



man pages section 5: Standards, Environments, and Macros

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 817-3946-10
September 2004

Copyright 2004 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2004 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Certaines parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux États-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux États-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux États-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



040714@9061



Contents

Preface 7

Introduction 13

Intro(5) 14

Standards, Environments, and Macros 15

ascii(5) 16

attributes(5) 18

charmap(5) 26

crypt_bsdbf(5) 29

crypt_bsdmd5(5) 30

crypt_sunmd5(5) 31

crypt_unix(5) 32

dhcp(5) 33

dhcp_modules(5) 35

environ(5) 37

eqnchar(5) 42

extensions(5) 43

filesystem(5) 44

fnmatch(5) 63

fns(5) 67

fns_dns(5) 69

fns_files(5) 71

fns_initial_context(5) 73

fns_nis+(5) 77

fns_nis(5) 79
fns_policies(5) 81
fns_references(5) 85
fns_x500(5) 88
formats(5) 91
fsattr(5) 96
iconv_1250(5) 107
iconv_1251(5) 113
iconv(5) 122
iconv_646(5) 126
iconv_852(5) 129
iconv_8859-1(5) 135
iconv_8859-2(5) 141
iconv_8859-5(5) 147
iconv_dhn(5) 155
iconv_koi8-r(5) 159
iconv_mac_cyr(5) 167
iconv_maz(5) 175
iconv_pc_cyr(5) 179
iconv_unicode(5) 185
isalist(5) 190
largefile(5) 192
lf64(5) 196
lfcompile(5) 203
lfcompile64(5) 206
live_upgrade(5) 208
locale(5) 212
man(5) 238
mansun(5) 242
me(5) 246
mm(5) 251
ms(5) 258
nfssec(5) 263
pam_authtok_check(5) 265
pam_authtok_get(5) 267
pam_authtok_store(5) 269
pam_dhkeys(5) 270
pam_dial_auth(5) 272

pam_krb5(5) 273
pam_ldap(5) 278
pam_passwd_auth(5) 283
pam_projects(5) 285
pam_rhosts_auth(5) 286
pam_roles(5) 287
pam_sample(5) 289
pam_smartcard(5) 291
pam_unix(5) 293
pam_unix_account(5) 296
pam_unix_auth(5) 297
pam_unix_session(5) 299
prof(5) 300
rbac(5) 301
regex(5) 304
regexp(5) 313
SEAM(5) 320
sgml(5) 322
smartcard(5) 326
standards(5) 328
sticky(5) 333
term(5) 334
vgrindefs(5) 338
wbem(5) 341

Index 345

Preface

Both novice users and those familiar with the SunOS operating system can use online man pages to obtain information about the system and its features. A man page is intended to answer concisely the question “What does it do?” The man pages in general comprise a reference manual. They are not intended to be a tutorial.

Overview

The following contains a brief description of each man page section and the information it references:

- Section 1 describes, in alphabetical order, commands available with the operating system.
- Section 1M describes, in alphabetical order, commands that are used chiefly for system maintenance and administration purposes.
- Section 2 describes all of the system calls. Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible returned value.
- Section 3 describes functions found in various libraries, other than those functions that directly invoke UNIX system primitives, which are described in Section 2.
- Section 4 outlines the formats of various files. The C structure declarations for the file formats are given where applicable.
- Section 5 contains miscellaneous documentation such as character-set tables.
- Section 6 contains available games and demos.
- Section 7 describes various special files that refer to specific hardware peripherals and device drivers. STREAMS software drivers, modules and the STREAMS-generic set of system calls are also described.

- Section 9 provides reference information needed to write device drivers in the kernel environment. It describes two device driver interface specifications: the Device Driver Interface (DDI) and the Driver/Kernel Interface (DKI).
- Section 9E describes the DDI/DKI, DDI-only, and DKI-only entry-point routines a developer can include in a device driver.
- Section 9F describes the kernel functions available for use by device drivers.
- Section 9S describes the data structures used by drivers to share information between the driver and the kernel.

Below is a generic format for man pages. The man pages of each manual section generally follow this order, but include only needed headings. For example, if there are no bugs to report, there is no BUGS section. See the `intro` pages for more information and detail about each section, and `man(1)` for more information about man pages in general.

NAME	This section gives the names of the commands or functions documented, followed by a brief description of what they do.								
SYNOPSIS	<p>This section shows the syntax of commands or functions. When a command or file does not exist in the standard path, its full path name is shown. Options and arguments are alphabetized, with single letter arguments first, and options with arguments next, unless a different argument order is required.</p> <p>The following special characters are used in this section:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 10px;">[]</td> <td>Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.</td> </tr> <tr> <td style="padding-right: 10px;">. . .</td> <td>Ellipses. Several values can be provided for the previous argument, or the previous argument can be specified multiple times, for example, "filename . . .".</td> </tr> <tr> <td style="padding-right: 10px;"> </td> <td>Separator. Only one of the arguments separated by this character can be specified at a time.</td> </tr> <tr> <td style="padding-right: 10px;">{ }</td> <td>Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.</td> </tr> </table>	[]	Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.	. . .	Ellipses. Several values can be provided for the previous argument, or the previous argument can be specified multiple times, for example, "filename . . .".		Separator. Only one of the arguments separated by this character can be specified at a time.	{ }	Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.
[]	Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.								
. . .	Ellipses. Several values can be provided for the previous argument, or the previous argument can be specified multiple times, for example, "filename . . .".								
	Separator. Only one of the arguments separated by this character can be specified at a time.								
{ }	Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.								

PROTOCOL	This section occurs only in subsection 3R to indicate the protocol description file.
DESCRIPTION	This section defines the functionality and behavior of the service. Thus it describes concisely what the command does. It does not discuss OPTIONS or cite EXAMPLES. Interactive commands, subcommands, requests, macros, and functions are described under USAGE.
IOCTL	This section appears on pages in Section 7 only. Only the device class that supplies appropriate parameters to the <code>ioctl(2)</code> system call is called <code>ioctl</code> and generates its own heading. <code>ioctl</code> calls for a specific device are listed alphabetically (on the man page for that specific device). <code>ioctl</code> calls are used for a particular class of devices all of which have an <code>io</code> ending, such as <code>mtio(7I)</code> .
OPTIONS	This section lists the command options with a concise summary of what each option does. The options are listed literally and in the order they appear in the SYNOPSIS section. Possible arguments to options are discussed under the option, and where appropriate, default values are supplied.
OPERANDS	This section lists the command operands and describes how they affect the actions of the command.
OUTPUT	This section describes the output – standard output, standard error, or output files – generated by the command.
RETURN VALUES	If the man page documents functions that return values, this section lists these values and describes the conditions under which they are returned. If a function can return only constant values, such as 0 or -1, these values are listed in tagged paragraphs. Otherwise, a single paragraph describes the return values of each function. Functions declared void do not return values, so they are not discussed in RETURN VALUES.
ERRORS	On failure, most functions place an error code in the global variable <code>errno</code> indicating why they failed. This section lists alphabetically all error codes a function can generate and describes the conditions that cause each error. When more than

	one condition can cause the same error, each condition is described in a separate paragraph under the error code.
USAGE	This section lists special rules, features, and commands that require in-depth explanations. The subsections listed here are used to explain built-in functionality: Commands Modifiers Variables Expressions Input Grammar
EXAMPLES	This section provides examples of usage or of how to use a command or function. Wherever possible a complete example including command-line entry and machine response is shown. Whenever an example is given, the prompt is shown as <code>example%</code> , or if the user must be superuser, <code>example#</code> . Examples are followed by explanations, variable substitution rules, or returned values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS, and USAGE sections.
ENVIRONMENT VARIABLES	This section lists any environment variables that the command or function affects, followed by a brief description of the effect.
EXIT STATUS	This section lists the values the command returns to the calling program or shell and the conditions that cause these values to be returned. Usually, zero is returned for successful completion, and values other than zero for various error conditions.
FILES	This section lists all file names referred to by the man page, files of interest, and files created or required by commands. Each is followed by a descriptive summary or explanation.
ATTRIBUTES	This section lists characteristics of commands, utilities, and device drivers by defining the attribute type and its corresponding value. See <code>attributes(5)</code> for more information.
SEE ALSO	This section lists references to other man pages, in-house documentation, and outside publications.

DIAGNOSTICS	This section lists diagnostic messages with a brief explanation of the condition causing the error.
WARNINGS	This section lists warnings about special conditions which could seriously affect your working conditions. This is not a list of diagnostics.
NOTES	This section lists additional information that does not belong anywhere else on the page. It takes the form of an aside to the user, covering points of special interest. Critical information is never covered here.
BUGS	This section describes known bugs and, wherever possible, suggests workarounds.

Introduction

Intro(5)

NAME	Intro – introduction to miscellany
DESCRIPTION	Among the topics presented in this section are: <ul style="list-style-type: none">Standards The POSIX (IEEE) Standards and the X/Open Specifications are described on the <code>standards</code> page.Environments The user environment (<code>environ</code>), the subset of the user environment that depends on language and cultural conventions (<code>locale</code>), the large file compilation environment (<code>lfcompile</code>), and the transitional compilation environment (<code>lfcompile64</code>) are described.Macros The macros to format Reference Manual pages (<code>man</code> and <code>mansun</code>) as well as other text format macros (<code>me</code>, <code>mm</code>, and <code>ms</code>) are described.Characters Tables of character sets (<code>ascii</code>, <code>charmap</code>, <code>eqnchar</code>, and <code>iconv</code>), file format notation (<code>formats</code>), file name pattern matching (<code>fnmatch</code>), and regular expressions (<code>regex</code> and <code>regexp</code>) are presented.FNS Topics concerning the Federated Naming Service (<code>fns</code>, <code>fns_initial_context</code>, <code>fns_policies</code>, and <code>fns_references</code>) are discussed.

Standards, Environments, and Macros

ascii(5)

NAME ascii – map of ASCII character set
SYNOPSIS cat /usr/pub/ascii
DESCRIPTION /usr/pub/ascii is a map of the ASCII character set, to be printed as needed. It contains octal and hexadecimal values for each character. While not included in that file, a chart of decimal values is also shown here.

Octal – Character

000 NUL	001 SOH	002 STX	003 ETX	004 EOT	005 ENQ	006 ACK	007 BEL
010 BS	011 HT	012 NL	013 VT	014 NP	015 CR	016 SO	017 SI
020 DLE	021 DC1	022 DC2	023 DC3	024 DC4	025 NAK	026 SYN	027 ETB
030 CAN	031 EM	032 SUB	033 ESC	034 FS	035 GS	036 RS	037 US
040 SP	041 !	042 "	043 #	044 \$	045 %	046 &	047 '
050 (051)	052 *	053 +	054 ,	055 -	056 .	057 /
060 0	061 1	062 2	063 3	064 4	065 5	066 6	067 7
070 8	071 9	072 :	073 ;	074 <	075 =	076 >	077 ?
100 @	101 A	102 B	103 C	104 D	105 E	106 F	107 G
110 H	111 I	112 J	113 K	114 L	115 M	116 N	117 O
120 P	121 Q	122 R	123 S	124 T	125 U	126 V	127 W
130 X	131 Y	132 Z	133 [134 \	135]	136 ^	137 _
140 `	141 a	142 b	143 c	144 d	145 e	146 f	147 g
150 h	151 i	152 j	153 k	154 l	155 m	156 n	157 o
160 p	161 q	162 r	163 s	164 t	165 u	166 v	167 w
170 x	171 y	172 z	173 {	174	175 }	176 ~	177 DEL

Hexadecimal – Character

00 NUL	01 SOH	02 STX	03 ETX	04 EOT	05 ENQ	06 ACK	07 BEL
08 BS	09 HT	0A NL	0B VT	0C NP	0D CR	0E SO	0F SI
10 DLE	11 DC1	12 DC2	13 DC3	14 DC4	15 NAK	16 SYN	17 ETB
18 CAN	19 EM	1A SUB	1B ESC	1C FS	1D GS	1E RS	1F US
20 SP	21 !	22 "	23 #	24 \$	25 %	26 &	27 '
28 (29)	2A *	2B +	2C ,	2D -	2E .	2F /
30 0	31 1	32 2	33 3	34 4	35 5	36 6	37 7
38 8	39 9	3A :	3B ;	3C <	3D =	3E >	3F ?
40 @	41 A	42 B	43 C	44 D	45 E	46 F	47 G
48 H	49 I	4A J	4B K	4C L	4D M	4E N	4F O
50 P	51 Q	52 R	53 S	54 T	55 U	56 V	57 W
58 X	59 Y	5A Z	5B [5C \	5D]	5E ^	5F _
60 `	61 a	62 b	63 c	64 d	65 e	66 f	67 g
68 h	69 i	6A j	6B k	6C l	6D m	6E n	6F o
70 p	71 q	72 r	73 s	74 t	75 u	76 v	77 w
78 x	79 y	7A z	7B {	7C	7D }	7E ~	7F DEL

Decimal – Character

0 NUL	1 SOH	2 STX	3 ETX	4 EOT	5 ENQ	6 ACK	7 BEL
8 BS	9 HT	10 NL	11 VT	12 NP	13 CR	14 SO	15 SI
16 DLE	17 DC1	18 DC2	19 DC3	20 DC4	21 NAK	22 SYN	23 ETB
24 CAN	25 EM	26 SUB	27 ESC	28 FS	29 GS	30 RS	31 US
32 SP	33 !	34 "	35 #	36 \$	37 %	38 &	39 '
40 (41)	42 *	43 +	44 ,	45 -	46 .	47 /
48 0	49 1	50 2	51 3	52 4	53 5	54 6	55 7
56 8	57 9	58 :	59 ;	60 <	61 =	62 >	63 ?
64 @	65 A	66 B	67 C	68 D	69 E	70 F	71 G
72 H	73 I	74 J	75 K	76 L	77 M	78 N	79 O

ascii(5)

80 P	81 Q	82 R	83 S	84 T	85 U	86 V	87 W
88 X	89 Y	90 Z	91 [92 \	93]	94 ^	95 _
96 `	97 a	98 b	99 c	100 d	101 e	102 f	103 g
104 h	105 i	106 j	107 k	108 l	109 m	110 n	111 o
112 p	113 q	114 r	115 s	116 t	117 u	118 v	119 w
120 x	121 y	122 z	123 {	124	125 }	126 ~	127 DEL

FILES /usr/pub/ascii

On-line chart of octal and hexadecimal values for the ASCII character set.

attributes(5)

NAME attributes, architecture, availability, CSI, stability, MT-Level – attributes of interfaces

DESCRIPTION The ATTRIBUTES section of a manual page contains a table (see below) defining attribute types and their corresponding values.

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC
Availability	SUNWcsu
CSI	Enabled
Interface Stability	Unstable
MT-Level	Safe

Architecture Architecture defines processor or specific hardware. (See -p option of uname(1)). In some cases, it may indicate required adapters or peripherals.

Availability This refers to the software package which contains the command or component being described on the man page. To be able to use the command, the indicated package must have been installed. For information on how to add a package see pkgadd(1M).

Code Set Independence (CSI) OS utilities and libraries which are free of dependencies on the properties of any code sets are said to have Code Set Independence (CSI). They have the attribute of being CSI enabled. This is in contrast to many commands and utilities in Solaris, for example, that work only with Extended Unix Codesets (EUC), an encoding method that allows concurrent support for up to four code sets and is commonly used to represent Asian character sets.

