-260, 302 definition of ... 26 [ENOTTY] ... 26, 145-149, 213, 274-275, 302 definition of ... 26 environ definition of ... 42 Environment Access ... 77, 250 Environment Description ... 27, 216 [ENXIO] ... 26, 90, 302 definition of ... 26 EOF ... 132, 134, 139 EOL ... 132, 134, 139 [EPERM] ... 26, 57, 69-70, 73-74, 93, 97, 107-108, 110, 149, 242, 302 definition of ... 26 [EPIPE] ... 26, 121, 213, 302 definition of ... 26

Epoch ... 19, 76, 103, 109, 152, 200, 208, 284 definition of ... 13 [ERANGE] ... 26, 87, 213, 257, 302 definition of ... 26 ERASE ... 132, 134, 139 [EROFS] ... 26, 90, 93, 95-97, 99, 101, 105, 107-108, 110, 213, 302 definition of ... 26 errno ... 211, 302 definition of ... 23 <errno.h> ... 23, 217, 259, 302 definition of ... 23 [ERRNO] ... 9 Error Numbers ... 23, 211 Error Reporting ... 161, 290 [ESPIPE] ... 26, 128, 302 definition of ... 26 [ESRCH] ... 26, 57, 74, 241-242, 302 definition of ... 26 /etc/passwd ... 209 [ETXTBSY] ... 230 [EWOULDBLOCK] ... 264 Examine and Change Blocked Signals ... 60, 244Examine and Change Signal Action ... 58, 243Examine Pending Signals ... 62, 244 example profile ... 318 [EXDEV] ... 26, 93, 101, 260, 302 definition of ... 26 exec ... 34-35, 42-43, 45-46, 59, 64, 67, 70, 73-74, 77, 84, 91, 96, 102-103, 115-116, 123, 127, 158-159, 226, 228-231, 246, 248-249, 254, 261, 266, 271, 287, 289 definition of ... 42 of shell scripts ... 228 rationale ... 228 execl() ... 42-44, 46, 228, 311 definition of ... 42 execle() ... 42-44, 46, 228, 311 definition of ... 42 execlp() ... 42-44, 228, 311 definition of ... 42 Execute a File ... 42, 228 Execution-Time Symbolic Constants for Portability Specifications ... 38, 225 Execution-Time Symbolic Constants ... 39 execv() ... 42, 44, 46, 228, 311 definition of ... 42 execve() ... 42, 44, 46, 228, 311 definition of ... 42

execvp() ... 42, 44, 228, 311 definition of ... 42 exit() ... 46, 48, 152, 158, 160-161, 189, 232-233, 240, 266, 288-289, 291, 307 definition of ... 151 rationale ... 291 EXIT FAILURE ... 152 EXIT_SUCCESS ... 152, 233 _exit() ... 46, 48-50, 73, 158-159, 189, 232-233, 240, 288-289, 311 definition of ... 49 rationale ... 232 exp() ... 305 definition of ... 151 Extended cpic Format ... 173, 297 extended security controls definition of ... 21, 208 Extended tar Format ... 169, 296 Extensions to setlocale() Function ... 154, 285 Extensions to Time Functions ... 152, 283

\mathbf{F}

fabs() ... 305 definition of ... 151 Fast File System ... 262 fclose() ... 158, 160-161, 189, 289, 291, 307 definition of ... 151 rationale ... 289 fcntl() ... 33, 42, 46, 91, 104, 113-116, 118, 121, 123-128, 131, 158, 176, 214, 265, 268, 270-271, 275, 287, 303 definition of ... 121 rationale ... 270 Return Values ... 126 <fcntl.h> ... 30, 88, 91, 121, 123-124, 127, 302 definition of ... 121 fdopen() ... 33, 156-158, 188, 288, 307 definition of ... 157 rationale ... 287 FD_CLOEXEC ... 44, 88, 113, 122-123, 158, 254, 270, 302 F_DUPFD ... 114, 122-123, 126, 270, 302 feature test macro definition of ... 13 feature test macros suggested ... 222 feof() ... 159, 307 definition of ... 151

ferror() ... 307 definition of ... 151 fflush() ... 159-161, 288-290, 307 definition of ... 151 rationale ... 290 fgetc() ... 160, 290, 307 definition of ... 151 rationale ... 290 F GETFD ... 122-123, 125, 270, 302 F_GETFL ... 122-123, 125, 270, 302 F_GETLK ... 122, 124-126, 302 fgetpos() ... 214, 272, 307 fgets() ... 160, 290, 307 definition of ... 151 rationale ... 290 FIFO special file ... 201 definition of ... 13 FIFO ... 13-14, 17, 26, 88-90, 95, 102, 111, 116-117, 119-121, 128, 170, 172-173, 175-176, 201, 207, 253, 257, 259, 267-269 make ... 95 File Access Modes Used for open() and fcntl() ... 123 file access permissions definition of ... 21, 208 File Accessibility ... 261 File Characteristics: Header and Data Structure ... 101, 260 File Characteristics ... 101, 260 file classes ... 201 File Control ... 121, 270 file description definition of ... 14 File Descriptor Deassignment ... 115, 266 File Descriptor Flags Used for *fcntl()* ... 122 File Descriptor Manipulation ... 114, 265 file descriptor definition of ... 14 duplicate ... 114, 122 map a stream pointer ... 156 open stream ... 157 file group class definition of ... 14 file hierarchy definition of ... 22, 208 file mode definition of ... 14 file offset definition of ... 14 file other class definition of ... 14

file owner class definition of ... 14 file permission bits ... 14, 21, 88, 91, 94-95, 103, 105-106, 261 definition of ... 15 file permissions ... 13-16, 20-21, 24, 27, 88, 90-92, 94-95, 103, 105-106, 169, 171-172, 174-175, 208, 246, 261, 263, 277 definition of ... 209 File Removal ... 96, 259 file serial number ... 101 definition of ... 15 File Status Flags Used for open() and fcntl() ... 122 file system ... 201 definition of ... 15 mounted ... 206 read-only ... 18 root ... 208 root of ... 208 file times update definition of ... 22, 210 file type (see *file*) file ... 201 access ... 104 access permissions ... 21 binary ... 152, 187 block special ... 12 change modes ... 106 change owner and group ... 107 character special ... 12 characteristics header ... 101 close ... 115 contiguous ... 294, 297-298 control ... 121 create or rewrite ... 91 definition of ... 14 execute ... 42 FIFO special ... 13 get status ... 103 hierarchy ... 22 high performance ... 294, 297-298 link to ... 92 locking ... 124, 270 make FIFO special ... 95 mode ... 14 offset ... 14 open ... 88 passwd ... 207 permission bits ... 15 read ... 116 regular ... 19 remove ... 162 rename ... 99 reposition read/write offset ... 127 serial number ... 15 set access and modification times ... 108

set creation mask ... 91 text ... 152, 187 times ... 108 times update ... 22 write ... 118 filename portability definition of ... 22, 209 filename ... 201 definition of ... 14 fileno() ... 33, 156, 158, 188, 206, 307 definition of ... 156 rationale ... 287 Files and Directories ... 83, 253 FILE ... 156-158, 160, 287-288, 307 <float.h> ... 303 flock ... 124, 302 Structure ... 125 *floor()* ... 305 definition of ... 151 fmod() ... 305 definition of ... 151 F_OK ... 37-38, 105, 311 fopen() ... 157-158, 160-161, 201, 288-289, 291, 307 definition of ... 151 rationale ... 289 {FOPEN_MAX} ... 35 foreground process group ID definition of ... 15 foreground process group definition of ... 15 foreground process definition of ... 15 foreground ... 12, 15, 129-131, 133-134, 136, 140, 143, 147-148, 202-206, 248, 275-278, 282-283 fork() ... 17-18, 41-42, 46, 49, 64, 67, 73, 77, 84, 116, 125, 130, 156, 158-159, 203, 214, 226-227, 229-231, 248, 254, 270-271, 276, 287-289, 311 definition of ... 41 rationale ... 226 Format of Directory Entries ... 83, 253 fpathconf() ... 39, 110-112, 263, 311 definition of ... 110 rationale ... 263 fpos_t ... 306 fprintf() ... 161, 272, 307 definition of ... 151 fputc() ... 160, 290, 307 definition of ... 151 rationale ... 290

fputs() ... 160, 290, 307 definition of ... 151 rationale ... 290 F_RDLCK ... 122, 124-126, 302 fread() ... 160, 290, 307 definition of ... 151 rationale ... 290 free() ... 188, 240, 256, 307 definition of ... 151 freopen() ... 158, 160, 290, 307 definition of ... 151 rationale ... 290 frexp() ... 305 definition of ... 151 fscanf() ... 160, 290, 307 definition of ... 151 rationale ... 290 fseek() ... 152, 159, 161, 290, 307 definition of ... 151 rationale ... 290 F_SETFD ... 122-123, 126, 270, 302 F_SETFL ... 122, 124, 126, 270, 302 F_SETLK ... 122, 124-126, 271, 302 F_SETLKW ... 122, 124-127, 271, 302 fsetpos() ... 214, 272, 307 fstat() ... 21-22, 33, 101, 103-104, 261, 308 definition of ... 103 rationale ... 260 ftell() ... 161, 290, 307 definition of ... 151 rationale ... 290 function prototypes ... 32 F_UNLCK ... 122, 124-125, 302 fwrite() ... 160, 272, 290, 307 definition of ... 151 rationale ... 290 F_WRLCK ... 122, 124-126, 302