However, for practical reasons, this independence is not absolute. Certain assumptions are still applied to the current CSI implementation:

- File code is a superset of ASCII.
- To support multi-byte characters and null-terminated UNIX file names, the NULL and / (slash) characters cannot be part of any multi-byte characters.
- Only "stateless" file code encodings are supported. Stateless encoding avoids shift, locking shift, designation, invocation, and so forth, although single shift is not excluded.
- Process code (wchar_t values) is implementation dependent and can change over time or between implementations or between locales.
- Not every object in Solaris 2 and Solaris 7 can have names composed of arbitrary characters. The names of the following objects must be composed of ASCII characters:
 - User names, group name, and passwords
 - System name
 - Names of printers and special devices

- Names of terminals (/dev/tty*)
- Process ID numbers
- Message queues, semaphores, and shared memory labels.
- The following may be composed of ISO Latin-1 or EUC characters:
 - File names
 - Directory names
 - Command names
 - Shell variables and environmental variable names
 - Mount points for file systems
 - NIS key names and domain names
- The names of NFS shared files should be composed of ASCII characters. Although files and directories may have names and contents composed of characters from non-ASCII code sets, using only the ASCII codeset allows NFS mounting across any machine, regardless of localization. For the commands and utilities that are CSI enabled, all can handle single-byte and multi-byte locales released in 2.6. For applications to get full support of internationalization services, dynamic binding has to be applied. Statically bound programs will only get support for C and POSIX locales.

Interface Stability

Sun often provides developers with early access to new technologies, which allows developers to evaluate with them as soon as possible. Unfortunately, new technologies are prone to changes and standardization often results in interface incompatibility from previous versions.

To make reasonable risk assessments, developers need to know how likely an interface is to change in future releases. To aid developers in making these assessments, interface stability information is included on some manual pages for commands, entry-points, and file formats.

The more stable interfaces can safely be used by nearly all applications, because Sun will endeavor to ensure that these continue to work in future minor releases. Applications that depend only on Standard and Stable interfaces should reliably continue to function correctly on future minor releases (but not necessarily on earlier major releases).

The less stable interfaces allow experimentation and prototyping, but should be used only with the understanding that they might change incompatibly or even be dropped or replaced with alternatives in future minor releases.

“Interfaces” that Sun does not document (for example, most kernel data structures and some symbols in system header files) may be implementation artifacts. Such internal interfaces are not only subject to incompatible change or removal, but we are unlikely to mention such a change in release notes.

Release Levels

Products are given release levels, as well as names, to aid compatibility discussions. Each release level may also include changes suitable for lower levels.

attributes(5)

Release	Version	Significance
Major	x.0	Likely to contain major feature additions; adhere to different, possibly incompatible Standard revisions; and though unlikely, could change, drop, or replace Standard or Stable interfaces. Initial product releases are usually 1.0.
Minor	x.y	Compared to an x.0 or earlier release (y!=0), it's likely to contain: minor feature additions, compatible Standard and Stable interfaces, possibly incompatible Evolving interfaces, or likely incompatible Unstable interfaces.
Micro	x.y.z	Intended to be interface compatible with the previous release (z!=0), but likely to add bug fixes, performance enhancements, and support for additional hardware.

Classifications

The following table summarizes how stability level classifications relate to release level. The first column lists the Stability Level. The second column lists the Release Level for Incompatible Changes, and the third column lists other comments. For a complete discussion of individual classifications, see the appropriate subsection below.

Stability	Release	Comments
Standard	Major (x.0)	Actual or de facto.
Stable	Major (x.0)	Incompatibilities are exceptional.
Evolving	Minor (x.y)	Migration advice might accompany an incompatibility.
Unstable	Minor (x.y)	Experimental or transitional: incompatibilities are common.
External	Micro (x.y.z)	Not controlled by Sun: intrarelease incompatibilities are common.

Stability	Release	Comments
Obsolete	Minor (x.y)	Deprecated interface: likely to be removed in a future minor release.

The interface stability levels described in this manual page apply to both source and binary interfaces unless otherwise stated. The stability level of each interface is unknown unless explicitly stated.

Standard[: [*organization_name*,] *standard_name*, *version*]

The documented interface complies with the standard listed. If a standard is not specified the interface is defined by several standards. This is usually the hierarchy built up from the C Language (defined by ISO/IEC or K&R), SVID 3 and associated ABIs (defined by AT&T), the POSIX standards (defined by IEEE and ISO/IEC), and the Single UNIX Specifications (defined by The Open Group). See `standards(5)` for a complete list of these standards.

Most of these interfaces are defined by a formal standard, and controlled by a standards development organization. Changes will usually be made in accordance with approved changes to that standard. This stability level can also apply to interfaces that have been adopted (without a formal standard) by an "industry convention."

Support is provided for only the specified version(s) of a standard; support for later versions is not guaranteed. If the standards development organization approves a non-upward-compatible change to a Standard interface that Sun decides to support, Sun will announce a compatibility and migration strategy.

Programmers producing portable applications should rely on the interface descriptions present in the standard or specification to which the application is intended to conform, rather than the manual page descriptions of Standard interfaces. When the standard or specification allows alternative implementation choices, the manual page usually only describes the alternative implemented by Sun. The manual page also describes any compatible extensions to the base definition of Standard interfaces provided by Sun.

Stable

A Stable interface is a mature interface under Sun's control. Sun will try to avoid non-upwards-compatible changes to these interfaces, especially in minor or micro releases.

If support of a Stable interface must be discontinued, Sun will attempt to provide notification and the stability level changes to Obsolete.

Evolving

An Evolving interface may eventually become Standard or Stable but is still in transition.

attributes(5)

Sun will make reasonable efforts to ensure compatibility with previous releases as it evolves. When non-upwards compatible changes become necessary, they will occur in minor and major releases; such changes will be avoided in micro releases whenever possible. If such a change is necessary, it will be documented in the release notes for the affected release, and when feasible, Sun will provide migration aids for binary compatibility and continued source development.

External

An External interface is controlled by an entity other than Sun. At Sun's discretion, Sun can deliver as part of any release updated and possibly incompatible versions of such interfaces, subject to their availability from the controlling entity. This classification is typically applied to publicly available "freeware" and similar objects.

For External interfaces, Sun makes no claims regarding either source or binary compatibility between any two releases. Applications based on these interfaces might not work in future releases, including patches that contain External interfaces.

Unstable

An Unstable interface is provided to give developers early access to new or rapidly changing technology or as an interim solution to a problem for which a more stable solution is anticipated in the future.

For Unstable interfaces, Sun no claims about either source or binary compatibility from one minor release to another. Applications developed based on these interfaces may not work in future minor releases.

Obsolete: Scheduled for removal after *event*

An Obsolete interface is supported in the current release, but is scheduled to be removed in a future (minor) release. When support of an interface is to be discontinued, Sun will attempt to provide notification before discontinuing support. Use of an Obsolete interface may produce warning messages.

MT-Level

Libraries are classified into four categories which define their ability to support multiple threads. Manual pages containing routines that are of multiple or differing levels show this within their NOTES or USAGE section.

Safe

Safe is an attribute of code that can be called from a multithreaded application. The effect of calling into a Safe interface or a safe code segment is that the results are valid even when called by multiple threads. Often overlooked is the fact that the result of this Safe interface or safe code segment can have global consequences that affect all threads. For example, the action of opening or closing a file from one thread is visible by all the threads within a process. A multi-threaded application has the responsibility for using these interfaces in a safe manner, which is different from whether or not the interface is Safe. For example, a multi-threaded application that closes a file that is still in use by other threads within the application is not using the `close(2)` interface safely.

Unsafe

An Unsafe library contains global and static data that is not protected. It is not safe to use unless the application arranges for only one thread at a time to execute within the library. Unsafe libraries may contain routines that are Safe; however, most of the library's routines are unsafe to call.

The following table contains reentrant counterparts for Unsafe functions. This table is subject to change by Sun.

Reentrant functions for libc:

Unsafe Function	Reentrant counterpart
ctime	ctime_r
localtime	localtime_r
asctime	asctime_r
gmtime	gmtime_r
ctermid	ctermid_r
getlogin	getlogin_r
rand	rand_r
readdir	readdir_r
strtok	strtok_r
tmpnam	tmpnam_r

MT-Safe

An MT-Safe library is fully prepared for multithreaded access. It protects its global and static data with locks, and can provide a reasonable amount of concurrency. Note that a library can be safe to use, but not MT-Safe. For example, surrounding an entire library with a monitor makes the library Safe, but it supports no concurrency so it is not considered MT-Safe. An MT-Safe library must permit a reasonable amount of concurrency. (This definition's purpose is to give precision to what is meant when a library is described as Safe. The definition of a Safe library does not specify if the library supports concurrency. The MT-Safe definition makes it clear that the library is Safe, and supports some concurrency. This clarifies the Safe definition, which can mean anything from being single threaded to being any degree of multithreaded.)

Async-Signal-Safe

Async-Signal-Safe refers to particular library routines that can be safely called from a signal handler. A thread that is executing an Async-Signal-Safe routine will not deadlock with itself if interrupted by a signal. Signals are only a problem for MT-Safe routines that acquire locks.

attributes(5)

Signals are disabled when locks are acquired in Async-Signal-Safe routines. This prevents a signal handler that might acquire the same lock from being called. The list of Async-Signal-Safe functions includes:

_exit	access	aio_error
aio_return	aio_suspend	alarm
cfgetispeed	cfgetospeed	cfsetispeed
cfsetospeed	chdir	chmod
chown	clock_gettime	close
creat	dup	dup2
execle	execve	fcntl
fdatasync	fork	fstat
fsync	getegid	geteuid
getgid	getgroups	getpgrp
getpid	getppid	getuid
kill	link	lseek
mkdir	mkfifo	open
pathconf	pause	pipe
read	rename	rmdir
sem_post	sema_post	setgid
setpgid	setsid	setuid
sigaction	sigaddset	sigdelset
sigemptyset	sigfillset	sigismember
sigpending	sigprocmask	sigqueue
sigsuspend	sleep	stat
sysconf	tcdrain	tcflow
tcflush	tcgetattr	tcgetpgrp
tcsendbreak	tcsetattr	tcsetpgrp
thr_kill	thr_sigsetmask	time
timer_getoverrun	timer_gettime	timer_settime
times	umask	uname

unlink	utime	wait
waitpid	write	

MT-Safe with Exceptions

See the NOTES or USAGE sections of these pages for a description of the exceptions.

Safe with Exceptions

See the NOTES or USAGE sections of these pages for a description of the exceptions.

Fork1-Safe

A Fork1-Safe library releases the locks it had held whenever `fork1(2)` is called in a Solaris thread program, or `fork(2)` in a POSIX (see `standards(5)`) thread program. Calling `fork(2)` in a POSIX thread program has the same semantic as calling `fork1(2)` in a Solaris thread program. All system calls, `libpthread`, and `libthread` are Fork1-Safe. Otherwise, you should handle the locking clean-up yourself (see `pthread_atfork(3C)`).

Cancel-Safety

If a multi-threaded application uses `pthread_cancel(3THR)` to cancel (that is, kill) a thread, it is possible that the target thread is killed while holding a resource, such as a lock or allocated memory. If the thread has not installed the appropriate cancellation cleanup handlers to release the resources appropriately (see `pthread_cancel(3THR)`), the application is "cancel-unsafe", that is, it is not safe with respect to cancellation. This unsafety could result in deadlocks due to locks not released by a thread that gets cancelled, or resource leaks; for example, memory not being freed on thread cancellation. All applications that use `pthread_cancel(3THR)` should ensure that they operate in a Cancel-Safe environment. Libraries that have cancellation points and which acquire resources such as locks or allocate memory dynamically, also contribute to the cancel-unsafety of applications that are linked with these libraries. This introduces another level of safety for libraries in a multi-threaded program: Cancel-Safety. There are two sub-categories of Cancel-Safety: Deferred-Cancel-Safety, and Asynchronous-Cancel-Safety. An application is considered to be Deferred-Cancel-Safe when it is Cancel-Safe for threads whose cancellation type is `PTHREAD_CANCEL_DEFERRED`. An application is considered to be Asynchronous-Cancel-Safe when it is Cancel-Safe for threads whose cancellation type is `PTHREAD_CANCEL_ASYNCHRONOUS`. Deferred-Cancel-Safety is easier to achieve than Asynchronous-Cancel-Safety, since a thread with the deferred cancellation type can be cancelled only at well-defined cancellation points, whereas a thread with the asynchronous cancellation type can be cancelled anywhere. Since all threads are created by default to have the deferred cancellation type, it may never be necessary to worry about asynchronous cancel safety. Indeed, most applications and libraries are expected to always be Asynchronous-Cancel-Unsafe. An application which is Asynchronous-Cancel-Safe is also, by definition, Deferred-Cancel-Safe.

SEE ALSO `uname(1)`, `pkgadd(1M)`, `Intro(3)`, `standards(5)`

charmap(5)

NAME	charmap – character set description file																																				
DESCRIPTION	<p>A character set description file or <i>charmap</i> defines characteristics for a coded character set. Other information about the coded character set may also be in the file. Coded character set character values are defined using symbolic character names followed by character encoding values.</p> <p>The character set description file provides:</p> <ul style="list-style-type: none">■ The capability to describe character set attributes (such as collation order or character classes) independent of character set encoding, and using only the characters in the portable character set. This makes it possible to create generic <code>localedef(1)</code> source files for all codesets that share the portable character set.■ Standardized symbolic names for all characters in the portable character set, making it possible to refer to any such character regardless of encoding.																																				
Symbolic Names	<p>Each symbolic name is included in the file and is mapped to a unique encoding value (except for those symbolic names that are shown with identical glyphs). If the control characters commonly associated with the symbolic names in the following table are supported by the implementation, the symbolic names and their corresponding encoding values are included in the file. Some of the encodings associated with the symbolic names in this table may be the same as characters in the portable character set table.</p> <table border="1"><tr><td><ACK></td><td><DC2></td><td><ENQ></td><td><FS></td><td><IS4></td><td><SOH></td></tr><tr><td><BEL></td><td><DC3></td><td><EOT></td><td><GS></td><td><LF></td><td><STX></td></tr><tr><td><BS></td><td><DC4></td><td><ESC></td><td><HT></td><td><NAK></td><td><SUB></td></tr><tr><td><CAN></td><td></td><td><ETB></td><td><IS1></td><td><RS></td><td><SYN></td></tr><tr><td><CR></td><td><DLE></td><td><ETX></td><td><IS2></td><td><SI></td><td><US></td></tr><tr><td><DC1></td><td></td><td><FF></td><td><IS3></td><td><SO></td><td><VT></td></tr></table>	<ACK>	<DC2>	<ENQ>	<FS>	<IS4>	<SOH>	<BEL>	<DC3>	<EOT>	<GS>	<LF>	<STX>	<BS>	<DC4>	<ESC>	<HT>	<NAK>	<SUB>	<CAN>		<ETB>	<IS1>	<RS>	<SYN>	<CR>	<DLE>	<ETX>	<IS2>	<SI>	<US>	<DC1>		<FF>	<IS3>	<SO>	<VT>
<ACK>	<DC2>	<ENQ>	<FS>	<IS4>	<SOH>																																
<BEL>	<DC3>	<EOT>	<GS>	<LF>	<STX>																																
<BS>	<DC4>	<ESC>	<HT>	<NAK>	<SUB>																																
<CAN>		<ETB>	<IS1>	<RS>	<SYN>																																
<CR>	<DLE>	<ETX>	<IS2>	<SI>	<US>																																
<DC1>		<FF>	<IS3>	<SO>	<VT>																																
Declarations	<p>The following declarations can precede the character definitions. Each must consist of the symbol shown in the following list, starting in column 1, including the surrounding brackets, followed by one or more blank characters, followed by the value to be assigned to the symbol.</p> <table><tr><td><code_set_name></td><td>The name of the coded character set for which the character set description file is defined.</td></tr><tr><td><mb_cur_max></td><td>The maximum number of bytes in a multi-byte character. This defaults to 1.</td></tr><tr><td><mb_cur_min></td><td>An unsigned positive integer value that defines the minimum number of bytes in a character for the encoded character set.</td></tr></table>	<code_set_name>	The name of the coded character set for which the character set description file is defined.	<mb_cur_max>	The maximum number of bytes in a multi-byte character. This defaults to 1.	<mb_cur_min>	An unsigned positive integer value that defines the minimum number of bytes in a character for the encoded character set.																														
<code_set_name>	The name of the coded character set for which the character set description file is defined.																																				
<mb_cur_max>	The maximum number of bytes in a multi-byte character. This defaults to 1.																																				
<mb_cur_min>	An unsigned positive integer value that defines the minimum number of bytes in a character for the encoded character set.																																				

<code><escape_char></code>	The escape character used to indicate that the characters following will be interpreted in a special way, as defined later in this section. This defaults to backslash (<code>\</code>), which is the character glyph used in all the following text and examples, unless otherwise noted.
<code><comment_char></code>	The character that when placed in column 1 of a charmap line, is used to indicate that the line is to be ignored. The default character is the number sign (#).

Format The character set mapping definitions will be all the lines immediately following an identifier line containing the string `CHARMAP` starting in column 1, and preceding a trailer line containing the string `END CHARMAP` starting in column 1. Empty lines and lines containing a `<comment_char>` in the first column will be ignored. Each non-comment line of the character set mapping definition (that is, between the `CHARMAP` and `END CHARMAP` lines of the file) must be in either of two forms:

```
"%s %s %s\n" , <symbolic-name>, <encoding>, <comments>
```

or

```
"%s. . .%s %s %s\n" , <symbolic-name>, <symbolic-name>, <encoding>, <comments>
```

In the first format, the line in the character set mapping definition defines a single symbolic name and a corresponding encoding. A character following an escape character is interpreted as itself; for example, the sequence `<\i>` represents the symbolic name `\` enclosed between angle brackets.

In the second format, the line in the character set mapping definition defines a range of one or more symbolic names. In this form, the symbolic names must consist of zero or more non-numeric characters, followed by an integer formed by one or more decimal digits. The characters preceding the integer must be identical in the two symbolic names, and the integer formed by the digits in the second symbolic name must be equal to or greater than the integer formed by the digits in the first name. This is interpreted as a series of symbolic names formed from the common part and each of the integers between the first and the second integer, inclusive. As an example, `<j0101>. . . <j0104>` is interpreted as the symbolic names `<j0101>`, `<j0102>`, `<j0103>`, and `<j0104>`, in that order.

A character set mapping definition line must exist for all symbolic names and must define the coded character value that corresponds to the character glyph indicated in the table, or the coded character value that corresponds with the control character symbolic name. If the control characters commonly associated with the symbolic

charmap(5)

names are supported by the implementation, the symbolic name and the corresponding encoding value must be included in the file. Additional unique symbolic names may be included. A coded character value can be represented by more than one symbolic name.

The encoding part is expressed as one (for single-byte character values) or more concatenated decimal, octal or hexadecimal constants in the following formats:

```
"%cd%d" , <escape_char>,<decimal byte value>  
"%cx%x" , <escape_char>,<hexadecimal byte value>  
"%co" , <escape_char>,<octal byte value>
```

Decimal Constants

Decimal constants must be represented by two or three decimal digits, preceded by the escape character and the lower-case letter d; for example, \d05, \d97, or \d143. Hexadecimal constants must be represented by two hexadecimal digits, preceded by the escape character and the lower-case letter x; for example, \x05, \x61, or \x8f. Octal constants must be represented by two or three octal digits, preceded by the escape character; for example, \05, \141, or \217. In a portable charmap file, each constant must represent an 8-bit byte. Implementations supporting other byte sizes may allow constants to represent values larger than those that can be represented in 8-bit bytes, and to allow additional digits in constants. When constants are concatenated for multi-byte character values, they must be of the same type, and interpreted in byte order from first to last with the least significant byte of the multi-byte character specified by the last constant.

Ranges of Symbolic Names

In lines defining ranges of symbolic names, the encoded value is the value for the first symbolic name in the range (the symbolic name preceding the ellipsis). Subsequent symbolic names defined by the range will have encoding values in increasing order. For example, the line

```
<j0101>. . .<j0104>      \d129\d254
```

will be interpreted as:

```
<j0101>      \d129\d254  
<j0102>      \d129\d255  
<j0103>      \d130\d0  
<j0104>      \d130\d1
```

Note that this line will be interpreted as the example even on systems with bytes larger than 8 bits. The comment is optional.

SEE ALSO

locale(1) localedef(1) nl_langinfo(3C) extensions(5), locale(5)

crypt_bsdbf(5)

NAME crypt_bsdbf – password hashing module using Blowfish cryptographic algorithm

SYNOPSIS /usr/lib/security/\$ISA/crypt_bsdbf.so

DESCRIPTION The crypt_bsdbf module is a one-way password hashing module for use with crypt(3C) that uses the Blowfish cryptographic algorithm. The algorithm identifier for crypt.conf(4) and policy.conf(4) is 2a.

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO passwd(1), crypt(3C), crypt_genhash_impl(3C), crypt_gensalt(3C), crypt_gensalt_impl(3C), getpassphrase(3C), crypt.conf(4), passwd(4), policy.conf(4), attributes(5)

crypt_bsdmd5(5)

- NAME** crypt_bsdmd5 – password hashing module using MD5 message hash algorithm
- SYNOPSIS** /usr/lib/security/\$ISA/crypt_bsdmd5.so
- DESCRIPTION** The crypt_bsdmd5 module is a one-way password hashing module for use with crypt(3C) that uses the MD5 message hash algorithm. The algorithm identifier for crypt.conf(4) and policy.conf(4) is 1. The output is compatible with md5crypt on BSD and Linux systems.
- ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

- SEE ALSO** passwd(1), crypt(3C), crypt_genhash_impl(3C), crypt_gensalt(3C), crypt_gensalt_impl(3C), getpassphrase(3C), crypt.conf(4), passwd(4), policy.conf(4), attributes(5)

NAME	crypt_sunmd5 – password hashing module using MD5 message hash algorithm				
SYNOPSIS	/usr/lib/security/\$ISA/crypt_sunmd5.so				
DESCRIPTION	<p>The <code>crypt_sunmd5</code> module is a one-way password hashing module for use with <code>crypt(3C)</code> that uses the MD5 message hash algorithm. The algorithm identifier for <code>crypt.conf(4)</code> and <code>policy.conf(4)</code> is <code>md5</code>.</p> <p>This module is designed to make it difficult to crack passwords that use brute force attacks based on high speed MD5 implementations that use code inlining, unrolled loops, and table lookup.</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Safe				
SEE ALSO	<code>passwd(1)</code> , <code>crypt(3C)</code> , <code>crypt_genhash_impl(3C)</code> , <code>crypt_gensalt(3C)</code> , <code>crypt_gensalt_impl(3C)</code> , <code>getpassphrase(3C)</code> , <code>crypt.conf(4)</code> , <code>passwd(4)</code> , <code>policy.conf(4)</code> , <code>attributes(5)</code>				

crypt_unix(5)

NAME	crypt_unix – traditional UNIX crypt algorithm				
DESCRIPTION	<p>The <code>crypt_unix</code> algorithm is the traditional UNIX crypt algorithm. It is not considered sufficiently secure for current systems and is provided for backwards compatibility. The <code>crypt_sunmd5(5)</code>, <code>crypt_bsdmd5(5)</code>, or <code>crypt_bsdbf(5)</code> algorithm should be used instead.</p> <p>The algorithm identifier for <code>policy.conf(4)</code> is <code>__unix__</code>. There is no entry in <code>crypt.conf(4)</code> for this algorithm.</p> <p>The <code>crypt_unix</code> algorithm is internal to <code>libc</code> and provides the string encoding function used by <code>crypt(3C)</code> when the first character of the salt is not a "\$".</p> <p>This algorithm is based on a one-way encryption algorithm with variations intended (among other things) to frustrate use of hardware implementations of a key search. Only the first eight characters of the key passed to <code>crypt()</code> are used with this algorithm; the rest are silently ignored. The salt is a two-character string chosen from the set <code>[a-zA-Z0-9./]</code>. This string is used to perturb the hashing algorithm in one of 4096 different ways.</p>				
USAGE	The return value of the <code>crypt_unix</code> algorithm might not be portable among standard-conforming systems. See <code>standards(5)</code> .				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Safe				
SEE ALSO	<code>passwd(1)</code> , <code>crypt(3C)</code> , <code>crypt_genhash_impl(3C)</code> , <code>crypt_gensalt(3C)</code> , <code>crypt_gensalt_impl(3C)</code> , <code>getpassphrase(3C)</code> , <code>crypt.conf(4)</code> , <code>passwd(4)</code> , <code>policy.conf(4)</code> , <code>attributes(5)</code> , <code>crypt_bsdbf(5)</code> , <code>crypt_bsdmd5(5)</code> , <code>crypt_sunmd5(5)</code> , <code>standards(5)</code>				

NAME	dhcp – Dynamic Host Configuration Protocol	
DESCRIPTION	<p>Dynamic Host Configuration Protocol (DHCP) enables host systems in a TCP/IP network to be configured automatically for the network as they boot. DHCP uses a client/server mechanism: servers store configuration information for clients, and provide that information upon a client's request. The information can include the client's IP address and information about network services available to the client.</p> <p>This manual page provides a brief summary of the Solaris DHCP implementation.</p>	
Solaris DHCP Client	<p>The Solaris DHCP client is implemented as background daemon, <code>dhcpage</code>(1M). This daemon is started automatically during bootup if there exists at least one <code>dhcp.interface</code> file in <code>/etc</code>. Only interfaces with a corresponding <code>/etc/dhcp.interface</code> file are automatically configured during boot. Network parameters needed for system configuration during bootup are extracted from the information received by the daemon through the use of the <code>dhcpinfo</code>(1) command. The daemon's default behavior can be altered by changing the tunables in the <code>/etc/default/dhcpage</code> file. The daemon is controlled by the <code>ifconfig</code>(1M) utility. Check the status of the daemon using the <code>netstat</code>(1M) and <code>ifconfig</code>(1M) commands.</p>	
Solaris DHCP Server	<p>The Solaris DHCP server is implemented as a background daemon, <code>in.dhcpd</code>(1M). This daemon can deliver network configuration information to either BOOTP or DHCP clients. The Solaris DHCP service can be managed using the <code>dhcpcmgr</code>(1M) GUI or the command line utilities <code>dhcpconfig</code>(1M), <code>dhtadm</code>(1M), and <code>pntadm</code>(1M).</p>	
DHCP Configuration Tables	<p>The Solaris DHCP server stores client configuration information in the following two types of tables:</p>	
	dhcptab tables	<p>Contain macros and options (also known as symbols), used to construct a package of configuration information to send to each DHCP client. There exists only one <code>dhcptab</code> for the DHCP service. The <code>dhcptab</code>(4) can be viewed and modified using the <code>dhtadm</code>(1M) command or <code>dhcpcmgr</code>(1M) graphical utility. See <code>dhcptab</code>(4) for more information about the syntax of <code>dhcptab</code> records. See <code>dhcp_inittab</code>(4) for more information about the DHCP options and symbols.</p>
	DHCP network tables	<p>DHCP network tables, which contain mappings of client IDs to IP addresses and parameters associated with those addresses. Network tables are named with the IP address of the network, and can be created, viewed, and modified using the <code>pntadm</code> command or <code>dhcpcmgr</code> graphical utility. See <code>dhcp_network</code>(4) for more information about network tables.</p>

dhcp(5)

SEE ALSO | dhcpinfo(1), dhcpagent(1M), dhcpconfig(1M), dhcpmgr(1M), dhtadm(1M),
ifconfig(1M), in.dhcpd(1M), netstat(1M), pntadm(1M), syslog(3C),
dhcp_network(4), dhcptab(4), dhcpsvc.conf(4), dhcp_inittab(4),
dhcp_modules(5)

Solaris DHCP Service Developer's Guide

Alexander, S., and R. Droms. *RFC 2132, DHCP Options and BOOTP Vendor Extensions*. Silicon Graphics, Inc. Bucknell University. March 1997.

Droms, R. *RFC 1534, Interoperation Between DHCP and BOOTP*. Bucknell University. October 1993.

Droms, R. *RFC 2131, Dynamic Host Configuration Protocol*. Bucknell University. March 1997.

Wimer, W. *RFC 1542, Clarifications and Extensions for the Bootstrap Protocol*. Carnegie Mellon University. October 1993.