G

General Concepts ... 21, 208
General File Creation ... 88, 257
General Terminal Interface Control Functions ... 143, 281
General Terminal Interface ... 129, 274
General Terms ... 11, 198
General ... 1
Generate Terminal Pathname ... 78, 251
Get and Set State ... 143, 282
Get Configurable Pathname Variables ... 110, 263

Part 1: SYSTEM API [C LANGUAGE]

Get Configurable System Variables ... 80, 252Get Effective Group ID ... 68, 246 Get Effective User ID ... 68, 246 Get File Status ... 103, 260 Get Foreground Process Group ID ... 147, 282Get Password From User ... 252 Get Process and Parent Process IDs ... 67, 246Get Process Group ID ... 72, 247 Get Process Times ... 76, 250 Get Real Group ID ... 68, 246 Get Real User ID ... 68, 246 Get Supplementary Group IDs ... 70, 247 Get System Name 74 Get System Time ... 75, 249 Get User Name ... 71, 247 Get Working Directory Pathname ... 87 getc() ... 160, 290, 307 definition of ... 151 rationale ... 290 getchar() ... 160, 290, 307 definition of ... 151 rationale ... 290 getcwd() ... 86-87, 256, 311 definition of ... 87 rationale ... 256 getegid() ... 68, 247, 311 definition of ... 68 rationale ... 246 getenv() ... 77-78, 189, 229, 307 definition of ... 77 rationale ... 250 geteuid() ... 68, 311 definition of ... 68 rationale ... 246 getgid() ... 68, 311 definition of ... 68 rationale ... 246 getgrent() ... 293 rationale for omission ... 293 getgrgid() ... 33, 166, 293, 303 definition of ... 166 rationale ... 293 getgrnam() ... 33, 166, 211, 293, 303 definition of ... 166 rationale ... 293 getgroups() ... 70, 247, 311 definition of ... 70 rationale ... 247

gethostid() ... 249 gethostname() ... 249 getlogin() ... 71, 166-167, 247, 311 definition of ... 71 rationale ... 247 getpass() ... 252 rationale for omission ... 252 getpgrp() ... 72, 74, 204, 247, 311 definition of ... 72 rationale ... 247 getpid() ... 57, 67, 73, 241, 311 definition of ... 67 rationale ... 246 getppid() ... 67, 311 definition of ... 67 rationale ... 246 getpwent() ... 293 rationale for omission ... 293 getpwnam() ... 33, 71, 167, 211, 247, 293, 305 definition of ... 167 rationale ... 293 getpwuid() ... 33, 71, 167, 293, 305 definition of ... 167 rationale ... 293 gets() ... 160, 290, 307 definition of ... 151 rationale ... 290 gettimeofday() ... 249-250 getty ... 276 getuid() ... 68, 70, 215, 241, 311 definition of ... 68 rationale ... 246 getwd() ... 256 gid_t ... 15, 20, 68, 70, 101, 107, 166-167, 214-215, 293, 309 definition of ... 27 GKS ... 196 gmtime() ... 152, 200, 310 definition of ... 152 Graphics Standards ... 181 gr_gid ... 166, 303 gr_mem ... 166, 303 gr_name ... 166, 303 Group Database Access ... 166, 293 group file ... 201 group ID definition of ... 15 effective ... 13, 68 get effective ... 68 get real ... 68 real ... 18,68 saved set- ... 19 set ... 68 supplementary ... 20, 70

group change file ... 107 group Structure ... 166 groups multiple (see supplementary group ID) <grp.h> ... 166, 173, 177, 215, 293, 303 definition of ... 166

Η

handle definition of ... 158 harmonized definition of ... 314 Header Contents Samples ... 301 header definition of ... 29 Headers and Function Prototypes ... 32, 223 hierarchy file ... 22 Historical Documentation and Introductory Texts ... 182 historical implementations ... 201 HOME definition of ... 27 variable ... 27 hosted implementation ... 201 HUPCL ... 135, 138

I

ICANON ... 134, 139-141 ICRNL ... 134, 136-137 IEXTEN ... 135, 139-140 IFS variable ... 217 IGNBRK ... 136 IGNCR ... 134, 136-137 IGNPAR ... 136-137 Implementation Conformance ... 2, 190 implementation defined ... 3-4, 11, 14, 16-18, 21, 24, 37-38, 42-43, 49-50, 53, 62, 75, 89, 91, 94-95, 103-104, 106-107, 117, 119, 121, 127, 129-132, 134-141, 145, 152, 159, 162, 165, 169, 171, 175-177, 185, 198-199, 201, 206-209, 211-213, 218, 230, 232, 235-236, 238, 241, 243-244, 256-259, 272, 284-288, 291, 294, 296, 299 definition of ... 10 rationale ... 197

implementation ... 201 historical ... 182, 201 hosted ... 187, 201 native ... 206 specific ... 201 incomplete pathname ... 202 init ... 205, 233, 241 INLCR ... 136-137 ino t ... 101, 309 definition of ... 27 INPCK ... 136-137 Input and Output Primitives ... 113, 264 Input and Output ... 116, 266 Input Modes ... 136, 279 Input Processing and Reading Data ... 131, 278 Interactions of Other FILE-Type C Functions ... 158, 288 interchange cpio ... 173 multiple volumes ... 177 tar ... 169 Interface Characteristics ... 129, 275 International Standardized Profile (ISP) definition of ... 313 inter-process channel ... 113 INTR ... 133-134, 139-140 {INT_MAX} ... 34, 304 {INT_MIN} ... 34, 304 Invariant Value ... 37 Invariant Values ... 37 *iocntl()* ... 274 ioctl() ... 213, 273-275, 282-283 IPC (see inter-process channel) TRV abbreviation ... 20 definition of ... 20 isalnum() ... 301 definition of ... 151 isalpha() ... 301 definition of ... 151 isascii() ... 188 isatty() ... 79, 188, 212, 311 definition of ... 79 rationale ... 251 *iscntrl()* ... 301 definition of ... 151 *isdigit*() ... 301 definition of ... 151 isgraph() ... 301 definition of ... 151

Part 1: SYSTEM API [C LANGUAGE]

ISIG ... 133-134, 139-140 islower() ... 301 definition of ... 151 ISO/IEC 10646 ... 182 ISO/IEC 646 ... 2, 20, 169, 171-174, 176, 207, 295 rationale for selection ... 295 ISO/IEC 9899 ... 2, 20 ISO/IEC 9945-2 ... 181, 226, 228, 286, 289, 294, 315, 318-319 ISO/IEC Conforming POSIX.1 Application ... 4 *isprint()* ... 301 definition of ... 151 ISP definition of ... 313 *ispunct()* ... 301 definition of ... 151 isspace() ... 301 definition of ... 151 ISTRIP ... 136-137 *isupper()* ... 301 definition of ... 151 *isxdigit()* ... 301 definition of ... 151 IXOFF ... 134, 136, 140, 146, 224 IXON ... 134, 136-137, 140

J

jmp_buf ... 305
Job Control Signals ... 52
job control ... 4, 15, 38, 47, 50-52, 73, 118,
121, 129-130, 134-135, 139-140, 143, 197,
202-207, 212, 231, 233-234, 236, 239, 242,
247-249, 251, 273, 275, 277
definition of ... 15
implementing applications ... 204
implementing shells ... 202
implementing systems ... 205
set process group ID ... 73

Κ

kernel ... 206
kill ... 235
kill() ... 25, 33, 42, 51, 54, 56-57, 60-61, 65,
 67, 73, 187, 189, 214, 235-237, 239-242, 283,
 306
 definition of ... 56
 rationale ... 241

KILL ... 132, 134, 139 KornShell ... 216

L

labels ANSI ... 300 labs() ... 307 LANG definition of ... 28 variable ... 28, 155 Language Standards ... 181 Language-Dependent Services for the C Programming Language ... 5, 192 Language-Specific Services for the C Programming Language ... 151, 283 lconv ... 304 LC ALL definition of ... 28 variable ... 28, 154-155, 286-287, 304 LC * variable ... 216 LC COLLATE definition of ... 28 variable ... 28, 154, 286, 304 LC_CTYPE definition of ... 28 variable ... 28, 154, 286, 304 LC_MESSAGES variable ... 320 **LC_MONETARY** definition of ... 28 variable ... 28, 154, 286, 304 LC_NUMERIC definition of ... 28 variable ... 28, 154, 286, 304 LC_TIME definition of ... 28 variable ... 28, 154, 286, 304, 320 L_ctermid ... 78, 251 L_cuserid ... 30, 306 *ldexp()* ... 305 definition of ... 151 ldiv() ... 307 *ldiv_t* ... 307 library routine ... 206 {LIMIT} ... 9, 34 limits ... 34 Clanguage ... 34 invariant values ... 37 minimum values ... 34 pathname values ... 36 run-time increasable values ... 34 run-time invariant values ... 35

imits.h> ... 3, 34-37, 80, 110-111, 188, 190, 223-225, 236, 251, 263, 304, 318 definition of ... 34 Line Control Functions ... 145, 282 link (see directory entry) link count definition of ... 16 Link to a File ... 92, 258 link definition of ... 15 symbolic ... 294, 297-298 link() ... 24, 92-93, 97, 101, 103, 177, 199, 258-259, 311 definition of ... 92 rationale ... 258 {LINK_MAX} ... 25, 93-94, 100, 111, 304 definition of ... 36 l_len ... 124-125, 303 Local Modes ... 139, 280 Locale Definitions ... 319 locale C ... 154-155, 286 POSIX ... 154-155, 286 localeconv() ... 304 <locale.h> ... 154,304 localtime() ... 152-153, 162, 200, 284, 310 definition of ... 152 lockf() ... 271 locking ... 270 advisory ... 124, 272 exclusive ... 271 mandatory ... 272 log10() ... 305 definition of ... 151 log() ... 305 definition of ... 151 logical device ... 206 login name definition of ... 16 login shell ... 228 login ... 228,246 login definition of ... 16 LOGNAME definition of ... 28 variable ... 28, 216 longjmp() ... 66, 162, 189, 213, 240, 245, 291-292, 305 definition of ... 151 rationale ... 291 _longjmp() ... 291

{LONG_MAX} ... 34, 304 {LONG_MIN} ... 34, 304 *l_pid* ... 125, 303 lread() ... 214, 266 rationale for omission ... 266 lseek() ... 26, 37, 91, 118, 121, 127-128, 152, 158-161, 193, 214, 272-273, 290, 311 definition of ... 127 rationale ... 272 symbolic constants ... 37 *l_start* ... 124-125, 303 *l_sysid* ... 271 *l_type* ... 126, 303 Values for Record Locking With fcntl() ... 122 *l_whence* ... 124-125, 303 *lwrite()* ... 266 rationale for omission ... 266

Μ

machine ... 74, 309 macro feature test ... 13 MAIL variable ... 217 main() ... 43, 46, 48, 53, 156, 228, 232 definition of ... 43 Make a Directory ... 94, 258 Make a FIFO Special File ... 95, 259 make ... 233 malloc() ... 188, 240, 256, 291-292, 307 definition of ... 151 Manipulate Signal Sets ... 57, 242 Map a Stream Pointer to a File Descriptor ... 156, 287 Mask for Use With File Access Modes ... 123 mask file creation ... 91 <math.h> ... 305 {MAX_CANON} ... 111, 132, 304 definition of ... 36 {MAX_CHAR} ... 224 {MAX_INPUT} ... 111, 131-132, 224, 304 definition of ... 36 mav definition of ... 10 rationale ... 197 mblen() ... 307 {MB_CUR_MAX} ... 315
{MB_LEN_MAX} ... 34, 304, 315 mbstowcs() ... 307 mbtowc() ... 307 memchr() ... 308 memcmp() ... 308 memcpy() ... 308 memmove() ... 308 memory management omitted functions ... 292 memset() ... 308 Minimum Values ... 34-35, 223 MIN ... 132-133, 139 mkdir() ... 21, 33, 91-92, 94, 99, 103, 107, 177, 253, 258-259, 308 definition of ... 94 rationale ... 258 $mkfifo() \ldots 21, 33, 91-92, 95-96, 103, 107,$ 201, 253, 259, 308 definition of ... 95 rationale ... 259 mknod() ... 201, 259 mktime() ... 152-153, 162, 200, 310 definition of ... 152 mode bit ... 44, 172, 257 mode definition of ... 16 Modem Disconnect ... 135, 279 mode_t ... 88, 91, 94-95, 101-102, 106, 214, 258, 260, 309 definition of ... 27 modf() ... 305 definition of ... 151 modtime ... 109, 312 mount point ... 206 mount() ... 206 mounted file system ... 206 multibyte character ... 12, 135, 201, 278, 280, 295 multiple groups (see supplementary group ID) Multiple Volumes ... 177, 299

Ν

name
 login ... 16
 system ... 74
 user ... 20
{NAME_MAX} ... 14, 22-23, 25, 39, 46, 83, 85 86, 90, 93, 95-98, 100, 104-106, 108-109,
 111-112, 225, 262-263, 304
 definition of ... 36

namespace pollution ... 218 <National Body> Conforming POSIX.1 Application ... 5 native implementation ... 206 NDEBUG ... 301 Networking Standards ... 179 new() ... 292 {NGROUPS_MAX} ... 4, 20, 70, 80, 247, 304 nlink_t ... 101, 214, 309 definition of ... 27 NL ... 134, 137, 139 nodename ... 74, 309 NOFLSH ... 139-140 nohup ... 229 Noncanonical Mode Input Processing ... 132, 278Nonlocal Jumps ... 162, 291 Normative References ... 2, 189 null character definition of ... 29 null string definition of ... 13, 16 NULL ... 304, 306-307, 310-311 definition of ... 29 Numerical Limits ... 34, 223

0

O_ACCMODE ... 123, 302 O_APPEND ... 88, 119, 122, 161, 288, 302 obsolescent ... 81, 143-144, 197, 252, 265, 281 definition of ... 10 rationale ... 197 O_CREAT ... 88-90, 122, 302 O_EXCL ... 88-90, 122, 302 off_t ... 101, 125, 127, 214, 309 definition of ... 27 oflag Values for open() ... 122 Omitted Memory Management ... 292 O_NDELAY ... 258, 264, 269-270 O_NOCTTY ... 89, 122, 130, 257, 302 O_NONBLOCK ... 89-90, 113, 117-120, 122, 129, 131, 133, 138, 258, 264, 267-270, 302 Open a File ... 88, 257 Open a Stream on a File Descriptor ... 157, 287 open file description ... 206 definition of ... 16

open file definition of ... 16 Open System Environment (OSE) definition of ... 314 open() ... 21, 33, 45, 88-92, 97, 104, 114-116, 118, 121, 127-131, 137-140, 156-157, 160, 176, 193, 201, 253, 257-258, 260-261, 265, 277, 287, 291, 303 definition of ... 88 rationale ... 257 opendir() ... 33, 83-85, 176, 254-255, 302 definition of ... 83 rationale ... 254 Opening a Terminal Device File ... 129, 275 {OPEN_MAX} ... 25, 80, 84, 114-115, 126, 304 definition of ... 36 OPOST ... 137 O_RDONLY ... 88-89, 123, 257-258, 302 O_RDWR ... 88-90, 123, 257-258, 302 orphaned process group ... 206, 233, 239 definition of ... 16 OSE (see Open System Environment) Other C Language Functions ... 162, 291 Other C Language-Related Specifications ... 6, 194 Other Language-Related Specifications ... 7, 195 Other Standards ... 181 O_TRUNC ... 39, 89-91, 111, 122, 257, 302 Output Modes ... 137, 280 owner change file ... 107 O_WRONLY ... 88-90, 123, 258, 302

P

Parameters That Can Be Set ... 135, 279 PARENB ... 138 parent directory ... 207 definition of ... 16 parent process ID definition of ... 16 parent process definition of ... 16 PARMRK ... 136-137 PARODD ... 138 passwd file ... 207 passwd ... 167 Structure ... 167 <passwd.h> ... 215 path prefix definition of ... 17 pathconf() ... 37, 39, 110-112, 202, 226, 262-263, 311 definition of ... 110 rationale ... 263 pathname component definition of ... 17 pathname resolution definition of ... 22, 211 Pathname Variable Values ... 36, 224 pathname absolute ... 11 definition of ... 17 get configurable values ... 110 get working directory ... 87 incomplete ... 202 relative ... 19 PATH definition of ... 28 variable ... 28, 43, 46, 216, 230 {PATH_MAX} ... 17, 25, 46, 85-86, 90, 93, 95-98, 100, 104-106, 108-109, 111-112, 225, 252, 256-257, 263, 297-298, 304 definition of ... 36 pause() ... 49, 63-66, 238, 241, 245, 311 definition of ... 64 rationale ... 245 pclose() ... 231 permission file ... 15 permissions file access ... 21 perror() ... 161, 212, 290, 307 definition of ... 151 rationale ... 290 {PID_MAX} ... 214 $pid_t \dots 18, 41, 47, 56, 67, 72-73, 124, 147$ 148, 214-215, 223, 309 definition of ... 27 pipe ... 14, 26, 34, 36, 51, 55, 102, 111, 113, 116-117, 119-121, 128, 202-203, 207, 227, 238, 257, 265, 267-269, 289, 311 definition of ... 17 pipe() ... 17, 96, 103-104, 113-116, 118, 121, 265, 287, 311 definition of ... 113 rationale ... 265 {PIPE_BUF} ... 111, 119-120, 267-269, 304 definition of ... 36 {PIPE_MAX} ... 269 Pipes ... 113, 265

popen() ... 231, 289 portability constants compile-time ... 38 execution-time ... 38 portable filename character set ... 207 definition of ... 17 portable filenames ... 22 <posix.h> ... 225 POSIX locale ... 154-155, 286 **POSIX Open System Environment** definition of ... 314 POSIX.1 and the C Standard ... 186 POSIX.1 Symbols ... 30, 217 POSIX.1 abbreviation ... 21 definition of ... 21 {_POSIX_ARG_MAX} ... 304 definition of ... 34 (_POSIX_CHILD_MAX} ... 304 definition of ... 34 (_POSIX_CHOWN_RESTRICTED} ... 4, 107-108, 111, 263, 311 definition of ... 39 (_POSIX_JOB_CONTROL) ... 4, 15, 73, 80, 129, 147-148, 311 definition of ... 38 {_POSIX_LINK_MAX} ... 304 definition of ... 34 [POSIX_MAX_CANON] ... 304 definition of ... 34 (_POSIX_MAX_INPUT) ... 304 definition of ... 34 (_POSIX_NAME_MAX) ... 304 definition of ... 34 (_POSIX_NGROUPS_MAX) ... 304 definition of ... 34 (_POSIX_NO_TRUNC) ... 22-23, 25, 46, 85-86, 90, 93, 95-98, 100, 104-106, 108-109, 111-112, 225, 311 definition of ... 39 [_POSIX_OPEN_MAX} ... 304 definition of ... 34 {_POSIX_PATH_MAX} ... 304 definition of ... 34 [_POSIX_PIPE_BUF] ... 304 definition of ... 34 (_POSIX_SAVED_IDS) ... 44, 56, 69-70, 80, 229, 311 definition of ... 38 _POSIX_SOURCE ... 32, 218, 221 definition of ... 31