NAME	dhcp_modules – data storage modules for the DHCP service
DESCRIPTION	<p>This man page describes the characteristics of data storage modules (public modules) for use by the Solaris Dynamic Host Configuration Protocol (DHCP) service.</p> <p>Public modules are the part of the DHCP service architecture that encapsulate the details of storing DHCP service data in a data storage service. Examples of data storage services are NIS+, Oracle, and <code>ufs</code> file systems.</p> <p>Public modules are dynamic objects which can be shipped separately from the Solaris DHCP service. Once installed, a public module is visible to the DHCP service, and can be selected for use by the service through the DHCP service management interfaces (<code>dhcpcmgr(1M)</code>, <code>dhcpconfig(1M)</code>, <code>dhtadm(1M)</code>, and <code>pntadm(1M)</code>).</p> <p>Public modules may be provided by Sun Microsystems, Inc or by third parties.</p> <p>The Solaris DHCP service management architecture provides a mechanism for plugging in public module-specific administration functionality into the <code>dhcpcmgr(1M)</code> and <code>dhcpconfig(1M)</code> utilities. This functionality is in the form of a Java Bean, which is provided by the public module vendor. This Java Bean collects public module-specific configuration from the user (you) and provides it to the Solaris DHCP service.</p> <p>The Solaris DHCP service bundles three modules with the service, which are described below. There are three <code>dhcpsvc.conf(4)</code> DHCP service configuration parameters pertaining to public modules: <code>RESOURCE</code>, <code>PATH</code>, and <code>RESOURCE_CONFIG</code>. See <code>dhcpsvc.conf(4)</code> for more information about these parameters.</p>
SUNWfiles	<p>This module stores its data in ASCII files. Although the format is ASCII, hand-editing is discouraged. It is useful for DHCP service environments that support several hundred to a couple thousand of clients and lease times are a few hours or more.</p> <p>This module's data may be shared between DHCP servers through the use of NFS.</p>
SUNWbinfiles	<p>This module stores its data in binary files. It is useful for DHCP service environments with many networks and many thousands of clients. This module provides an order of magnitude increase in performance and capacity over <code>SUNWfiles</code>.</p> <p>This module's data cannot be shared between DHCP servers.</p>
SUNWnisplus	<p>This module stores its data within a NIS+ domain. It is useful in environments where NIS+ is already deployed and facilitates sharing among multiple DHCP servers. This module supports several hundred to a few thousand clients with lease times of several hours or more.</p> <p>The NIS+ service should be hosted on a machine with ample CPU power, memory, and disk space, as the load on NIS+ is significant when it is used to store DHCP data. Periodic checkpointing of the NIS+ service is necessary in order to roll the transaction logs and keep the NIS+ service operating at its highest efficiency. See <code>nisping(1M)</code> and <code>crontab(1)</code> for more information.</p>

dhcp_modules(5)

SEE ALSO | crontab(1), dhcpconfig(1M), dhcprm(1M), dhtadm(1M), nisping(1M),
pntadm(1M), dhcpsvc.conf(4), dhcp(5)

Solaris DHCP Service Developer's Guide

NAME	environ – user environment
DESCRIPTION	<p>When a process begins execution, one of the <code>exec</code> family of functions makes available an array of strings called the environment; see <code>exec(2)</code>. By convention, these strings have the form <i>variable=value</i>, for example, <code>PATH=/sbin:/usr/sbin</code>. These environmental variables provide a way to make information about a program's environment available to programs.</p> <p>A name may be placed in the environment by the <code>export</code> command and <i>name=value</i> arguments in <code>sh(1)</code>, or by one of the <code>exec</code> functions. It is unwise to conflict with certain shell variables such as <code>MAIL</code>, <code>PS1</code>, <code>PS2</code>, and <code>IFS</code> that are frequently exported by <code>.profile</code> files; see <code>profile(4)</code>.</p> <p>The following environmental variables can be used by applications and are expected to be set in the target run-time environment.</p> <p>HOME The name of the user's login directory, set by <code>login(1)</code> from the password file; see <code>passwd(4)</code>.</p> <p>LANG The string used to specify internationalization information that allows users to work with different national conventions. The <code>setlocale(3C)</code> function checks the <code>LANG</code> environment variable when it is called with "" as the <code>locale</code> argument. <code>LANG</code> is used as the default locale if the corresponding environment variable for a particular category is unset or null. If, however, <code>LC_ALL</code> is set to a valid, non-empty value, its contents are used to override both the <code>LANG</code> and the other <code>LC_*</code> variables. For example, when invoked as <code>setlocale(LC_CTYPE, "")</code>, <code>setlocale()</code> will query the <code>LC_CTYPE</code> environment variable first to see if it is set and non-null. If <code>LC_CTYPE</code> is not set or null, then <code>setlocale()</code> will check the <code>LANG</code> environment variable to see if it is set and non-null. If both <code>LANG</code> and <code>LC_CTYPE</code> are unset or <code>NULL</code>, the default "C" locale will be used to set the <code>LC_CTYPE</code> category.</p> <p>Most commands will invoke <code>setlocale(LC_ALL, "")</code> prior to any other processing. This allows the command to be used with different national conventions by setting the appropriate environment variables.</p> <p>The following environment variables correspond to each category of <code>setlocale(3C)</code>:</p> <p>LC_ALL If set to a valid, non-empty string value, override the values of <code>LANG</code> and all the other <code>LC_*</code> variables.</p> <p>LC_COLLATE This category specifies the character collation sequence being used. The information corresponding to this category is stored in a database created by the <code>localedef(1)</code> command. This environment variable affects <code>strcoll(3C)</code> and <code>strxfrm(3C)</code>.</p>

environ(5)

LC_CTYPE

This category specifies character classification, character conversion, and widths of multibyte characters. When `LC_CTYPE` is set to a valid value, the calling utility can display and handle text and file names containing valid characters for that locale; Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide; and EUC characters of 1, 2, or 3 column widths. The default "C" locale corresponds to the 7-bit ASCII character set; only characters from ISO 8859-1 are valid. The information corresponding to this category is stored in a database created by the `localedef()` command. This environment variable is used by `ctype(3C)`, `mblen(3C)`, and many commands, such as `cat(1)`, `ed(1)`, `ls(1)`, and `vi(1)`.

LC_MESSAGES

This category specifies the language of the message database being used. For example, an application may have one message database with French messages, and another database with German messages. Message databases are created by the `mkmsgs(1)` command. This environment variable is used by `exstr(1)`, `gettext(1)`, `srchtxt(1)`, `gettext(3C)`, and `gettext(3C)`.

LC_MONETARY

This category specifies the monetary symbols and delimiters used for a particular locale. The information corresponding to this category is stored in a database created by the `localedef(1)` command. This environment variable is used by `localeconv(3C)`.

LC_NUMERIC

This category specifies the decimal and thousands delimiters. The information corresponding to this category is stored in a database created by the `localedef()` command. The default C locale corresponds to "." as the decimal delimiter and no thousands delimiter. This environment variable is used by `localeconv(3C)`, `printf(3C)`, and `strtod(3C)`.

LC_TIME

This category specifies date and time formats. The information corresponding to this category is stored in a database specified in `localedef()`. The default C locale corresponds to U.S. date and time formats. This environment variable is used by many commands and functions; for example: `at(1)`, `calendar(1)`, `date(1)`, `strftime(3C)`, and `getdate(3C)`.

MSGVERB

Controls which standard format message components `fmtmsg` selects when messages are displayed to `stderr`; see `fmtmsg(1)` and `fmtmsg(3C)`.

NETPATH

A colon-separated list of network identifiers. A network identifier is a character string used by the Network Selection component of the system to provide application-specific default network search paths. A network identifier must consist of non-null characters and must have a length of at least 1. No maximum length is specified. Network identifiers are normally chosen by the system administrator. A network identifier is also the first field in any `/etc/netconfig` file entry. `NETPATH` thus provides a link into the `/etc/netconfig` file and the information

about a network contained in that network's entry. `/etc/netconfig` is maintained by the system administrator. The library routines described in `getnetpath(3NSL)` access the `NETPATH` environment variable.

NLSPATH

Contains a sequence of templates which `catopen(3C)` and `gettext(3C)` use when attempting to locate message catalogs. Each template consists of an optional prefix, one or more substitution fields, a filename and an optional suffix. For example:

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

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

Substitution fields consist of a `%` symbol, followed by a single-letter keyword. The following keywords are currently defined:

- `%N` The value of the *name* parameter passed to `catopen()`.
- `%L` The value of `LANG` or `LC_MESSAGES`.
- `%l` The language element from `LANG` or `LC_MESSAGES`.
- `%t` The territory element from `LANG` or `LC_MESSAGES`.
- `%c` The codeset element from `LANG` or `LC_MESSAGES`.
- `%%` A single `%` character.

An empty string is substituted if the specified value is not currently defined. The separators `"_"` and `"."` are not included in `%t` and `%c` substitutions.

Templates defined in `NLSPATH` are separated by colons (`:`). A leading colon or two adjacent colons (`::`) is equivalent to specifying `%N`. For example:

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

indicates to `catopen()` that it should look for the requested message catalog in *name*, *name.cat* and `/nlslib/$LANG/name.cat`. For `gettext()`, `%N` automatically maps to "messages".

If `NLSPATH` is unset or `NULL`, `catopen()` and `gettext()` call `setlocale(3C)`, which checks `LANG` and the `LC_*` variables to locate the message catalogs.

`NLSPATH` will normally be set up on a system wide basis (in `/etc/profile`) and thus makes the location and naming conventions associated with message catalogs transparent to both programs and users.

PATH

The sequence of directory prefixes that `sh(1)`, `time(1)`, `nice(1)`, `nohup(1)`, and other utilities apply in searching for a file known by an incomplete path name. The prefixes are separated by colons (`:`). `login(1)` sets `PATH=/usr/bin`. For more detail, see `sh(1)`.

environ(5)

SEV_LEVEL

Define severity levels and associate and print strings with them in standard format error messages; see `addseverity(3C)`, `fmtmsg(1)`, and `fmtmsg(3C)`.

TERM

The kind of terminal for which output is to be prepared. This information is used by commands, such as `vi(1)`, which may exploit special capabilities of that terminal.

TZ

Timezone information. The contents of this environment variable are used by the functions `ctime(3C)`, `localtime(3C)`, `strftime(3C)`, and `mktime(3C)` to override the default timezone. If TZ is not in the following form, it designates a path to a timezone database file relative to `/usr/share/lib/zoneinfo/`, ignoring the first character if it is a colon (:). Otherwise, TZ has the form:

```
stdoffset [dst [offset] [, start [/time] , end [/time] ]]
```

std and *dst*

Three or more bytes that are the designation for the standard (*std*) and daylight savings time (*dst*) timezones. Only *std* is required. If *dst* is missing, then daylight savings time does not apply in this locale. Upper- and lower-case letters from the portable character set are explicitly allowed. Any graphic characters from the portable character set except a leading colon (:) or digits, the comma (,), the minus (-), the plus (+), and the null character are permitted to appear in these fields, but their meaning is unspecified.

offset

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*) is required and may be a single digit. The *offset* following *std* is required. If no *offset* follows *dst*, daylight savings 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 must be between 0 and 24, and the minutes (and seconds), if present, must be between 0 and 59. Out of range values may cause unpredictable behavior. If preceded by a “-”, the timezone is east of the Prime Meridian. Otherwise, it is west of the Prime Meridian (which may be indicated by an optional preceding “+” sign).

start/time, end/time

Indicate when to change to and back from daylight savings time, where *start/time* describes when the change from standard time to daylight savings time occurs, and *end/time* describes when the change back happens. Each *time* field describes when, in current local time, the change is made.

The formats of *start* and *end* are one of the following:

- Jn The Julian day n ($1 \leq n \leq 365$). Leap days are not counted. That is, in all years, February 28 is day 59 and March 1 is day 60. It is impossible to refer to the occasional February 29.
- n The zero-based Julian day ($0 \leq n \leq 365$). Leap days are counted, and it is possible to refer to February 29.
- $Mm.n.d$ The d^{th} day, ($0 \leq d \leq 6$) of week n of month m of the year ($1 \leq n \leq 5$, $1 \leq m \leq 12$), where week 5 means “the last d -day in month m ” which may occur in either the fourth or the fifth week). Week 1 is the first week in which the d^{th} day occurs. Day zero is Sunday.

Implementation specific defaults are used for *start* and *end* if these optional fields are not given.

The *time* has the same format as *offset* except that no leading sign (“-” or “+” is allowed. The default, if *time* is not given is 02:00:00.

SEE ALSO

cat(1), date(1), ed(1), fmtmsg(1), localedef(1), login(1), ls(1), mkmsgs(1), nice(1), nohup(1), sh(1), sort(1), time(1), vi(1), exec(2), addseverity(3C), catopen(3C), ctime(3C), ctype(3C), fmtmsg(3C), getdate(3C), getnetpath(3NSL), gettext(3C), gettxt(3C), localeconv(3C), mblen(3C), mktime(3C), printf(3C), setlocale(3C), strcoll(3C), strftime(3C), strtod(3C), strxfrm(3C), TIMEZONE(4), netconfig(4), passwd(4), profile(4)

eqnchar(5)

NAME eqnchar – special character definitions for eqn

SYNOPSIS eqn /usr/share/lib/pub/eqnchar *filename* | troff *options*
 neqn /usr/share/lib/pub/eqnchar *filename* | troff *options*

DESCRIPTION The eqnchar command contains nroff(1) and troff(1) character definitions for constructing characters that are not available on the Graphic Systems typesetter. These definitions are primarily intended for use with eqn(1) and neqn(1). It contains definitions for the following characters:

<i>ciplus</i>	\oplus	$ $	$ $	<i>square</i>	\square
<i>citimes</i>	\otimes	<i>langle</i>	\langle	<i>circle</i>	\circ
<i>wig</i>	\sim	<i>rangle</i>	\rangle	<i>blot</i>	\blacksquare
<i>-wig</i>	\approx	<i>hbar</i>	\hbar	<i>bullet</i>	\bullet
<i>>wig</i>	\gtrsim	<i>ppd</i>	\perp	<i>prop</i>	\propto
<i><wig</i>	\lesssim	<i><-></i>	\leftrightarrow	<i>empty</i>	\emptyset
<i>=wig</i>	\equiv	<i><=></i>	\Leftrightarrow	<i>member</i>	\in
<i>star</i>	$*$	<i> <</i>	\leftarrow	<i>nomem</i>	\notin
<i>bigstar</i>	\ast	<i> ></i>	\rightarrow	<i>cup</i>	\cup
<i>=dot</i>	\doteq	<i>ang</i>	\sphericalangle	<i>cap</i>	\cap
<i>orsign</i>	\vee	<i>rang</i>	\sphericalcap	<i>incl</i>	\subseteq
<i>andsign</i>	\wedge	<i>3dot</i>	\vdots	<i>subset</i>	\subset
<i>=del</i>	$\overset{\Delta}{=}$	<i>thf</i>	\therefore	<i>supset</i>	\supset
<i>oppA</i>	∇	<i>quarter</i>	$\frac{1}{4}$	<i>!subset</i>	$\not\subset$
<i>oppE</i>	\equiv	<i>3quarter</i>	$\frac{3}{4}$	<i>!supset</i>	$\not\supset$
<i>angstrom</i>	\AA	<i>degree</i>	$^{\circ}$		

FILES /usr/share/lib/pub/eqnchar

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWdoc

SEE ALSO eqn(1), nroff(1), troff(1), attributes(5)

NAME	extensions – localedef extensions description file
DESCRIPTION	<p>A localedef extensions description file or <i>extensions</i> file defines various extensions for the <code>localedef(1)</code> command.</p> <p>The localedef extensions description file provides:</p> <ul style="list-style-type: none">■ EUC code set width information via the <code>cwidth</code> keyword: <code>cwidth bc1 : sw1, bc2 : sw2, bc3 : sw3</code> where <code>bc1</code>, <code>bc2</code>, and <code>bc3</code> indicate the number of bytes (byte count) per character for EUC codesets 1, 2, and 3, respectively. <code>sw1</code>, <code>sw2</code>, and <code>sw3</code> indicate screen width for EUC codesets 1, 2, and 3, respectively.■ Other extensions which will be documented in a future release.
SEE ALSO	<code>locale(1)</code> , <code>localedef(1)</code> , <code>environ(5)</code> , <code>locale(5)</code>

filesystem(5)

NAME	filesystem – file system organization
SYNOPSIS	/
	/usr
DESCRIPTION	<p>The file system tree is organized for administrative convenience. Distinct areas within the file system tree are provided for files that are private to one machine, files that can be shared by multiple machines of a common platform, files that can be shared by all machines, and home directories. This organization allows sharable files to be stored on one machine but accessed by many machines using a remote file access mechanism such as NFS. Grouping together similar files makes the file system tree easier to upgrade and manage.</p> <p>The file system tree consists of a root file system and a collection of mountable file systems. The <code>mount(2)</code> program attaches mountable file systems to the file system tree at mount points (directory entries) in the root file system or other previously mounted file systems. Two file systems, <code>/</code> (the root) and <code>/usr</code>, must be mounted in order to have a completely functional system. The root file system is mounted automatically by the kernel at boot time; the <code>/usr</code> file system is mounted by the system start-up script, which is run as part of the booting process.</p> <p>Certain locations, noted below, are approved installation locations for bundled Foundation Solaris software. In some cases, the approved locations for bundled software are also approved locations for add-on system software or for applications. The following descriptions make clear where the two locations differ. For example, <code>/etc</code> is the installation location for platform&hyphen;dependent configuration files that are bundled with Solaris software. The analogous location for applications is <code>/etc/opt/packageName</code>.</p> <p>In the following descriptions, <i>subsystem</i> is a category of application or system software, such as a window system (dt) or a language (java1.2)</p> <p>The following descriptions make use of the terms <i>platform</i>, <i>platform&hyphen;dependent</i>, <i>platform&hyphen;independent</i>, and <i>platform&hyphen;specific</i>. Platform refers to a machine's Instruction Set Architecture or processor type, such as is returned by <code>uname -i</code>. <i>Platform&hyphen;dependent</i> refers to a file that is installed on all platforms and whose contents vary depending on the platform. Like a platform&hyphen;dependent file, a <i>platform&hyphen;independent</i> file is installed on all platforms. However, the contents of the latter type remains the same on all platforms. An example of a platform&hyphen;dependent file is a compiled, executable program. An example of a platform&hyphen;independent file is a standard configuration file, such as <code>/etc/hosts</code>. Unlike a platform&hyphen;dependent or a platform&hyphen;independent file, the <i>platform&hyphen;specific</i> file is installed only on a subset of supported platforms. Most platform-specific files are gathered under <code>/platform</code> and <code>/usr/platform</code>.</p>
Root File System	The root file system contains files that are unique to each machine. It contains the following directories:

/
Root of the overall file system name space.