(_POSIX_SSIZE_MAX) ... 304 definition of ... 34 (POSIX_STREAM_MAX) ... 304 definition of ... 34 { POSIX_TZNAME_MAX} ... 304, 318 definition of ... 34 (_POSIX_VDISABLE) ... 111, 135, 141, 311 definition of ... 39 {_POSIX_VERSION} ... 80, 249, 311 definition of ... 38 pow() ... 305 definition of ... 151 Primitive System Data Types ... 27, 213 printf() ... 160-161, 193, 290, 307 definition of ... 151 rationale ... 290 privileges (see appropriate privileges) Process Creation and Execution ... 41, 226 Process Creation ... 41, 226 Process Environment ... 67, 246 process group ID ... 15, 18, 27, 41, 44, 47, 56, 72-74, 129, 147-149, 203, 247-248, 275-276 definition of ... 18 foreground ... 15 get ... 72 get foreground ... 147 set ... 72 set for job control ... 73 set foreground ... 148 process group leader definition of ... 18 process group lifetime definition of ... 18 process group background ... 12 definition of ... 17 foreground ... 15 leader ... 18 lifetime ... 18 orphaned ... 16, 206, 233, 239 terminal ... 129 Process Groups ... 72, 129, 247, 275 process groups concepts in job control ... 203 Process Identification ... 67, 246 process ID 1 ... 233 definition of ... 18 get ... 67 get parent ... 67 parent ... 16 rationale ... 214

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

process lifetime ... 242 definition of ... 18 Process Primitives ... 41, 226 Process Termination ... 46, 230 process background ... 12 child ... 12 concurrent execution ... 227 controlling ... 12 create inter-process channel ... 113 creation ... 41 definition of ... 17 delay execution ... 65 foreground ... 15 parent ... 16 send signal to ... 56 single-threaded ... 227 suspend execution ... 64 system ... 20 terminate ... 49 times ... 76 wait for termination ... 47 profile example ... 318 **Profiles** ... 313 prototypes function ... 32 PS1 variable ... 217 PS2 variable ... 217 ptrace() ... 212, 231-232, 243 *ptrdiff_t* ... 306 putc() ... 160, 290, 307 definition of ... 151 rationale ... 290 putchar() ... 160, 290, 307 definition of ... 151 rationale ... 290 putenv() ... 250 puts() ... 160, 290, 307 definition of ... 151 rationale ... 290 pwd ... 256 <pwd.h> ... 27, 167, 173, 177, 215, 293, 305 definition of ... 167 pw_dir ... 167, 305 pw_gid ... 167, 305 pw_name ... 167, 305 pw_shell ... 167, 305 pw_uid ... 167, 305

INFORMATION TECHNOLOGY-POSIX

Q

```
qsort() ... 188, 307
definition of ... 151
QUIT ... 133-134, 139-140
```

R

raise() ... 54, 61, 189, 306 rand() ... 307 definition of ... 151 raw mode ... 273 Read from a File ... 116, 267 read() ... 91, 103, 114, 116-119, 128, 130-135, 158, 160, 177, 202, 213, 215, 238-240, 253, 264-267, 269, 277-278, 299-300, 311 definition of ... 116 rationale ... 267 readdir() ... 33, 83-85, 254-255, 302 definition of ... 83 rationale ... 254 read-only file system definition of ... 18 readv() ... 269 real group ID ... 15, 44, 68-70, 105 definition of ... 18 real user ID ... 20, 34-35, 44, 68-69, 105, 241, 261 definition of ... 18 realloc() ... 307 definition of ... 151 Referenced C Language Routines ... 151, 283 regular file ... 207 definition of ... 19 Related Activities ... 315 Related Functions by Both ... 189 Related Open Systems Standards ... 179 Related Standards ... 315 Relationship to IEEE Draft Project 1003.0 ... 315 relative pathname definition of ... 19 release ... 74, 309 Remove a Directory ... 98, 259 Remove Directory Entries ... 96, 259 remove() ... 103, 162, 307 definition of ... 151 Rename a File ... 99, 259 rename() ... 93, 97, 99-100, 151, 189, 259, 307 definition of ... 99 rationale ... 259

Reposition Read/Write File Offset ... 127, 272Required Signals ... 52 Requirements ... 2, 190 Reserved Header Symbols ... 31 rewind() ... 161, 290, 307 definition of ... 151 rationale ... 290 rewinddir() ... 33, 83-84, 255, 302 definition of ... 83 rationale ... 254 rmdir() ... 97-98, 101, 213, 259, 311 definition of ... 98 rationale ... 259 R OK ... 37-38, 105, 311 root directory ... 207 definition of ... 19 root file system ... 208 root of a file system ... 208 Run-Time Increasable Values ... 34-35, 224 Run-Time Invariant Values (Possibly Indeterminate) ... 35-36, 224

S

sa_flags ... 59, 306 sa_handler ... 59, 306 sa_mask ... 59, 243, 306 samefile() ... 214 Sample National Profile ... 317 SA_NOCLDSTOP ... 53, 59, 203, 243, 305 saved set-group-ID ... 15, 38, 44, 69-70, 246 definition of ... 19 saved set-user-ID ... 19-20, 38, 44, 56, 69, 246 definition of ... 19 sbrk() ... 188, 292 rationale for omission ... 292 scanf() ... 160, 193, 290, 307 definition of ... 151 rationale ... 290 {SCHAR_MAX} ... 34, 304 {SCHAR_MIN} ... 34, 304 Schedule Alarm ... 63, 244 Scope and Normative References ... 185 Scope ... 1, 185 seconds since the Epoch ... 76, 103, 109, 152, 200, 208, 284 definition of ... 19 security considerations ... 4, 21, 56, 199, 201, 205, 208-209, 216, 230, 242, 248, 252, 260, 262, 277, 282-283, 293, 295-296

security controls additional ... 21 alternate ... 21 extended ... 21 security access control lists ... 209 mandatory ... 209 monolithic privileges ... 199 seek [see lseek()] seekdir() ... 255-256 rationale for omission ... 255 SEEK_CUR ... 30, 38, 124, 127, 303, 311 SEEK_END ... 30, 38, 124, 127, 303, 311 SEEK_SET ... 30, 38, 124-125, 127, 289, 303, 311 select() ... 266 Send a Signal to a Process ... 56, 241 session leader definition of ... 19 session lifetime definition of ... 19 session ... 12, 15-16, 19, 44, 50, 56, 72-74, 130, 148-149, 179, 204, 206-207, 233, 242, 248, 276, 283 create ... 72 definition of ... 19 leader ... 19 lifetime ... 19 Set File Access and Modification Times ... 108, 262 Set File Creation Mask ... 91, 258 Set Foreground Process Group ID ... 148, 283 Set Position of Directory Stream ... 255 Set Process Group ID for Job Control ... 73, 248 Set Time Zone ... 162, 292 Set User and Group IDs ... 68, 246 setbuf() ... 307 definition of ... 151 setcooked() ... 273 setgid() ... 68-71, 311 definition of ... 68 rationale ... 246 setgrent() ... 293 rationale for omission ... 293 set-group-ID ... 44, 107, 229-230, 272 setgroups() ... 247 sethostid() ... 249 sethostname() ... 249 setjmp() ... 162, 189, 245, 291-292, 305 definition of ... 151 rationale ... 291

IEEE Std 1003.1-1990 <setjmp.h> ... 162,305 _setjmp() ... 291 setlocale() ... 154-155, 189, 285-287, 304 definition of ... 154 rationale ... 285 setpgid() ... 18-19, 72-74, 148, 197, 203-204, 233, 248-249, 276, 311 definition of ... 73 rationale ... 248 setpgrp() ... 248 setpwent() ... 293 rationale for omission ... 293 setraw() ... 273 setsid() ... 18, 57, 72-74, 130, 148, 233, 248, 276.311 definition of ... 72 rationale ... 248 setuid() ... 44, 46, 68-69, 311 definition of ... 68 rationale ... 246 set-user-ID ... 44, 107, 229-230, 241, 256 setvbuf() ... 288, 307 sh ... 287 shall definition of ... 10 rationale ... 197 shell scripts execing ... 228 shell ... 1, 202-205, 216-217, 226, 228, 231, 233, 236, 238-239, 242, 247-249, 275, 277, 288 job control ... 202, 236, 239, 242 login ... 205, 228, 247 SHELL variable ... 217 should definition of ... 10 rationale ... 197 {SHRT_MAX} ... 34, 304 {SHRT_MIN} ... 34, 304 SIGABRT ... 46, 51, 235, 291, 305 sigaction ... 58, 306 sigaction() ... 33, 50, 53, 57-60, 62-64, 66, 72-73, 128, 162, 236-237, 241-244, 283, 291, 306 definition of ... 58 rationale ... 243 sigaddset() ... 33, 57-58, 306 definition of ... 57 rationale ... 242 SIGALRM ... 51, 64-66, 245, 305 sig_atomic_t ... 306