/dev
Primary location for special files. Typically, device files are built to match the kernel and hardware configuration of the machine.

/dev/cfg
Symbolic links to physical ap_ids.

/dev/cua
Device files for uucp.

/dev/dsk
Block disk devices.

/dev/fbs
Frame buffer device files.

/dev/fd
File descriptors.

/dev/md
Logical volume management meta-disk devices.

/dev/printers
USB printer device files.

/dev/pts
Pseudo-terminal devices.

/dev/rdisk
Raw disk devices.

/dev/rmt
Raw tape devices.

/dev/sad
Entry points for the STREAMS Administrative driver.

/dev/sound
Audio device and audio device control files.

/dev/swap
Default swap device.

/dev/term
Terminal devices.

/devices
Physical device files.

/etc
Platform-dependent administrative and configuration files and databases that are not shared among systems. /etc may be viewed as the directory that

filesystem(5)

defines the machine's identity. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/etc/opt/packagename`.

`/etc/acct`

Accounting system configuration information.

`/etc/apache`

Apache configuration files.

`/etc/cron.d`

Configuration information for `cron(1M)`.

`/etc/default`

Defaults information for various programs.

`/etc/dfs`

Configuration information for shared file systems.

`/etc/dhcp`

Dynamic Host Configuration Protocol (DHCP) configuration files.

`/etc/dmi`

Solstice Enterprise Agents configuration files.

`/etc/fn`

Federated Naming Service and X.500 support files.

`/etc/fs`

Binaries organized by file system types for operations required before `/usr` is mounted.

`/etc/gss`

Generic Security Service (GSS) Application Program Interface configuration files.

`/etc/gtk`

GNOME (GNU Network Object Model Environment) configuration files.

`/etc/inet`

Configuration files for Internet services.

`/etc/init.d`

Shell scripts for transitioning between run levels.

`/etc/iplanet`

iPlanet configuration files.

`/etc/krb5`

Kerberos configuration files.

`/etc/lib`

Shared libraries needed during booting.

`/etc/lp`

Configuration information for the printer subsystem.

`/etc/llc2`
Logical link control (llc2) driver configuration files.

`/etc/lp`
Configuration information for the printer subsystem.

`/etc/lu`
Solaris Live Upgrade configuration files.

`/etc/lvm`
Solaris Logical Volume Manager configuration files.

`/etc/mail`
Mail subsystem configuration.

`/etc/nca`
Solaris Network Cache and Accelerator (NCA) configuration files.

`/etc/net`
Configuration information for transport independent network services.

`/etc/nfs`
NFS server logging configuration file.

`/etc/openwin`
OpenWindows configuration files.

`/etc/opt`
Configuration information for optional packages.

`/etc/ppp`
Solaris PPP configuration files.

`/etc/rc0.d`
Scripts for entering or leaving run level 0. See `init(1M)`.

`/etc/rc1.d`
Scripts for entering or leaving run level 1. See `init(1M)`.

`/etc/rc2.d`
Scripts for entering or leaving run level 2. See `init(1M)`.

`/etc/rc3.d`
Scripts for entering or leaving run level 3. See `init(1M)`.

`/etc/rcS.d`
Scripts for bringing the system up in single user mode.

`/etc/rcm`
Directory for reconfiguration manager (RCM) custom scripts.

`/etc/rpcsec`
This directory might contain an NIS+ authentication configuration file.

`/etc/saf`
Service Access Facility files.

filesystem(5)

`/etc/security`
Basic Security Module (BSM) configuration files.

`/etc/sfw`
Samba configuration files.

`/etc/skel`
Default profile scripts for new user accounts. See `useradd(1M)`.

`/etc/smartcard`
Solaris SmartCard configuration files.

`/etc/snmp`
Solstice Enterprise Agents configuration files.

`/etc/ssh`
Secure Shell configuration files. See `ssh(1)`

`/etc/sysevent`
`syseventd` configuration files.

`/etc/subsystem`
Platform‐dependent *subsystem* configuration files that are not shared among systems. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/etc/opt/packageName`.

`/etc/tm`
Trademark files; contents displayed at boot time.

`/etc/usb`
USB configuration information.

`/etc/uucp`
UUCP configuration information. See `uucp(1C)`.

`/etc/wrsm`
WCI Remote Shared Memory (WRSM) configuration information. See `wrsmconf(1M)`

`/export`
Default root of the shared file system tree.

`/home`
Default root of a subtree for user directories.

`/kernel`
Subtree of platform‐dependent loadable kernel modules required as part of the boot process. It includes the generic part of the core kernel that is platform‐independent, `/kernel/genunix`. See `kernel(1M)` An approved installation location for bundled Solaris software and for add-on system software.

`/kernel/drv`
32-bit device drivers.

`/kernel/drv/sparcv9`
64-bit SPARC device drivers.

`/kernel/genunix`
Platform-independent kernel.

`/kernel/subsystem/ia64`
64-bit x86 platform-dependent modules required for boot. An approved installation location for bundled Solaris software and for add-on system software. Note that `ia64` is an example name; the actual name might be different.

`/kernel/subsystem/sparcv9`
64-bit SPARC platform-dependent modules required for boot. An approved installation location for bundled Solaris software and for add-on system software.

`/mnt`
Default temporary mount point for file systems. This is an empty directory on which file systems can be temporarily mounted.

`/opt`
Root of a subtree for add-on application packages.

`/platform`
Subtree of platform-specific objects which need to reside on the root filesystem. It contains a series of directories, one per supported platform. The semantics of the series of directories is equivalent to `/` (root).

`/platform/`uname -i`/kernel`
Platform-specific modules required for boot. These modules have semantics equivalent to `/kernel`. It includes the file `unix`, the core kernel. See `kernel(1M)`. An approved installation location for bundled Solaris software and for add-on system software.

`/platform/`uname -m`/kernel`
Hardware class-specific modules required for boot. An approved installation location for bundled Solaris software and for add-on system software.

`/platform/`uname -i`/kernel/subsystem/ia64`
x86 64-bit, platform-dependent modules required for boot. Note that `ia64` is an example name; the actual name might be different. An approved installation location for bundled Solaris software.

`/platform/`uname -i`/kernel/subsystem/sparcv9`
SPARC 64-bit platform-specific modules required for boot. An approved installation location for bundled Solaris software.

`/platform/`uname -i`/kernel/sparcv9/unix`
64-bit platform-dependent kernel.

`/platform/`uname -i`/kernel/unix`
32-bit platform-dependent kernel.

filesystem(5)

`/platform/`uname -i`/lib`
Platform‐specific shared objects required for boot. Semantics are equivalent to `/lib`. An approved installation location for bundled Solaris software and for add-on system software.

`/platform/`uname -i`/sbin`
Platform‐specific administrative utilities required for boot. Semantics are equivalent to `/sbin`. An approved installation location for bundled Solaris software and for add-on system software.

`/proc`
Root of a subtree for the process file system.

`/sbin`
Essential executables used in the booting process and in manual system recovery. The full complement of utilities is available only after `/usr` is mounted. `/sbin` is an approved installation location for bundled Solaris software.

`/tmp`
Temporary files; cleared during the boot operation.

`/usr`
Mount point for the `/usr` file system. See description of `/usr` file system, below.

`/var`
Root of a subtree for varying files. Varying files are files that are unique to a machine but that can grow to an arbitrary (that is, variable) size. An example is a log file. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/var/opt/packagename`.

`/var/adm`
System logging and accounting files.

`/var/apache`
Scripts, icons, logs, and cache pages for Apache web server.

`/var/audit`
Basic Security Module (BSM) audit files.

`/var/crash`
Default depository for kernel crash dumps.

`/var/cron`
Log files for `cron(1M)`.

`/var/dmi`
Solstice Enterprise Agents (SEA) Desktop Management Interface (DMI) run-time components.

`/var/dt`
`dtlogin` configuration files.

`/var/ftp`
FTP server directory.

`/var/inet`
IPv6 router state files.

`/var/krb5`
Database and log files for Kerberos.

`/var/ld`
Configuration files for runtime linker.

`/var/ldap`
LDAP client configuration files.

`/var/log`
System log files.

`/var/lp`
Line printer subsystem logging information.

`/var/mail`
Directory where users' mail is kept.

`/var/news`
Community service messages. This is not the same as USENET-style news.

`/var/nfs`
NFS server log files.

`/var/nis`
NIS+ databases.

`/var/ntp`
Network Time Protocol (NTP) server state directory.

`/var/opt`
Root of a subtree for varying files associated with optional software packages. An approved installation location for add-on system software and applications.

`/var/preserve`
Backup files for `vi(1)` and `ex(1)`.

`/var/run`
Temporary files which are not needed across reboots. Only root may modify the contents of this directory.

`/var/sadm`
Databases maintained by the software package management utilities.

`/var/sadm/system/logs`
Status log files produced by software management functions and/or applications. For example, log files produced for product installation. An approved installation location for bundled Solaris software and for add-on system software and applications.

filesystem(5)

`/var/saf`
Service access facility logging and accounting files.

`/var/samba`
Log and lock files for Samba.

`/var/snmp`
SNMP status and configuration information.

`/var/spool`
Contains directories for files used in printer spooling, mail delivery, `cron(1M)`, `at(1)`, and so forth.

`/var/spool/clientmqueue`
sendmail(1M) client files.

`/var/spool/cron`
cron(1M) and at(1) spooling files.

`/var/spool/locks`
Spooling lock files.

`/var/spool/lp`
Line printer spool files. See `lp(1)`.

`/var/spool/mqueue`
Mail queued for delivery.

`/var/spool/pkg`
Spooled packages.

`/var/spool/print`
LP print service client-side request staging area.

`/var/spool/samba`
Samba print queue.

`/var/spool/uucp`
Queued uucp(1C) jobs.

`/var/spool/uucppublic`
Files deposited by uucp(1C).

`/var/statmon`
Network status monitor files.

`/var/tmp`
Files that vary in size or presence during normal system operations. This directory is *not* cleared during the boot operation. An approved installation location for bundled Solaris software and for add-on system software and applications.

`/var/uucp`
uucp(1C) log and status files.

/usr File System	<p><code>/var/yp</code> Databases needed for backwards compatibility with NIS and <code>ypbind(1M)</code>; unnecessary after full transition to NIS+.</p> <p>Because it is desirable to keep the root file system small and not volatile, on disk-based systems larger file systems are often mounted on <code>/home</code>, <code>/opt</code>, <code>/usr</code>, and <code>/var</code>.</p> <p>The file system mounted on <code>/usr</code> contains platform-dependent and platform-independent sharable files. The subtree rooted at <code>/usr/share</code> contains platform-independent sharable files; the rest of the <code>/usr</code> tree contains platform-dependent files. By mounting a common remote file system, a group of machines with a common platform may share a single <code>/usr</code> file system. A single <code>/usr/share</code> file system can be shared by machines of any platform. A machine acting as a file server can share many different <code>/usr</code> file systems to support several different architectures and operating system releases. Clients usually mount <code>/usr</code> read-only so that they do not accidentally change any shared files.</p> <p>The <code>/usr</code> file system contains the following subdirectories:</p> <p><code>/usr/4lib</code> a.out libraries for the Binary Compatibility Package.</p> <p><code>/usr/5bin</code> Symbolic link to the <code>/usr/bin</code> directory.</p> <p><code>/usr/x</code> Symbolic link to the <code>/usr/openwin</code> directory.</p> <p><code>/usr/adm</code> Symbolic link to the <code>/var/adm</code> directory.</p> <p><code>/usr/apache</code> Apache executables, loadable modules, and documentation.</p> <p><code>/usr/aset</code> Directory for Automated Security Enhancement Tools (ASET) programs and files.</p> <p><code>/usr/bin</code> Platform-dependent, user-invoked executables. These are commands users expect to be run as part of their normal <code>\$PATH</code>. For executables that are different on a 64-bit system than on a 32-bit system, a wrapper that selects the appropriate executable is placed here. See <code>isaexec(3C)</code>. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is <code>/opt/package/bin</code>.</p> <p><code>/usr/bin/ia64</code> x86 64-bit, platform-dependent, user-invoked executables. Note that <code>ia64</code> is an example name; the actual name might be different. This directory should not be part of a user's <code>\$PATH</code>. A wrapper in <code>/usr/bin</code> should invoke the executable in this directory. See <code>isaexec(3C)</code>. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is <code>/opt/package/bin/ia64</code>.</p>
-------------------------	---

filesystem(5)

`/usr/bin/sparcv9`

SPARC 64-bit, platform‐dependent, user-invoked executables. This directory should not be part of a user's `$PATH`. A wrapper in `/usr/bin` should invoke the executable in this directory. See `isaexec(3C)`. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/bin/sparcv9`.

`/usr/bin/subsystem`

Platform‐dependent user-invoked executables that are associated with *subsystem*. These are commands users expect to be run as part of their normal `$PATH`. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/bin`.

`/usr/bin/subsystem/ia64`

x86 64-bit, platform‐dependent, user-invoked executables. Note that `ia64` is an example name; the actual name might be different. This directory should not be part of a user's `$PATH`. A wrapper in `/usr/bin` should invoke the executable in this directory. See `isaexec(3C)`. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/bin/ia64`.

`/usr/bin/subsystem/sparcv9`

SPARC 64-bit, platform‐dependent, user-invoked executables. This directory should not be part of a user's `$PATH`. A wrapper in `/usr/bin` should invoke the executable in this directory. See `isaexec(3C)`. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/bin/sparcv9`.

`/usr/subsystem/bin`

Platform‐dependent user-invoked executables that are associated with *subsystem*. These are commands users expect to be run as part of their normal `$PATH`. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/bin`.

`/usr/subsystem/bin/ia64`

x86 64-bit, platform‐dependent, user-invoked executables. Note that `ia64` is an example name; the actual name might be different. This directory should not be part of a user's `$PATH`. A wrapper in `/usr/bin` should invoke the executable in this directory. See `isaexec(3C)`. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/bin/ia64`.

`/usr/subsystem/bin/sparcv9`

SPARC 64-bit, platform‐dependent, user-invoked executables. This directory should not be part of a user's `$PATH`. A wrapper in `/usr/bin` should invoke the executable in this directory. See `isaexec(3C)`. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/bin/sparcv9`.

`/usr/ccs`
C compilation system.

`/usr/ccs/bin`
C compilation commands and system utilities.

`/usr/ccs/lib`
Symbolic link to `/usr/lib`.

`/usr/demo`
Demo programs and data.

`/usr/dict`
Symbolic link to the `/usr/share/lib/dict` directory, which contains the dictionary file used by the UNIX spell program.

`/usr/dt`
root of a subtree for CDE software.

`/usr/dt/bin`
Primary location for CDE system utilities.

`/usr/dt/include`
Header files for CDE software.

`/usr/dt/lib`
Libraries for CDE software.

`/usr/dt/share/man`
On-line reference manual pages for CDE software.

`/usr/games`
An empty directory, a remnant of the SunOS 4.0/4.1 software.

`/usr/include`
Include headers (for C programs).

`/usr/iplanet`
Directory server executables, loadable modules, and documentation.

`/usr/j2se`
Java 2 SDK executables, loadable modules, and documentation.

`/usr/java*`
Directories containing Java programs and libraries.

`/usr/kernel`
Subtree of platform-dependent loadable kernel modules, not needed in the root filesystem. An approved installation location for bundled Solaris software.

`/usr/kvm`
A mount point, retained for backward compatibility, that formerly contained platform-specific binaries and libraries.

filesystem(5)

`/usr/lib`
Platform‐dependent libraries, various databases, commands and daemons not invoked directly by a human user. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/packagename/lib`.

`/usr/lib/64`
Symbolic link to the most portable 64-bit Solaris interfaces.

`/usr/lib/acct`
Accounting scripts and binaries. See `acct(1M)`.

`/usr/lib/class`
Scheduling‐class-specific directories containing executables for `priocntl(1)` and `dispadmin(1M)`.

`/usr/lib/dict`
Database files for `spell(1)`.

`/usr/lib/font`
`troff(1)` font description files.

`/usr/lib/fs`
File system type dependent modules; generally not intended to be invoked directly by the user.

`/usr/lib/ia64`
x86 64-bit, platform‐dependent libraries, various databases, commands and daemons not invoked directly by a human user. Note that `ia64` is an example name; the actual name might be different. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/packagename/lib/ia64`.

`/usr/lib/iconv`
Conversion tables for `iconv(1)`.

`/usr/lib/libp`
Profiled libraries.

`/usr/lib/locale`
Localization databases.

`/usr/lib/lp`
Line printer subsystem databases and back-end executables.

`/usr/lib/mail`
Auxiliary programs for the `mail(1)` subsystem.

`/usr/lib/netsvc`
Internet network services.

`/usr/lib/nfs`
Auxiliary NFS-related programs and daemons.

`/usr/lib/pics`
 Position Independent Code (PIC) archives needed to rebuild the run-time linker.

`/usr/lib/refer`
 Auxiliary programs for `refer(1)`.

`/usr/lib/sa`
 Scripts and commands for the system activity report package. See `sar(1)`.

`/usr/lib/saf`
 Auxiliary programs and daemons related to the service access facility.

`/usr/lib/sparcv9`
 SPARC 64-bit, platform‐dependent libraries, various databases, commands and daemons not invoked directly by a human user. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/lib/sparcv9`.

`/usr/lib/spell`
 Auxiliary programs and databases for `spell(1)`. This directory is only present when the Binary Compatibility Package is installed.

`/usr/lib/uucp`
 Auxiliary programs and daemons for `uucp(1C)`.

`/usr/lib/subsystem`
 Platform‐dependent libraries, various databases, commands and daemons that are associated with `subsystem` and that are not invoked directly by a human user. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/lib`.

`/usr/lib/subsystem/ia64`
 x86 64-bit, platform‐dependent libraries, various databases, commands and daemons that are associated with `subsystem` and that are not invoked directly by a human user. Note that `ia64` is an example name; the actual name might be different. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/lib/ia64`.

`/usr/lib/subsystem/sparcv9`
 SPARC 64-bit, platform‐dependent libraries, various databases, commands and daemons that are associated with `subsystem` and that are not invoked directly by a human user. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/lib/sparcv9`.

`/usr/subsystem/lib`
 Platform‐dependent libraries, various databases, commands and daemons not invoked directly by a human user. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/lib`.

filesystem(5)

`/usr/subsystem/lib/ia64`
x86 64-bit, platform‐dependent libraries, various databases, commands and daemons that are associated with *subsystem* and that are not invoked directly by a human user. Note that *ia64* is an example name; the actual name might be different. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/ia64`.

`/usr/subsystem/lib/sparcv9`
SPARC 64-bit, platform‐dependent libraries, various databases, commands and daemons that are associated with *subsystem* and that are not invoked directly by a human user. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/sparcv9`.

`/usr/local`
Not part of the SVR4-based Solaris distribution. The `/usr` directory is exclusively for software bundled with the Solaris operating system. If needed for storing machine-local add-on software, create the directory `/opt/local` and make `/usr/local` a symbolic link to `/opt/local`. The `/opt` directory or filesystem is for storing add-on software to the system.

`/usr/mail`
Symbolic link to the `/var/mail` directory.

`/usr/man`
Symbolic link to the `/usr/share/man` directory.

`/usr/net/servers`
Entry points for foreign name service requests relayed using the network listener. See `listen(1M)`.

`/usr/news`
Symbolic link to the `/var/news` directory.

`/usr/oasys`
Commands and files related to the Form and Menu Language Interpreter (FMLI) execution environment. See `face(1)`.

`/usr/old`
Programs that are being phased out.

`/usr/openwin`
Installation or mount point for the OpenWindows software.

`/usr/perl5`
Perl 5 programs and documentation

`/usr/platform`
Subtree of platform‐specific objects which does not need to reside on the root filesystem. It contains a series of directories, one per supported platform. The

semantics of the series of directories is equivalent to `/platform`, except for subdirectories which do not provide utility under one or the other (for example, `/platform/include` is not needed).

`/usr/platform/`uname -i`/include`

Symbolic link to `../`uname -i`/include`. Platform-specific system (`sys`, `vm`) header files with semantics equivalent to `/usr/include`. An approved installation location for bundled Solaris software and for add-on system software.

`/usr/platform/`uname -i`/kernel`

Platform-specific modules with semantics equivalent to `/usr/kernel`. An approved installation location for bundled Solaris software and for add-on system software.

`/usr/platform/`uname -i`/lib`

Platform-specific daemon and shared objects with semantics equivalent to `/usr/lib`. An approved installation location for bundled Solaris software and for add-on system software.

`/usr/platform/`uname -i`/lib/ia64`

x86 64-bit, platform-specific daemon and shared objects. Note that `ia64` is an example name; the actual name might be different. An approved installation location for bundled Solaris software and for add-on system software.

`/usr/platform/`uname -i`/lib/sparcv9`

SPARC 64-bit, platform-specific daemon and shared objects. An approved installation location for bundled Solaris software and for add-on system software.

`/usr/platform/`uname -i`/[s]mannum`

Where *num* can be one of `3x`, `1m`, `4`, `7d`, or `9e`. Platform-specific system manual pages for documenting platform-specific, shared objects, administration utilities, configuration files, special files/modules, and header files. An approved installation location for bundled Solaris software and for add-on system software.

`/usr/platform/`uname -i`/sbin`

Platform-specific system administration utilities with semantics equivalent to `/usr/sbin`. An approved installation location for bundled Solaris software and for add-on system software.

`/usr/preserve`

Symbolic link to the `/var/preserve` directory.

`/usr/proc`

Directory for the `proc` tools.

`/usr/proc/bin`

Contains links to SPARC Version 8 binaries in `/usr/bin`.

`/usr/pub`

Files for online man page and character processing.

`/usr/sadm`

System administration files and directories.

filesystem(5)

`/usr/sadm/bin`
Binaries for the Form and Menu Language Interpreter (FMLI) scripts. See `fml i(1)`.

`/usr/sadm/install`
Executables and scripts for package management.

`/usr/sbin`
Platform‐dependent executables for system administration, expected to be run only by system administrators. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/sbin`.

`/usr/sbin/install.d`
Custom Jumpstart scripts and executables.

`/usr/sbin/static`
Statically linked version of selected programs from `/usr/bin` and `/usr/sbin`. These are used to recover from broken dynamic linking and before all pieces necessary for dynamic linking are present.

`/usr/sbin/sparc7` and `sparc9`
32-bit and 64-bit versions of commands.

`/usr/sfw`
GNU and open source executables, libraries, and documentation.

`/usr/sbin/subsystem`
Platform‐dependent executables for system administration, expected to be run only by system administrators, and associated with *subsystem*. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/sbin`.

`/usr/subsystem/sbin`
Platform‐dependent executables for system administration, expected to be run only by system administrators, and associated with *subsystem*. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/package/sbin`.

`/usr/share`
Platform‐independent sharable files. An approved installation location for bundled Solaris software.

`/usr/share/admserv5.1`
iPlanet Console and Administration Server documentation.

`/usr/share/audio`
Sample audio files.

`/usr/share/ds5`
iPlanet Server documentation.

`/usr/share/lib`
Platform‐independent sharable databases. An approved installation location for bundled Solaris software.

`/usr/share/lib/dict`
Contains word list for `spell(1)`.

`/usr/share/lib/keytables`
Keyboard layout description tables.

`/usr/share/lib/mailx`
Help files for `mailx(1)`.

`/usr/share/lib/nterm`
`nroff(1)` terminal tables.

`/usr/share/lib/pub`
Character set data files.

`/usr/share/lib/tabset`
Tab setting escape sequences.

`/usr/share/lib/terminfo`
Terminal description files for `terminfo(4)`.

`/usr/share/lib/tmac`
Macro packages and related files for text processing tools, for example, `nroff(1)` and `troff(1)`.

`/usr/share/lib/zoneinfo`
Time zone information.

`/usr/share/[s]man`
Platform‐independent sharable manual pages. An approved installation location for bundled Solaris software. The analogous location for add-on system software or for applications is `/opt/packageName/[s]man`.

`/usr/share/src`
Source code for kernel, utilities, and libraries.

`/usr/snadm`
Files related to system and network administration.

`/usr/spool`
Symbolic link to the `/var/spool` directory.

`/usr/src`
Symbolic link to the `/usr/share/src` directory.

`/usr/tmp`
Symbolic link to the `var/tmp` directory.

`/usr/ucb`
Berkeley compatibility package binaries.

`/usr/ucbinclude`
Berkeley compatibility package headers.

`/usr/ucb/lib`
Berkeley compatibility package libraries.

filesystem(5)

/usr/vmsys

Commands and files related to the Framed Access Command Environment (FACE) programs. See [face\(1\)](#).

/usr/xpg4

Directory for POSIX-compliant utilities.

SEE ALSO [at\(1\)](#), [ex\(1\)](#), [face\(1\)](#), [fmli\(1\)](#), [iconv\(1\)](#), [lp\(1\)](#), [isainfo\(1\)](#), [mail\(1\)](#), [mailx\(1\)](#), [nroff\(1\)](#), [priocntl\(1\)](#), [refer\(1\)](#), [sar\(1\)](#), [sh\(1\)](#), [spell\(1\)](#), [troff\(1\)](#), [uname\(1\)](#), [uucp\(1C\)](#), [vi\(1\)](#), [acct\(1M\)](#), [cron\(1M\)](#), [dispadmin\(1M\)](#), [fsck\(1M\)](#), [init\(1M\)](#), [kernel\(1M\)](#), [mknod\(1M\)](#), [mount\(1M\)](#), [useradd\(1M\)](#), [ypbind\(1M\)](#), [mount\(2\)](#), [intro\(4\)](#), [terminfo\(4\)](#)

NAME	fnmatch – file name pattern matching
DESCRIPTION	<p>The pattern matching notation described below is used to specify patterns for matching strings in the shell. Historically, pattern matching notation is related to, but slightly different from, the regular expression notation. For this reason, the description of the rules for this pattern matching notation is based on the description of regular expression notation described on the <code>regex(5)</code> manual page.</p>
Patterns Matching a Single Character	<p>The following <i>patterns matching a single character</i> match a single character: <i>ordinary characters</i>, <i>special pattern characters</i> and <i>pattern bracket expressions</i>. The pattern bracket expression will also match a single collating element.</p> <p>An ordinary character is a pattern that matches itself. It can be any character in the supported character set except for NUL, those special shell characters that require quoting, and the following three special pattern characters. Matching is based on the bit pattern used for encoding the character, not on the graphic representation of the character. If any character (ordinary, shell special, or pattern special) is quoted, that pattern will match the character itself. The shell special characters always require quoting.</p> <p>When unquoted and outside a bracket expression, the following three characters will have special meaning in the specification of patterns:</p> <ul style="list-style-type: none"> ? A question-mark is a pattern that will match any character. * An asterisk is a pattern that will match multiple characters, as described in <i>Patterns Matching Multiple Characters</i>, below. [The open bracket will introduce a pattern bracket expression. <p>The description of basic regular expression bracket expressions on the <code>regex(5)</code> manual page also applies to the pattern bracket expression, except that the exclamation-mark character (!) replaces the circumflex character (^) in its role in a <i>non-matching list</i> in the regular expression notation. A bracket expression starting with an unquoted circumflex character produces unspecified results.</p> <p>The restriction on a circumflex in a bracket expression is to allow implementations that support pattern matching using the circumflex as the negation character in addition to the exclamation-mark. A portable application must use something like <code>[\^!]</code> to match either character.</p> <p>When pattern matching is used where shell quote removal is not performed (such as in the argument to the <code>find -name</code> primary when <code>find</code> is being called using one of the <code>exec</code> functions, or in the <i>pattern</i> argument to the <code>fnmatch(3C)</code> function, special characters can be escaped to remove their special meaning by preceding them with a backslash character. This escaping backslash will be discarded. The sequence <code>\\</code> represents one literal backslash. All of the requirements and effects of quoting on ordinary, shell special and special pattern characters will apply to escaping in this context.</p>

fnmatch(5)

Both quoting and escaping are described here because pattern matching must work in three separate circumstances:

- Calling directly upon the shell, such as in pathname expansion or in a case statement. All of the following will match the string or file abc:

abc	"abc"	a"b"c	a\bc	a[b]c
a["b"]c	a[\b]c	a["\b"]c	a?c	a*c

The following will not:

"a?c"	a*c	a\[b]c
-------	-----	--------

- Calling a utility or function without going through a shell, as described for `find(1)` and the function `fnmatch(3C)`
- Calling utilities such as `find`, `cpio`, `tar` or `pax` through the shell command line. In this case, shell quote removal is performed before the utility sees the argument. For example, in:

`find /bin -name e\c[\h]o -print` after quote removal, the backslashes are presented to `find` and it treats them as escape characters. Both precede ordinary characters, so the `c` and `h` represent themselves and `echo` would be found on many historical systems (that have it in `/bin`). To find a file name that contained shell special characters or pattern characters, both quoting and escaping are required, such as:

```
pax -r . . . "*a\ ( \?"to extract a filename ending with a (?.
```

Conforming applications are required to quote or escape the shell special characters (sometimes called metacharacters). If used without this protection, syntax errors can result or implementation extensions can be triggered. For example, the KornShell supports a series of extensions based on parentheses in patterns; see `ksh(1)`

Patterns Matching Multiple Characters

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

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

Since each asterisk matches zero or more occurrences, the patterns `a*b` and `a**b` have identical functionality.