ISO/IEC 9945-1: 1990

SIGBUS ... 235-236 SIGCHLD ... 50-51, 53-54, 59, 203, 236-239, 243, 305 SIGCLD ... 235-236, 238, 243 SIGCONT ... 50, 52-53, 55-56, 203, 233-234, 236, 238-239, 242, 305 sigdelset() ... 33, 57-58, 306 definition of ... 57 rationale ... 242 sigemptyset() ... 33, 57, 242-243, 306 definition of ... 57 rationale ... 242 SIGEMT ... 235-236 sigfillset() ... 33, 57, 242-243, 306 definition of ... 57 rationale ... 242 SIGFPE ... 51, 54, 61, 235-236, 238, 305 SIGHUP ... 50-51, 135, 233-234, 239, 305 SIGILL ... 51, 54, 61, 235-236 SIGINT ... 24, 51, 133, 136, 205, 226, 235, 305 SIGIOT ... 235-236 sigismember() ... 33, 57-58, 306 definition of ... 57 rationale ... 242 sigjmp_buf ... 305 SIGKILL ... 51, 53-54, 59, 61, 235, 237-239, 241, 243, 305 siglongjmp() ... 33, 66, 162, 189, 213, 240, 291, 305 definition of ... 162 rationale ... 291 Signal Actions ... 53, 238 Signal Concepts ... 51, 235 Signal Effects on Other Functions ... 55, 240 Signal Generation and Delivery ... 51, 237 Signal Names ... 51, 235 signal ... 208 BSD differences ... 234 concepts ... 51 definition of ... 19 examine and change action ... 58 examine and change blocked ... 60 examine pending ... 62 manipulate sets ... 57 mask ... 60 queueing ... 235 send to process ... 56 stacks ... 234 wait ... 62 signal() ... 25, 60, 189, 234-235, 237, 241-244, 283, 306 <signal.h> ... 24, 46, 49, 51, 56-64, 91, 162, 305 definition of ... 51

Signals ... 51, 234 sigpending() ... 33, 46, 58, 62-63, 234, 242, 306 definition of ... 62 rationale ... 244 SIGPIPE ... 51, 121, 213, 305 sigprocmask() ... 33, 46, 53, 58-63, 162, 237, 242, 244-245, 291, 306 definition of ... 60 rationale ... 244 SIGQUIT ... 24, 51, 133, 305 sigreturn() ... 234 SIG_BLOCK ... 61, 305 SIG_DFL ... 44, 51, 53-54, 59-60, 205, 234-235, 237-238, 305 SIG_ERR ... 305 SIG_IGN ... 44, 51, 53-54, 59, 203, 205, 229, 237-238, 305 SIG_SETMASK ... 61, 305 SIG_UNBLOCK ... 61, 305 SIGSEGV ... 51, 54, 61, 235-236, 305 sigsetjmp() ... 33, 162, 189, 291, 305 definition of ... 162 rationale ... 291 sigset_t ... 51, 57, 59-60, 62, 235, 306 sigstack() ... 234 SIGSTOP ... 52-54, 59, 61, 236, 239, 243, 305 sigsuspend() ... 33, 53, 58-60, 62-63, 162, 238, 241, 244-245, 291, 306 definition of ... 62 rationale ... 244 SIGSYS ... 235-236 SIGTERM ... 52, 235, 305 SIGTRAP ... 235-236 SIGTSTP ... 51, 53, 134, 205, 236, 239, 305 SIGTTIN ... 51, 53, 118, 130, 204-205, 236, 239, 305 SIGTTOU ... 51, 53, 121, 131, 140, 143, 203, 205, 236, 239, 277, 282, 305 SIGUSR1 ... 52, 235-236, 305 SIGUSR2 ... 52, 235-236, 305 sigvec() ... 243 sin() ... 305 definition of ... 151 sinh() ... 305 definition of ... 151 S_IRGRP ... 30, 103, 303, 308 S_IROTH ... 30, 103, 303, 308 S_IRUSR ... 30, 103, 303, 308 S_IRWXG ... 30, 103, 303, 308

S_IRWXO ... 30, 103, 303, 308 S_IRWXU ... 30, 103, 303, 308 S_ISBLK ... 30, 102, 303, 308 S_ISCHR ... 30, 102, 303, 308 S_ISDIR ... 30, 102, 303, 308 S_ISFIFO ... 30, 102, 303, 308 S_ISGID ... 30, 103, 106-107, 260-262, 303, 308 S_ISREG ... 30, 102, 303, 308 S_ISUID ... 30, 103, 106-107, 260-262, 303, 308 S_IWGRP ... 30, 103, 303, 308 S_IWOTH ... 30, 103, 155, 303, 308 S_IWUSR ... 30, 103, 303, 308 S_IXGRP ... 30, 103, 303, 308 S_IXOTH ... 30, 103, 303, 308 S_IXUSR ... 30, 103, 303, 308 size_t ... 87, 116, 215, 266, 306-307, 309-311 definition of ... 27 slash definition of ... 20 sleep() ... 63, 65-66, 240-241, 244-246, 311 definition of ... 65 rationale ... 245 socket() ... 287 sockets ... 299 Solely by POSIX.1 ... 188 Solely by the C Standard ... 188 Special Characters ... 133, 279 Special Control Characters ... 140, 280 Special File Creation ... 94, 258 Special Symbol (CLK_TCK) ... 81, 252 specific implementation ... 201 speed_t ... 141, 310 definition of ... 142 sprintf() ... 307 definition of ... 151 sqrt() ... 305 definition of ... 151 srand() ... 307 definition of ... 151 sscanf() ... 307 definition of ... 151 {SSIZE_MAX} ... 27, 117, 119, 215, 304 definition of ... 37 ssize_t ... 33-34, 37, 116, 118, 215, 309, 311 definition of ... 27 standards database ... 181 graphics ... 181 language ... 181 networking ... 179 open systems ... 179

START ... 134, 137, 141, 146 stat ... 103, 308 Structure ... 101 stat() ... 21-22, 33, 46, 91, 95-96, 101, 103-105, 107, 172-173, 176-177, 193, 199, 253, 260, 262, 308 definition of ... 103 rationale ... 260 <stat.h> (see <sys/stat.h>) st_atime ... 22, 45, 84, 89, 94-95, 101, 103, 113, 117, 160, 211, 254, 262, 267, 288, 308 definition of ... 101 st_ctime ... 22, 89-90, 92, 94-95, 97-98, 100-101, 103, 106-107, 109, 113, 120, 160-161, 288, 308 definition of ... 101 <stdarg.h> ... 306 <stddef.h> ... 306 STDERR_FILENO ... 156, 311 st_dev ... 101, 260, 308 definition of ... 101 STDIN_FILENO ... 156, 311 <stdio.h> ... 30, 78, 156-157, 306 <stdlib.h> ... 77, 194-195, 307 STDOUT_FILENO ... 156, 311 st_gid ... 101, 308 definition of ... 101 st_ino ... 101, 260, 308 definition of ... 101 st_mode ... 101-102 definition of ... 101 st_mtime ... 22, 89-90, 92, 94-95, 97-98, 100-101, 103, 113, 120, 160-161, 262, 288, 308 definition of ... 101 st_nlink ... 101, 214, 308 definition of ... 101 STOP ... 134, 137, 141, 146 strcat() ... 308 definition of ... 151 *strchr()* ... 308 definition of ... 151 strcmp() ... 308 definition of ... 151 strcoll() ... 308 strcpy() ... 308 definition of ... 151 strcspn() ... 308 definition of ... 151 st_rdev ... 219 stream binary ... 152 open ... 157 pointer ... 156

text ... 152 {STREAM_MAX} ... 80, 304 definition of ... 35 streams Eighth Edition ... 198 strerror() ... 308 strftime() ... 152-153, 162, 310 definition of ... 152 Strictly Conforming POSIX.1 Application ... 4, 192 string definition of ... 29 <string.h> ... 307 strlen() ... 308 definition of ... 151 strncat() ... 308 definition of ... 151 *strncmp()* ... 308 definition of ... 151 *strncpy*() ... 308 definition of ... 151 strpbrk() ... 308 definition of ... 151 strrchr() ... 308 definition of ... 151 strspn() ... 308 definition of ... 151 strstr() ... 308 definition of ... 151 strtod() ... 307 strtok() ... 308 definition of ... 151 strtol() ... 307 strtoul() ... 307 struct direct ... 253 dirent ... 83, 85, 253-254, 302 flock ... 124, 302 group ... 166, 303 lconv ... 304 passwd ... 167, 305 sigaction ... 58, 306 stat ... 103, 308 termios ... 141, 143, 310 tm ... 310 tms ... 76, 308 utimbuf ... 108, 311 utsname ... 74, 309 structures additions to ... 218 strxfrm() ... 308 st_size ... 308 definition of ... 101

Part 1: SYSTEM API [C LANGUAGE]

st uid ... 101, 308 definition of ... 101 su ... 242,246 subshells ... 204 Suggested Feature Test Macros ... 222 super-user ... 199, 208, 241, 246, 258-259, 261-262, 295 supplementary group ID definition of ... 20 supplementary groups ... 14-15, 20, 34, 39, 44-45, 69-70, 106-107, 208, 247, 262 supported definition of ... 11 rationale ... 197 Suspend Process Execution ... 64, 245 SUSP ... 134, 139-140 Symbolic Constant for the lseek() Function 37 Symbolic Constants for the access() Function ... 37-38, 225 Symbolic Constants for the lseek() Function ... 38, 225 Symbolic Constants ... 37, 225 symbolic link ... 294, 297-298 Symbols From the C Standard ... 29, 217 symbols C Standard ... 29 POSIX.1 ... 30 sysconf() ... 34-35, 77, 80-81, 202, 236, 252, 263, 311 definition of ... 80 rationale ... 252 <sys/dir.h> ... 253 sysname ... 74, 309 <sys/stat.h> ... 46, 88, 91-92, 94-96, 101, 103-104, 106-108, 110, 173, 177, 214-215, 257, 260, 308 definition of ... 101 File Modes ... 102 File Types ... 102 Time Entries ... 103 system call ... 208 restarting ... 234 System Databases ... 165, 293 system documentation definition of ... 11 System Identification ... 74, 249 System III ... 207, 213, 222, 234, 249, 262, 264, 267, 273 System Name ... 249 system process definition of ... 20

System V ... 199, 205-206, 222, 225, 230, 233-238, 241, 243-244, 246-249, 253, 258-259, 262-265, 267, 270-271, 273-274, 282-283, 291, 297 system definition of ... 20 get configurable variables ... 80 name ... 74 system() ... 231, 288-289, 307 <sys/times.h> ... 76,308 definition of ... 76 <sys/types.h> ... 27, 41, 47, 56, 67-68, 70, 72-73, 83, 88, 91, 94-95, 103, 106-108, 121, 127, 147-148, 166-167, 213, 215, 217, 221, 250, 260, 293, 309 definition of ... 27 <sys/utsname.h> ... 74,309 definition of ... 74 <sys/wait.h> ... 47-48,309 definition of ... 47

Т

tan() ... 305 definition of ... 151 tanh() ... 305 definition of ... 151 tapecntl() ... 274 tar ... 169, 294-297 format ... 169 Header Block ... 170 <tar.h> ... 170 tc*() ... 275 tcdrain() ... 33, 145-146, 279, 310 definition of ... 145 rationale ... 282 *tcflag_t* ... 310 definition of ... 136 tcflow() ... 33, 145-147, 310 definition of ... 145 rationale ... 282 tcflush() ... 33, 145-146, 310 definition of ... 145 rationale ... 282 tcgetattr() ... 33, 39, 141, 143-145, 204, 281-282, 310 definition of ... 143 rationale ... 282 *tcgetpgrp()* ... 147-148, 204, 276, 282, 311 definition of ... 147 rationale ... 282 tcsendbreak() ... 33, 145-146, 310 definition of ... 145 rationale ... 282

tcsetattr() ... 33, 39, 141-145, 204, 275, 281-282, 310 definition of ... 143 rationale ... 282 tcsetpgrp() ... 74, 130, 148-149, 203-204, 283, 311 definition of ... 148 rationale ... 283 telldir() ... 255-256 rationale for omission ... 255 termcntl() ... 274 Terminal Access Control ... 130, 277 terminal device (see terminal) Terminal Identification ... 78, 251 terminal access control ... 130 baud rates ... 141 canonical mode input ... 132 closing ... 135 control functions ... 143 control modes ... 138 controlling ... 12, 130 definition of ... 20 determine device name ... 79 general interface ... 129 generate pathname ... 78 get and set state ... 143 input ... 131 input modes ... 136 line control functions ... 145 local modes ... 139 modem disconnect ... 135 noncanonical mode input ... 132 opening ... 129 output ... 133 output modes ... 137 process group ... 129 setting parameters ... 135 special characters ... 133 special control characters ... 140 Terminate a Process ... 49, 232 Terminology and General Requirements ... 9 Terminology ... 10, 197 termios c_cc Special Control Characters ... 140 termios c_cflag Field ... 138 termios c_iflag Field ... 136 termios c_lflag Field ... 139 termios ... 141, 143, 310 Baud Rate Values ... 141 definition of ... 129 get and set state ... 143 line control functions ... 145 setting parameters ... 135 Structure ... 135-136, 279

<termios.h> ... 135-147,309 definition of ... 135 TERM definition of ... 28 variable ... 28 terms ... 11 text stream ... 152 time zone ... 29, 152-153, 162, 189, 283-285 Time ... 75, 249 time extensions to C Standard functions ... 152 get process ... 76 time() ... 33, 75-77, 189, 212, 249-250, 310 definition of ... 75 rationale ... 249 <time.h> ... 75, 81, 152, 162, 217, 250, 310 Timer Operations ... 63, 244 TIME ... 132-133, 139 times() ... 33, 42, 46, 49, 76-77, 199, 212, 244, 249, 308 definition of ... 76 rationale ... 250 <times.h> (see <sys/times.h>) time_t ... 29, 75-76, 101, 109, 200, 217, 249-250, 310 definition of ... 27 tm ... 310 tm_hour ... 19, 310 tm_isdst ... 310 tm_mday ... 200, 310 tm_min ... 19, 310 tm_mon ... 200, 310 tmpfile() ... 161, 290, 307 definition of ... 151 rationale ... 290 *tmpnam()* ... 307 definition of ... 151 tms ... 76, 308 tms_cstime ... 41, 45, 76-77, 308 tms_cutime ... 41, 45, 76-77, 308 tm_sec ... 19, 310 tms_stime ... 45, 76-77, 308 tms_utime ... 45, 76, 308 tm_wday ... 310 tm_yday ... 19, 200, 310 tm_year ... 19, 200, 310 toascii() ... 188 tolower() ... 301 definition of ... 151

tolower() ... 188 TOSTOP ... 131, 139-140, 203 toupper() ... 301 definition of ... 151 _toupper() ... 188 TRAILER!!! ... 176 trojan horse ... 199 ttyname() ... 79, 188, 251, 311 definition of ... 79 rationale ... 251 Types of Conformance ... 5, 194 <types.h> (see <sys/types.h>) tzname ... 153 {TZNAME_MAX} ... 80, 152, 304 definition of ... 35 T7. definition of ... 29 rationale ... 283 variable ... 29, 34-35, 152-153, 162, 216, 283-285, 319 tzset() ... 33, 153, 162-163, 310 definition of ... 162

U

{UCHAR_MAX} ... 34, 304 {UID_MAX} ... 215 uid_t ... 68, 101, 107, 167, 214-215, 293, 309 definition of ... 27 {UINT_MAX} ... 34, 245-246, 304 {ULONG_MAX} ... 34, 304 umask() ... 33, 46, 91-92, 95-96, 258, 308 definition of ... 91 rationale ... 258 umount() ... 206 uname() ... 33, 74-75, 249, 309 definition of ... 74 rationale ... 249 Structure Members ... 75 undefined ... 3, 6-7, 11, 19, 25, 27, 29-30, 32, 49, 54-55, 58, 61, 70, 84, 87-89, 117, 119, 124, 127, 132, 135, 137, 144, 153, 158-159, 170, 176, 190, 193, 197-198, 212, 232, 240, 251, 257, 288, 298-299, 319 definition of ... 11 rationale ... 198 ungetc() ... 160, 307 definition of ... 151 <unistd.h> ... 3, 27, 29-30, 32, 37-39, 80, 104-105, 110-111, 121, 127-128, 156, 173, 190, 215, 222-223, 225, 236, 251, 311 definition of ... 37

unlink() ... 93, 96-97, 99, 101, 103, 116, 162, 213, 258-259, 311 definition of ... 96 rationale ... 259 unsafe functions ... 240 unspecified ... 3-4, 11-12, 16, 20, 22-24, 27, 32, 42, 45, 47, 53-54, 56, 60, 62, 65-66, 71, 75-79, 83-84, 88-89, 92, 98, 101, 111, 116, 118, 121, 124, 129-130, 135, 141-142, 152, 154, 156-157, 159, 161, 165-167, 169, 172-174, 176, 185, 190, 193, 197-198, 211, 233, 237, 243-245, 250, 253, 276 definition of ... 11 rationale ... 198 User Database Access ... 167, 293 User Identification ... 68, 246 user ID definition of ... 20 effective ... 13, 68 get effective ... 68 get real ... 68 real ... 18,68 saved set- ... 19 set ... 68 user name definition of ... 20 get ... 71 USER variable ... 216 ushort_t ... 213 {USHRT_MAX} ... 34, 304 ustat() ... 260 utimbuf ... 108, 311 utime() ... 33, 103, 108-109, 262, 312 definition of ... 108 rationale ... 262 <utime.h> ... 108-109, 311 utsname ... 74, 309 <utsname.h> (see <sys/utsname.h>)

V

Values for cpio c_mode Field ... 175
VEOF ... 140
VEOL ... 140
VERASE ... 140
Version 7 ... 211, 222, 234, 241, 244, 249-250, 253, 262, 265, 267, 273, 296
version ... 74, 309
vfork() ... 248

ISO/IEC 9945-1: 1990 IEEE Std 1003.1-1990

vfprintf() ... 160-161, 290, 307
rationale ... 290
vhangup() ... 206, 233
VINTR ... 140
VKILL ... 140
VMIN ... 140-141
vprintf() ... 160-161, 290, 307
rationale ... 290
VQUIT ... 140
vsprintf() ... 307
VSTART ... 140-141
VSTOP ... 140-141
VSUSP ... 140
VTIME ... 140-141