Examples:

<code>a [bc]</code>	matches the strings <code>ab</code> and <code>ac</code> .
<code>a*d</code>	matches the strings <code>ad</code> , <code>abd</code> and <code>abcd</code> , but not the string <code>abc</code> .
<code>a*d*</code>	matches the strings <code>ad</code> , <code>abcd</code> , <code>abcdef</code> , <code>aaaad</code> and <code>adddd</code> .
<code>*a*d</code>	matches the strings <code>ad</code> , <code>abcd</code> , <code>efabcd</code> , <code>aaaad</code> and <code>adddd</code> .

Patterns Used for Filename Expansion

The rules described so far in *Patterns Matching Multiple Characters* and *Patterns Matching a Single Character* are qualified by the following rules that apply when pattern matching notation is used for filename expansion.

1. The slash character in a pathname must be explicitly matched by using one or more slashes in the pattern; it cannot be matched by the asterisk or question-mark special characters or by a bracket expression. Slashes in the pattern are identified before bracket expressions; thus, a slash cannot be included in a pattern bracket expression used for filename expansion. For example, the pattern `a [b/c] d` will not match such pathnames as `abd` or `a/d`. It will only match a pathname of literally `a [b/c] d`.
2. If a filename begins with a period (`.`), the period must be explicitly matched by using a period as the first character of the pattern or immediately following a slash character. The leading period will not be matched by:
 - the asterisk or question-mark special characters
 - a bracket expression containing a non-matching list, such as:

`[!a]` a range expression, such as:

`[%-0]` or a character class expression, such as:

`[[:punct:]]` It is unspecified whether an explicit period in a bracket expression matching list, such as:

`[.abc]` can match a leading period in a filename.

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

`/foo/bar/x*/bam` search permission is needed for directories `/` and `foo`, search and read permissions are needed for directory `bar`, and search permission is

fnmatch(5)

needed for each `x*` directory. If the pattern matches any existing filenames or pathnames, the pattern will be replaced with those filenames and pathnames, sorted according to the collating sequence in effect in the current locale. If the pattern contains an invalid bracket expression or does not match any existing filenames or pathnames, the pattern string is left unchanged.

SEE ALSO `find(1)`, `ksh(1)`, `fnmatch(3C)`, `regex(5)`

NAME	fns – overview of FNS
DESCRIPTION	<p>Federated Naming Service (FNS) provides a method for federating multiple naming services under a single, simple interface for the basic naming operations. The service supports resolution of <i>composite</i> names, names that span multiple naming systems, through the naming interface. In addition to the naming interface, FNS also specifies <i>policies</i> for composing names in the enterprise namespace. See <code>fns_policies(5)</code> and <code>fns_initial_context(5)</code>.</p> <p>Fundamental to the FNS model are the notions of composite names and <i>contexts</i>. A context provides operations for:</p> <ul style="list-style-type: none"> ■ associating (binding) names to objects ■ resolving names to objects ■ removing bindings, listing names, renaming and so on. <p>A context contains a set of names to reference bindings. A reference contains a list of communication end-points. Every naming operation in the FNS interface is performed on a context object.</p> <p>The federated naming system is formed by contexts from one naming system being bound in the contexts of another naming system. Resolution of a composite name proceeds from contexts within one naming system to those in the next, until the name is resolved.</p>
XFN	XFN is <i>X/Open Federated Naming</i> . The programming interface and policies that FNS supports are specified by XFN. See <code>xfn(3XFN)</code> and <code>fns_policies(5)</code> .
Composite Names	<p>A composite name is a name that spans multiple naming systems. It consists of an ordered list of components. Each component is a name from the namespace of a single naming system. FNS defines the syntax for constructing a composite name using names from component naming systems. Individual naming systems are responsible for the syntax of each component.</p> <p>The syntax for composite names is that components are composed left to right using the slash character ('/') as the component separator. For example, the composite name <code>. . . /Wiz.Com/site/Oceanview.East</code> consists of four components: <code>. . .</code>, <code>Wiz.COM</code>, <code>site</code>, and <code>Oceanview.East</code>. See <code>fns_policies(5)</code> and <code>fns_initial_context(5)</code> for more examples of composite names.</p>
Why FNS?	<p>FNS is useful for the following reasons:</p> <ul style="list-style-type: none"> ■ A single uniform naming interface is provided to clients for accessing naming services. Consequently, the addition of new naming services does not require changes to applications or existing naming services. Furthermore, applications that use FNS will be portable across platforms because the interface exported by FNS is XFN, a public, open interface endorsed by other vendors and by the X/Open Company.

fns(5)

- Names can be composed in a uniform way (that is, FNS supports a model in which composite names are constructed in a uniform syntactic way and can have any number of components).
- Coherent naming is encouraged through the use of shared contexts and shared names.

FNS and Naming Systems

FNS has support for NIS+, NIS, and files as enterprise-level naming services. This means that FNS implements the enterprise-level policies using NIS+, NIS, and files. FNS also supports DNS and X.500 (via DAP or LDAP) as global naming services, as well as support for federating NIS+ and NIS with DNS and X.500. See the corresponding individual man page for information about the implementation for a specific naming service.

SEE ALSO

`nis+(1)`, `xfn(3XFN)`, `fns_dns(5)`, `fns_files(5)`, `fns_initial_context(5)`, `fns_nis(5)`, `fns_nis+(5)`, `fns_policies(5)`, `fns_references(5)`, `fns_x500(5)`

NAME	fns_dns – overview of FNS over DNS implementation
DESCRIPTION	<p>Federated Naming Service (FNS) provides a method for federating multiple naming services under a single, simple interface for the basic naming operations. One of the naming services supported by FNS is the Internet Domain Name System, or DNS (see in.named(1M)). DNS is a hierarchical collection of name servers that provide the Internet community with host and domain name resolution. FNS uses DNS to name entities globally. Names can be constructed for any enterprise that is accessible on the Internet; consequently, names can also be constructed for objects exported by these enterprises.</p> <p>FNS provides the XFN interface for performing naming resolution on DNS domains and hosts. In addition, enterprise namespaces such as those served by NIS+ and NIS can be federated with DNS by adding TXT records to DNS. To federate an NIS+ or NIS namespace under DNS, you first obtain the root reference for the NIS+ hierarchy or NIS domain. This reference is referred to as the <i>next naming system reference</i> because it refers to the <i>next</i> naming system beneath the DNS domain. This reference contains information about how to communicate with the NIS+ or NIS servers and has the following format:</p> <pre><domainname> <server name> [<server address>]</pre> <p>where <domainname> is the fully qualified domain name. Notice that NIS+ and NIS have slightly different syntaxes for domain names. For NIS+, the fully qualified domain name is case-insensitive and terminated by a dot character ('.'). For NIS, the fully qualified domain name is case-sensitive and is <i>not</i> terminated by a dot character. For both NIS+ and NIS, <server address> is optional. If it is not supplied, a host name lookup will be performed to get the machine's address.</p> <p>For example, if the machine <code>wiz-nisplus-server</code> with address <code>133.33.33.33</code> serves the NIS+ domain <code>wiz.com.</code>, the reference would look like this:</p> <pre>wiz.com. wiz-nisplus-server 133.33.33.33</pre> <p>For NIS, the reference information is of the form:</p> <pre><domainname> <server name></pre> <p>For example, if the machine <code>woz-nis-server</code> serves the NIS domain <code>Woz.COM</code>, the reference would look like this:</p> <pre>Woz.COM woz-nis-server</pre>

fns_dns(5)

After obtaining this information, you then edit the DNS table (see `in.named(1M)`) and add a TXT record with this reference information. The TXT record must be associated with a DNS domain that includes an NIS record. For example, the reference information shown in the examples above would be entered as follows.

For NIS+:

```
TXT "XFNNISPLUS wiz.com. wiz-nisplus-server 133.33.33.33"
```

For NIS:

```
TXT "XFNNIS woz.com woz-nis-server"
```

Note the mandatory double quotes (" ') delimiting the contents of the TXT record. After making any changes to the DNS table, you must notify the server by either restarting it or sending it a signal to reread the table:

```
#kill -HUP `cat /etc/named.pid`
```

This update effectively adds the next naming system reference to DNS. You can look up this reference using `fnlookup(1)` to see if the information has been added properly. For example, the following command looks up the next naming system reference of the DNS domain `wiz.COM`:

```
#fnlookup -v .../Wiz.COM/
```

Note the mandatory trailing slash (/).

After this administrative step has been taken, clients outside of the NIS+ hierarchy or NIS domain can access and perform operations on the contexts in the NIS+ hierarchy or NIS domain. Foreign NIS+ clients access the hierarchy as unauthenticated NIS+ clients. Continuing the example above, and assuming that NIS+ is federated underneath the DNS domain `wiz.COM`, you can now list the root of the NIS+ enterprise using the command:

```
#fnlist .../Wiz.COM/
```

SEE ALSO `fnlist(1)`, `fnlookup(1)`, `nis+(1)`, `in.named(1M)`, `ypserv(1M)`, `xfn(3XFN)`, `fns(5)`, `fns_nis(5)`, `fns_nis+(5)`, `fns_references(5)`, `fns_x500(5)`

NAME	fns_files – overview of FNS over files implementation
DESCRIPTION	<p>The Federated Naming Service (FNS) provides a method for federating multiple naming services under a single, simple interface for the basic naming operations. One of the naming services supported by FNS is <code>/etc</code> files. FNS provides the XFN interface for performing naming and attribute operations on FNS enterprise objects (organization, site, user, host, and service objects), using files as the naming service. FNS stores bindings for these objects in files and uses them in conjunction with existing <code>/etc</code> files objects.</p>
FNS Policies and /etc Files	<p>FNS defines policies for naming objects in the federated namespace (see <code>fns_policies(5)</code>). At the enterprise level, FNS policies specify naming for organizations, hosts, users, sites, and services. The enterprise-level naming service provides contexts to allow other objects to be named relative to these objects.</p> <p>The organizational unit namespace provides a hierarchical namespace for naming subunits of an enterprise. In <code>/etc</code> files, there is no concept of an organization. Hence, with respect to <code>/etc</code> files as the naming service, there is a single organizational unit context that represents the entire system. Users in an FNS organizational unit correspond to the users in the <code>/etc/passwd</code> file. FNS provides a context for each user in the <code>/etc/passwd</code> file.</p> <p>Hosts in an FNS organizational unit correspond to the hosts in the <code>/etc/hosts</code> file. FNS provides a context for each host in the <code>/etc/hosts</code> file.</p>
Security Considerations	<p>Changes to the FNS information (using the commands <code>fncreate(1M)</code>, <code>fncreate_fs(1M)</code>, <code>fnbind(1)</code>, <code>fndestroy(1M)</code> and <code>fnunbind(1)</code>) can be performed only by the privileged users on the system that exports the <code>/var/fn</code> directory. Also, based on the UNIX user IDs, users are allowed to modify their own contexts, bindings, and attributes, from any machine that mounts the <code>/var/fn</code> directory.</p> <p>For example, the command <code>fncreate(1M)</code> creates FNS related files and directories in the system on which the command is executed. Hence, the invoker of the <code>fncreate(1M)</code> command must have super-user privileges in order to create the user, host, site, and service contexts. However, a user could use the <code>fnunbind(1)</code> command to create calendar bindings in the user's own context, as in this example:</p> <pre>fnbind -r thisuser/service/calendar onc_calendar onc_cal_str jsmith@beatrix</pre> <p>The files object name that corresponds to an FNS composite name can be obtained using <code>fnlookup(1)</code> and <code>fnlist(1)</code>.</p>
USAGE	<p>The files used for storing FNS information are placed in the directory <code>/var/fn</code>. The machine on which <code>/var/fn</code> is located has access to the FNS file. The FNS information can be made accessible to other machines by exporting <code>/var/fn</code>. Client machines that NFS mount the <code>/var/fn</code> directory would then be able to access the FNS information.</p>

fns_files(5)

SEE ALSO | fnbind(1), fnlist(1), fnlookup(1), fnunbind(1), fncreate(1M),
fncreate_fs(1M), fndestroy(1M), xfn(3XFN), fns(5),
fns_initial_context(5), fns_nis(5), fns_nis+(5), fns_policies(5),
fns_references(5)

NAME	fns_initial_context – overview of the FNS Initial Context																				
DESCRIPTION	<p>Every FNS name is interpreted relative to some context, and every FNS naming operation is performed on a context object. The FNS programming interface (XFN) provides a function that allows the client to obtain an <i>Initial Context</i> object. The Initial Context provides the initial pathway to other FNS contexts. FNS defines a set of bindings that the client can expect to find in this context,</p> <p>FNS assumes that for every process:</p> <ol style="list-style-type: none"> 1. There is a user associated with the process when <code>fn_ctx_handle_from_initial()</code> is invoked. This association is based on the effective uid of the process. In the following discussion this user is denoted by <i>U</i>. The association of user to process may change during the life of a process but does not affect the context handle originally returned by <code>fn_ctx_handle_from_initial()</code>. 2. The process is running on a host when <code>fn_ctx_handle_from_initial()</code> is invoked. In the following discussion this host is denoted by <i>H</i>. <p>The following atomic names can appear in the Initial Context:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="text-align: left;">. . .</td> <td style="text-align: left;">thishost</td> <td style="text-align: left;">thisorgunit</td> </tr> <tr> <td style="text-align: left;">thisens</td> <td style="text-align: left;">myself</td> <td style="text-align: left;">myorgunit</td> </tr> <tr> <td style="text-align: left;">myens</td> <td style="text-align: left;">orgunit</td> <td style="text-align: left;">site</td> </tr> <tr> <td style="text-align: left;">user</td> <td style="text-align: left;">host</td> <td></td> </tr> </table> <p>Except for . . . , these names with an added underscore ('_') prefix are also in the Initial Context and have the same binding as their counterpart (for example, <code>thishost</code> and <code>_thishost</code> have the same binding). In addition, <code>org</code> has the same binding as <code>orgunit</code>, and <code>thisuser</code> has the same binding as <code>myself</code>. The bindings for these names are summarized in the following table.</p> <p>Some of these names may not necessarily appear in all Initial Contexts. For example, a process owned by the super-user of a machine does not have any of the user-related bindings. Or, for another example, an installation that has not set up a site namespace will not have the site-related bindings.</p> <table border="0"> <tr> <td style="vertical-align: top;">...</td> <td>global context for resolving DNS or X.500 names. Synonym: / . . .</td> </tr> <tr> <td style="vertical-align: top;">thishost</td> <td><i>H</i>'s host context. Synonym: <code>_thishost</code></td> </tr> <tr> <td style="vertical-align: top;">thisens</td> <td>the enterprise root of <i>H</i>. Synonym: <code>_thisens</code></td> </tr> <tr> <td style="vertical-align: top;">thisorgunit</td> <td><i>H</i>'s distinguished organizational unit context. In Solaris, this is <i>H</i>'s NIS+ home domain. Synonym: <code>_thisorgunit</code></td> </tr> </table>	. . .	thishost	thisorgunit	thisens	myself	myorgunit	myens	orgunit	site	user	host		...	global context for resolving DNS or X.500 names. Synonym: / . . .	thishost	<i>H</i> 's host context. Synonym: <code>_thishost</code>	thisens	the enterprise root of <i>H</i> . Synonym: <code>_thisens</code>	thisorgunit	<i>H</i> 's distinguished organizational unit context. In Solaris, this is <i>H</i> 's NIS+ home domain. Synonym: <code>_thisorgunit</code>
. . .	thishost	thisorgunit																			
thisens	myself	myorgunit																			
myens	orgunit	site																			
user	host																				
...	global context for resolving DNS or X.500 names. Synonym: / . . .																				
thishost	<i>H</i> 's host context. Synonym: <code>_thishost</code>																				
thisens	the enterprise root of <i>H</i> . Synonym: <code>_thisens</code>																				
thisorgunit	<i>H</i> 's distinguished organizational unit context. In Solaris, this is <i>H</i> 's NIS+ home domain. Synonym: <code>_thisorgunit</code>																				

fns_initial_context(5)

myself	<i>U</i> 's user context. Synonyms: <code>_myself</code> , <code>thisuser</code>
myens	the enterprise root of <i>U</i> . Synonym: <code>_myens</code>
myorgunit	<i>U</i> 's distinguished organizational unit context. In Solaris, this is <i>U</i> 's NIS+ home domain. Synonym: <code>_myorgunit</code>
user	the context in which users in the same organizational unit as <i>H</i> are named. Synonym: <code>_user</code>
host	the context in which hosts in the same organizational unit as <i>H</i> are named. Synonym: <code>_host</code>
org	the root context of the organizational unit namespace in <i>H</i> 's enterprise. In Solaris, this corresponds to the NIS+ root domain. Synonyms: <code>orgunit</code> , <code>_orgunit</code>
site	the root context of the site namespace in <i>H</i> 's enterprise, if the site namespace has been configured. Synonym: <code>_site</code>