W

Wait for a Signal ... 62, 244 Wait for Process Termination ... 47, 231 wait3() ... 231 wait() ... 18, 24, 33, 46-50, 76-77, 212, 226, 231-232, 236, 238, 240-241, 288, 309 rationale ... 230 <wait.h> (see <sys/wait.h>) waitpid() ... 18, 24, 33, 46-50, 76, 203, 207, 214, 226, 231-232, 240, 242-243, 309 definition of ... 47 rationale ... 230 wchar_t ... 307 wcstombs() ... 307 wctomb() ... 307 WERASE ... 278 WEXITSTATUS ... 48, 309 WIFEXITED ... 48, 309 WIFSIGNALED ... 48, 309 WIFSTOPPED ... 48, 232, 309 WNOHANG ... 47, 243, 309 W_OK ... 37-38, 105, 311 {WORD_BIT} ... 223 Working Directory Pathname ... 256 Working Directory ... 86, 256 working directory change ... 86 definition of ... 20 Write to a File ... 118, 267 write() ... 26, 91, 103, 114, 118-120, 128, 131, 133, 135, 158, 160-161, 170, 202-203, 213, 215, 238-240, 264-266, 269, 272, 277, 290, 300, 311 definition of ... 118 rationale ... 267

writev() ... 269
Writing Data and Output Processing ... 133, 279
WSTOPSIG ... 48, 309
WTERMSIG ... 48, 309
WUNTRACED ... 47, 203, 231-232, 309

Χ

X_OK ... 37-38, 105, 261, 311

Acknowledgments

We wish to thank the following organizations for donating significant computer, printing, and editing resources to the production of the 1988 version of this standard, upon which this revision is primarily based: UniForum (formerly /usr/group), Amdahl Corporation, Digital Equipment Corporation, MASSCOMP, UniSoft Corporation, and X/Open.

We wish to thank the following organizations for donating significant computer, printing, and editing resources to the production of this revision: Hewlett-Packard Company and X/Open.

This document was also approved by ISO/IEC JTC-1 SC22/WG15 as ISO/IEC 9945-1:1990. The IEEE wishes to thank the advisory groups of the National Bodies participating in WG15 for their contributions: Austria, Belgium, Canada, Denmark, France, Germany, Japan, Netherlands, United Kingdom, USA, and USSR.

The IEEE also wishes to thank the delegates to WG15 for their contributions:

AUSTRIA Gerhard Schmitt Wolfgang Schwabl

CANADA

Joe Cote Patrick Dempster George Kriger Bernard Martineau Major Douglas J. Moore Arnie Powell Paul Renaud Richard Sniderman

CEC Phil Bertrand Manuel Carbajo Michel Colin

DENMARK Peter E. Cordsen Isak Korn Keld Simonsen Claus Tondering

FINLAND Jikka Haikala

FRANCE Pascal Beyles Christophe Binot Claude Bourstin Jean-Michel Cornu Yves Delarue Eric Dumas Maurice Fathi Gerald Krummeck Herve Schauer Hubert Zimmerman

GERMANY Ron Elliot Helmut Stiegler Claus Unger Rainer Zimmer

IRELAND Hans-Jurgen Kugler

JAPAN Uiromiak

Hiromichi Kogure Shigekatsu Makao Yasushi Nakahara Nobuo Saito

NETHERLANDS J. Van Katwijk Willem Wakker H.J. Weegenaar

SWEDEN Mat Linder UK Nigel Bevan Cornelia Boldyreff Dave Cannon Don Chacon Dominic Dunlop David Flint Don E. Follond Martin Kirk Neil Martin Brian Meek Kevin Murphy Ian Newman Philip Rushton

USA Robert Bismuth Steven L. Carter Terence S. Dowling Ron Elliott Dale Harris John Hill James D. Isaak Hal Jespersen Roger J. Martin Shane McCarron Barry Needham Donn S. Terry Alan Weaver

USSR V. Koukhar Ostapenko Georgy Pavlovich Also we wish to thank the organizations employing the members of the Working Group and the Balloting Group for both covering the expenses related to attending and participating in meetings, and donating the time required both in and out of meetings for this effort.

Absolut Software ACM ADDAMAX Adobe Systems, Inc. Aeon Technologies, Inc. AFNOR* AFIII **AGS Information Services** Air Force Institute of Technology Aktiengesellschaft Alcyon Corporation Alliant Computer Systems Alsup Amdahl Corporation* Analysts International ANFOR Anistics, Inc. ANSI Applicon Applied Network Technology **APT Data Services** Archives Ltd. ARIX **Associated Computer Experts** AT Computing/Toernooiveld Atlanta UNIX Users Group AT&T* AT&T Bell Laboratories **AT&T** Information Systems AT&T UNIX Europe Ltd. ATTAGE Australian Government Dept. of Science Australian UNIX Systems User Group Automation Resource Group Automation Technologies Axon Data **Baldwin Information Processing** Barrister Information Systems Corporation Batelle Columbus Labs **BBN** Communications Corporation **BBN** Laboratories Bell Communications Research Bell-Northern Research Ltd. Billy Shakespeare & Company **Boeing Aerospace Company Boeing Computer Services** BP America Research & Development Brake Systems, Inc. **British Airways** British Olivetti Ltd. British Standards Institute* British Telecom Brown & Sharpe Bull, Inc.

Bull Sems **Burroughs** Corporation Burtek, Inc. C. S. Draper Lab, Inc. California State University California University Calspan Corporation Carnegie-Mellon University CCTA, Riverwalk House **Celerity Computing** Central Computer & Telecommunications Central Intelligence Agency Centre National D'Etudes des Telecomm. CFI Charles River Data Systems* **Chemical Abstracts Service Chorus Systems** Citicorp Transaction Tech., Inc. **Classic Conferences Cleveland State University** CMC Ltd. Commission of the European Communities** CommUNIXations COMPASS **Compugraphic Corporation Computer Design Computer Systems Engineering Computer Systems News** Computer Systems Resources, Inc. Computer Task Group Computer Works Computer X, Inc. Computerworld Comtek Concord Data Systems Concurrent Computer Corporation* Control Data Corporation* **Convergent Technologies Convex Computer Corporation** COSMIC Cray Research, Inc.* Cullinet Software, Inc. **Custom Development Environments** Dana Computers Dansk Data Elektronik A/S DDE Dansk Standardiseringsradd Dartmouth College Data Connection **Data General Corporation** Data Logic Ltd. Data Systems Engineering DataBoard, Inc. Datamation Datamension Corporation

Datapoint Corporation DEC International **Defense Communication Agency** Deutsches Institut fur Normung DGM&S, Inc. **Digital Equipment Corporation**** Digital Equipment GmbH **Digital Sound Corporation Directorate Land Armament Douglas Aircraft Corporation** Dravo Automation Sciences (DAS) Eastman-Stuart Ltd. **Eclipse Systems Corporation EDS** Corporation EDS of Canada, Ltd. EDV Zentrum der TU Wien **EER Systems Corporation Electronic Data Systems Corporation Electronic Engineering Times** Electrospace Systems, Inc. **Emerald** City Emerging Technologies Group, Inc. **Encore** Computer ENEA DATA Svenska AB Epicom Epilogue Technology Corporation Ericsson Information Systems AB ETA Systems, Inc. **European Laboratory for Particle Physics** European UNIX Eurotherm International EUUG* Exxon Chemical Pakistan Ltd. Federal Judicial Center Fern Universitat FGAN/FFM Fidcom System Ltd. Flexible Automation **Flexible Computer Corporation** Floating Point Systems Ford Aerospace Corporation Ford Motor Company Fortune Systems Corporation Fourth Shift Corporation Free Software Foundation, Inc. Fujitsu America, Inc. Future Tech, Inc. FUUG (Finland) Gartner Group Geac Computer International GEC Telecomms Ltd. **General Dynamics** General Electric Corporation General Motors Corporation General Motors Technical Center George Mason University Georgia Institute of Technology Gilbert International Gould CSD

Gould Electronics Gould SEL* **Government** Computer News **Grebyn** Corporation Grumman Aircraft Grumman Data Grupo de Redes de Computadores GUUG (West Germany) Harrie Corporation Harris Corporation* Hayden Publishing HCR Corporation Hewlett-Packard Company** Hewlett-Packard/Apollo* High Tech Publishing Company, Inc. Hi-Tech Editorial, Inc. Honeywell Bull, Inc. Honeywell ISI Honeywell Ltd. Hughes Aircraft Company Human Computing Resources IAI MBT **IBM Corporation** IBM** Deutschland **IBM Federal Sector Division IBM Research Center** IBM Thomas J. Watson Research Center ICL IDC **IEEE** Computer Magazine **IEEE Computer Society IEEE Micro IEEE** Reflector **IEEE Standards Office Imperial College** Industrial Technology Institute InfoPro Systems Information Concepts Pty Ltd. Ing. C. Olivetti & C., SPA **Inside Information** Institute for Defense Analysis Institute of Software Academia Sinica Integrated Systems Design, Inc. Intel Corporation Interactive Systems Corporation Intergraph Corporation International Bureau of Software Test International Computers Ltd. Internet Systems Iowa State University Irish UNIX Systems Users Group Iskra Automatika ISO-OSCRAC **Israel Aircraft Industries** Italian UNIX Systems User Group i2u **Itom International** ITSCJ* ITT ITUSA

Japanese Industrial Standards Committee Johnson Controls, Inc. KAIST Kendall Square Research Corporation K.E.T.R.I. Key Tech King Abdulziz University Korean UNIX User Group Lachman Associates, Inc. Lawrence Livermore National Laboratory Liberty Mutual Research Center Lisp Machine, Inc. LM Ericsson Lockheed Lockheed S.O.C Logicon, Inc. Lowell University LSI Appl. Info. & Learning Center LTV Aerospace & Defense Maharishi International University Mallinckrodt Institute of Radiology Mark Williams Company Markor Martin Marietta Aerospace Martin Marietta Data Systems Martin Marietta Energy Systems MASSCOMP Maxim Technologies, Inc. M.B.F. Systems, Inc. McDonnell Douglas McDonnell Douglas Computer Company McGill University Mercury Computer Systems, Inc. Meridian Software Systems Microfocus Micrology, Inc. Microsoft Corporation Microtel Pacific Research Microtex Mindcraft, Inc. Mini-Micro Systems **MIPS** Computer Systems MIT-LCS Mitsubishi Electric Corporation Modcomp* Modular Systems Molecular Genetics, Inc. Monarch Data Systems Mortice Kern Systems, Inc. Motorola, Inc. Motorola Israel, Inc. Motorola MCD MT XINU NAPS, Inc. NASA NASA Ames Research Center NASA Johnson Space Center NASA Kennedy Space Center National Cancer Institute

National Computer Security Center National Electrical Manufacturers Assoc. National Institute of Standards & Technology National Research Council National UNIX Systems User Group Naval Ocean Systems Center Naval Postgraduate School Naval Surface Weapons Center Naval Underwater Systems Center NBI, Inc. NCR* NCR Cambridge NEC Info Systems Nederlands Normalisatie-instituut New Media Development Center New Zealand UNIX Systems User Group, Inc. Nicolet Instrument Corporation Nikkei Byte, Nikkei McGraw-Hill, Inc. Nippon Tel & Tel Corporation NIUIS 77 Nixdorf Computer AG Norsk Data Ltd. North Holland P&C Northern N. E. UNIX User Group Northwestern Bell Northwestern University Novell Novon Research Group NRAO OCLC Ohio State University **Open Software Foundation*** Open Systems Architects, Inc. **Oren Yuen Associates** Osterreichesches Normungsinstitut Oxford Systems, Inc. Pacific Marine Technology Palladium Data Systems Perennial Philadelphia Area Computer Society Philips and Picker Medical Systems Philips Data Systems Phillips Publishing, Inc. Planning Research Corporation Plum Hall, Inc. **Plus Five Computer Services** Polish Computer Society Politecnic Di Torino-DIP Automatic Politecnico Di Mikeno Portland Community College POSIX Software Group Prime Computer, Inc. Princeton University Programming Concepts, Inc. Purdue University Pyramid Technology Rabbit Software Corporation RAMTEC, Inc. **RCA Automated Systems**

RDA Logicon Relcom, Inc. Release 1.0 **Richmond Computerware** Ricoh Systems, Inc. **RJO Enterprises Rogers State University Rolm Corporation** Rolm Mil-Spec Computers S. J. Lipton, Inc. SAH Consulting Sandia National Laboratories* SAS Institute, Inc. Schlumberger Well Services SCI Systems, Inc. Scientific Computer Systems SCS SDRC Seattle/UNIX Group Seay Systems, Inc. SEI Information Technology SELENIA SP. A Sequent Computer Systems, Inc. Seybold Office Computing Group Shell Int. Pet. Mij. Shreerang Society Siemens AG Sigma, IPA Silicon Valley Net Simpact Associates, Inc. Singapore UNIX Association Softech, Inc. Softtech Software Engineering Company Software Laboratories Ltd. Software Magazine Software News Software Productivity Consortium Software Research Associates, Inc. Software-PEI SoHar Southwest Research Institute Spectra-Tek U.K. Ltd. Sperry Corporation Sperry Ltd. Education Centre Sphinx Ltd. SSC St. Lawrence College Standards Council of Canada Statskontoret Stellar Computer, Inc. Stewart Research Enterprises Stratus Computer Strong Consulting Structural Dynamics Research Corporation Structured Methods Summit Computer Systems, Inc. Sun Microsystems, Inc.* Syntactics

Syntek Systems System House, Inc.* Systems & Software Magazine Systems Development Corporation Tampere University of Technology Tandem Computers, Inc. Technical Solutions, Inc. Technical University of Delft Technische Universitat Berlin Teknekron Infoswitch Tektronix. Inc.* Telephone Organization of Thailand Teli Foretagssystem Tenis Software Consulting, Inc. Texas Instruments, Inc. **Texas Instruments-DSG Texas Internet Consulting** The Algoma Steel Corporation Ltd. The C Journal, InfoPro Systems The Charles Stark Draper Laboratory, Inc. The Foxboro Company The Instruction Set The MITRE Corporation The Santa Cruz Operation The Wollongong Group Thinking Machines Corporation TIS Torch Computers Ltd. Toshiba Corporation Treasury Board of Canada* TRW Tsinghua University U.K. UNIX Systems User Group UNI Karlsruhe, Informatik II Uni'C' Computer Systems UniForum* UniForum Canada Unigram:X Unigroup of New York, Inc. Uni-Ops **UNIQ** Digital Technologies **UniSoft Corporation** UniSoft Ltd. Unisys Corporation** UNIT-C Universitat Dortmund Universitat Zurich-Irchel University of California, Berkeley University of California, Irvine University of California, Los Angeles University of Colorado University of Copenhagen University of Evansville University of Hong Kong University of Indonesia University of Lowell University of Maryland University of Minnesota University of Nevada, Las Vegas*

University of New Mexico University of Portland University of Santa Cruz University of South Florida University of Surrey University of Technology University of Tennessee University of Texas at Arlington University of Texas at Austin University of Toronto University of Utah University of Victoria University of Vienna University of Wisconsin-Milwaukee University of Zagreb **UNIX Houston UNIX** International **UNIX** Review **UNIX Systems** UNIX Technologies UNIX User Group Austria UNIX Users of Minnesota UNIX World UNIX-C Club UNIX/WORLD **US Air Force** US Army US Army Ballistic Research Lab. US Army ISEC **US Department of Defense** US Department of H.U.D. US Dept. of Commerce NOAA/NOS US D.O.T. Systems Center **US Federal Judicial Center US West Advanced Technologies USENIX Association* USSR State Committee for Standards** Varian Venturcom Verdix Corporation Veritas Technology, Inc. Videoton Virginia Polytech & State University Wang Laboratories, Inc. West Virginia University Western Digital Whitesmiths, Ltd. Wind River Systems, Inc. Woods Hole Oceanographic Institution World UNIX & C **XIOS Systems Corporation** X/Open Company Ltd. Yates Ventures

In the preceding list, the organizations marked with an asterisk (*) have hosted 1003 Working Group meetings since the group's inception in 1985, providing useful logistical support for the ongoing work of the committees.







ISO/IEC POSIX.1 The Source for UNIX[®] Standardization Information

ISO/IEC 9945-1 : 1990 (POSIX .1) defines a standard operating system interface and environment based on the UNIX Operating System. It supports applications portability at the source-code level between multivendor computer systems.

POSIX.1 establishes a set of basic services fundamental to the efficient construction of applications programs. Access to these services is provided through an operating system using the C programming language. The OS interface establishes standard semantics and syntax, and allows for application developers to design portable applications.

Complementing existing standards related to computer languages, database management, and computer graphics, POSIX.1 is designed for and used by both application developers and system implementors. It provides a solid base for the systems procurement and evaluation processes. By specifying POSIX.1 conformance, system purchasers can productively manage their software environments to achieve the benefits of applications portability. Likewise, hardware and software suppliers and developers will have a clear specification to follow in designing their systems and applications for the open software environment.

POSIX.1 constitutes a major step in the industry toward providing a comprehensive standardized applications environment. The IEEE's POSIX Committee is continuing POSIX-related standards work in areas such as POSIX-based Open System Environment (IEEE Project 1003.0), Shell and Utilities (IEEE Project 1003.2), Test Methods (IEEE Project 1003.3), Real-Time Systems Interfaces (IEEE Project 1003.4), An Ada Language Binding for POSIX (IEEE Project 1003.5), and a FOR-TRAN Language Binding to POSIX (IEEE Project 1003.9).

UNIX is a registered trademark of UNIX System Laboratories in the U.S. and other countries

New POSIX Seminar

IEEE Seminars . . . Keeping You Competitive Through Standards

If you define architectures for large distribution systems, set standards policy for your organization, or evaluate functional, performance, and integrity requirements, be sure to register for the new IEEE POSIX Seminar.

In Spring 1991, IEEE is launching a 2-day seminar based on the IEEE Project 1003.0, Guide to POSIX Open Systems Environments, and other IEEE projects concerning open systems. You will learn about frameworks and profiles for Information Technology standards, and how they directly apply to your organization.

For more information, call toll free in the USA at 1 (800) 678-IEEE and ask for Seminars on Standards. Or write to IEEE Standards Seminars, 445 Hoes Lane, PO Box 1331, Piscataway, NJ 08855-1331 USA. ISBN 1-55937-061-0