EXAMPLES

EXAMPLE 1 Names beginning with the enterprise root

The types of objects that may be named relative to the enterprise root are user, host, service, organizational unit, file, and site. Here are some examples of names that begin with the enterprise root:

```
thisens/orgunit/multimedia.servers.engineering
names an organizational unit multimedia.servers.engineering in H's
enterprise.
```

```
thisens/site/northwing.floor3.admin
names the north wing site, on the third floor of the administrations building in H's
enterprise.
```

```
myens/user/hdiffie
names the user hdiffie in U's enterprise.
```

```
myens/service/teletax
names the teletax service of U's enterprise.
```

EXAMPLE 2 Names beginning with organizational unit names

The types of objects that may be named relative to an organizational unit name are: user, host, service, file, and site. Here are some examples of names that begin with organizational unit names (either explicitly via `org`, or implicitly via `thisorgunit` or `myorgunit`), and name objects relative to organizational unit names when resolved in the Initial Context:

```
org/accounts_payable.finance/site/videoconference.northwing
names a conference room videoconference located in the north wing of the site
associated with the organizational unit accounts_payable.finance.
```

EXAMPLE 2 Names beginning with organizational unit names (Continued)

org/finance/user/mjones

names a user mjones in the organizational unit finance.

org/finance/host/inmail

names a machine inmail belonging to the organizational unit finance.

org/accounts_payable.finance/fs/pub/blue-and-whites/FY92-124

names a file pub/blue-and-whites/FY92-124 belonging to the organizational unit accounts_payable.finance.

org/accounts_payable.finance/service/calendar

names the calendar service of the organizational unit accounts_payable.finance. This might manage the meeting schedules of the organizational unit.

thisorgunit/user/cmead

names the user cmead in *H*'s organizational unit.

myorgunit/fs/pub/project_plans/widget.ps

names the file pub/project_plans/widget.ps exported by *U*'s organizational unit's file system.

EXAMPLE 3 Names beginning with site names

The types of objects that may be named relative to a site name are users, hosts, services, and files. Here are some examples of names that begin with site names via site, and name objects relative to sites when resolved in the Initial Context:

site/b5.mtv/service/printer/speedy

names a printer speedy in the b5.mtv site.

site/admin/fs/usr/dist

names a file directory usr/dist available in the site admin.

EXAMPLE 4 Names beginning with user names

The types of objects that may be named relative to a user name are services and files. Here are some examples of names that begin with user names (explicitly via user or implicitly via thisuser), and name objects relative to users when resolved in the Initial Context:

user/jsmith/service/calendar

names the calendar service of the user jsmith.

user/jsmith/fs/bin/games/riddles

names the file bin/games/riddles of the user jsmith.

thisuser/service/printer

names the printer service of *U*.

fns_initial_context(5)

EXAMPLE 5 Names beginning with host names

The types of objects that may be named relative to a host name are services and files. Here are some examples of names that begin with host names (explicitly via `host` or implicitly via `thishost`), and name objects relative to hosts when resolved in the Initial Context:

`host/mailhop/service/mailbox`

names the mailbox service associated with the machine `mailhop`.

`host/mailhop/fs/pub/saf/archives.91`

names the directory `pub/saf/archives.91` found under the root directory of the machine `mailhop`.

`thishost/service/printer`

names the printer service of *H*.

SEE ALSO `nis+(1)`, `geteuid(2)`, `fn_ctx_handle_from_initial(3XFN)`, `xfn(3XFN)`, `fns(5)`, `fns_policies(5)`

NAME	fns_nis+ – overview of FNS over NIS+ implementation
DESCRIPTION	<p>Federated Naming Service (FNS) provides a method for federating multiple naming services under a single, simple interface for the basic naming operations. One of the naming services supported by FNS is NIS+, the enterprise-wide information service in Solaris (see <code>nis+(1)</code>). FNS provides the XFN interface for performing naming and attribute operations on FNS enterprise objects (organization, site, user, host, and service objects) using NIS+. FNS stores bindings for these objects in NIS+ and uses them in conjunction with existing NIS+ objects.</p>
FNS Policies and NIS+	<p>FNS defines policies for naming objects in the federated namespace (see <code>fns_policies(5)</code>). At the enterprise level, FNS policies specify naming for organizations, hosts, users, sites, and services. The enterprise-level naming service provides contexts to allow other objects to be named relative to these objects.</p> <p>The organizational unit namespace provides a hierarchical namespace for naming subunits of an enterprise. An organizational unit maps to an NIS+ domain. Organizational unit names can be either fully qualified NIS+ domain names or relatively NIS+ domain names. If a terminal dot is present in the name, it is treated as a fully qualified name. Otherwise, the name is resolved relative to the root NIS+ domain.</p> <p>Users in the NIS+ namespace are found in the <code>passwd.org_dir</code> table of an NIS+ domain. Users in an FNS organizational unit correspond to the users in the <code>passwd.org_dir</code> table of the corresponding NIS+ domain. FNS provides a context for each user in the <code>passwd.org_dir</code> table.</p> <p>Hosts in the NIS+ namespace are found in the <code>hosts.org_dir</code> table of an NIS+ domain. Hosts in an FNS organizational unit correspond to the hosts in the <code>hosts.org_dir</code> table of the corresponding NIS+ domain. FNS provides a context for each host in the <code>hosts.org_dir</code> table.</p> <p>In NIS+, users and hosts have a notion of a <i>home domain</i>. It is the primary NIS+ domain that maintains information associated with them. A user or host's home domain can be determined directly using its NIS+ principal name, which is composed of the atomic user (login) name or the atomic host name, and the name of the NIS+ home domain. For example, user <code>jsmith</code> with home domain <code>wiz.com</code> has an NIS+ principal name, <code>jsmith.wiz.com</code>.</p> <p>A user's NIS+ home domain corresponds to the user's FNS organizational unit and determines the binding for <code>myens</code> and <code>myorgunit</code>.</p> <p>A host's NIS+ home domain corresponds to the host's FNS organizational unit and determines the binding for <code>thisens</code>, <code>thisorgunit</code>, <code>user</code>, and <code>host</code>.</p>

fns_nis+(5)

Federating NIS+ with DNS or X.500

Federating NIS+ with the global naming systems DNS or X.500 makes NIS+ contexts accessible outside of an NIS+ hierarchy. To enable the federation, the administrator must first add address information in either DNS or X.500 (see `fns_dns(5)` and `fns_x500(5)`). After this administrative step has been taken, clients outside of the NIS+ hierarchy can access contexts and perform operations from outside the hierarchy as an unauthenticated NIS+ client.

NIS+ Security

The command `fncreate(1M)` creates NIS+ tables and directories in the NIS+ hierarchy associated with the domain of the host on which it executes. The invoker of `fncreate(1M)` and other FNS commands is expected to have the necessary NIS+ credentials. (See `nis+(1)` and `nisdefaults(1)`). The environment variable `NIS_GROUP` of the process specifies the group owner for the NIS+ objects thus created. In order to facilitate administration of the NIS+ objects, `NIS_GROUP` should be set to the name of the NIS+ administration group for the domain prior to executing `fncreate(1M)` and other FNS commands. Changes to NIS+-related properties, including default access control rights, could be effected using NIS+ administration tools and interfaces after the context has been created. The NIS+ object name that corresponds to an FNS composite name can be obtained using `fnlookup(1)` and `fnlist(1)`.

SEE ALSO

`fnlist(1)`, `fnlookup(1)`, `nis+(1)`, `nischgrp(1)`, `nischmod(1)`, `nischown(1)`, `nisdefaults(1)`, `nisls(1)`, `fncreate(1M)`, `xfn(3XFN)`, `fns(5)`, `fns_dns(5)`, `fns_files(5)`, `fns_initial_context(5)`, `fns_nis(5)`, `fns_policies(5)`, `fns_references(5)`, `fns_x500(5)`

NAME	fns_nis – overview of FNS over NIS (YP) implementation
DESCRIPTION	<p>Federated Naming Service (FNS) provides a method for federating multiple naming services under a single, simple interface for the basic naming operations. One of the naming services supported by FNS is NIS (YP), the enterprise-wide information services in Solaris (see <code>ypcat(1)</code>, <code>ypmatch(1)</code>, <code>ypfiles(4)</code>). FNS provides the XFN interface for performing naming and attribute operations on FNS enterprise objects (organization, site, user, host and service objects) using NIS. FNS stores bindings for these objects in NIS and uses them in conjunction with existing NIS objects.</p>
FNS Policies and NIS	<p>FNS defines policies for naming objects in the federated namespace (see <code>fns_policies(5)</code>). At the enterprise level, FNS policies specify naming for organizations, hosts, users, sites, and services. The enterprise-level naming service provides contexts to allow other objects to be named relative to these objects.</p> <p>The FNS organizational unit namespace provides a hierarchical namespace for naming subunits of an enterprise. However, NIS does not support a hierarchical organizational structure. Therefore, a NIS domain maps to a single organizational unit in the FNS namespace.</p>
Federating NIS with DNS or X.500	<p>Users in an FNS organizational unit correspond to the users in the <code>passwd.byname</code> map of the corresponding NIS domain. FNS provides a context for each user in the <code>passwd.byname</code> map.</p> <p>Hosts in an FNS organizational unit correspond to the hosts in the <code>hosts.byname</code> map of the corresponding NIS domain. FNS provides a context for each host in the <code>hosts.byname</code> map.</p> <p>Federating NIS with the global naming systems DNS or X.500 makes NIS contexts accessible outside of an NIS domain. To enable the federation, the administrator must first add address information in either DNS or X.500 (see <code>fns_dns(5)</code> and <code>fns_x500(5)</code>). After this administrative step has been taken, clients outside of the NIS domain can access contexts and perform operations.</p>
Security Considerations	<p>Changes to the FNS information (using the commands <code>fncreate(1M)</code>, <code>fncreate_fs(1M)</code>, <code>fncreate_printer(1M)</code>, <code>fnbind(1)</code>, <code>fndestroy(1M)</code>, <code>fncheck(1M)</code>, and <code>fnunbind(1)</code>) can be performed only by the privileged users on the NIS master server that maintains the FNS information.</p> <p>For example, the command <code>fncreate(1M)</code> creates the NIS map for the associated NIS domain in the system on which it is executed. Hence, the command must be run by a privileged user either on the NIS master server or on a system that will serve as a NIS master server for FNS.</p> <p>The NIS object name that corresponds to an FNS composite name can be obtained using <code>fnlookup(1)</code> and <code>fnlist(1)</code>.</p>

fns_nis(5)

SEE ALSO | fnbind(1), fnlist(1), fnlookup(1), fnunbind(1), ypcat(1), ypmatch(1),
fncheck(1M), fncreate(1M), fncreate_fs(1M), fncreate_printer(1M),
fndestroy(1M), xfn(3XFN), ypfiles(4), fns(5), fns_dns(5), fns_files(5),
fns_initial_context(5), fns_nis+(5), fns_policies(5), fns_references(5),
fns_x500(5)

NAME	fns_policies – overview of the FNS Policies
DESCRIPTION	<p>FNS defines policies for naming objects in the federated namespace. The goal of these policies is to allow easy and uniform composition of names. The policies use the basic rule that objects with narrower scopes are named relative to objects with wider scopes.</p> <p>FNS policies are described in terms of the following three categories: global, enterprise, and application.</p> <p><i>Global naming service</i></p> <p>A global naming service is a naming service that has world-wide scope. Internet DNS and X.500 are examples of global naming services. The types of objects named at this global level are typically countries, states, provinces, cities, companies, universities, institutions, and government departments and ministries. These entities are referred to as <i>enterprises</i>.</p> <p><i>Enterprise-level naming service</i></p> <p>Enterprise-level naming services are used to name objects within an enterprise. Within an enterprise, there are naming services that provide contexts for naming common entities such as organizational units, physical sites, human users, and computers. Enterprise-level naming services are bound below the global naming services. Global naming services provide contexts in which the root contexts of enterprise-level naming services can be bound.</p> <p><i>Application-level naming service</i></p> <p>Application-level naming services are incorporated in applications offering services such as file service, mail service, print service, and so on. Application-level naming services are bound below enterprise naming services. The enterprise-level naming services provide contexts in which contexts of application-level naming services can be bound.</p> <p>FNS has policies for global and enterprise naming. Naming within applications is left to individual applications or groups of related applications and not specified by FNS.</p> <p>FNS policy specifies that DNS and X.500 are global naming services that are used to name enterprises. The global namespace is named using the name A DNS name or an X.500 name can appear after the Support for federating global naming services is planned for a future release of FNS.</p>

fns_policies(5)

Within an enterprise, there are namespaces for organizational units, sites, hosts, users, files and services, referred to by the names `orgunit`, `site`, `host`, `user`, `fs`, and `service`. In addition, these namespaces can be named using these names with an added underscore ('_') prefix (for example, `host` and `_host` have the same binding). The following table summarizes the FNS policies.

Context Type	Subordinate Context	Parent Context
org unit	site user host file system service	enterprise root
site	user host file system service	enterprise root org unit
user	service file system	enterprise root org unit
host	service file system	enterprise root org unit
service	not specified	enterprise root org unit site user host
file system	none	enterprise root org unit site user host

In Solaris, an organizational unit name corresponds to an NIS+ domain name and is identified using either the fully-qualified form of its NIS+ domain name, or its NIS+ domain name relative to the NIS+ root. Fully-qualified NIS+ domain names have a terminal dot ('.'). For example, assume that the NIS+ root domain is "Wiz.COM." and "sales" is a subdomain of that. Then, the names `org/sales.Wiz.COM.` and `org/sales` both refer to the organizational unit corresponding to the same NIS+ domain `sales.Wiz.COM.`

User names correspond to names in the corresponding NIS+ `passwd.org_dir` table. The file system context associated with a user is obtained from his entry in the NIS+ `passwd.org_dir` table.

Host names correspond to names in the corresponding NIS+ `hosts.org_dir` table. The file system context associated with a host corresponds to the files systems exported by the host.

EXAMPLES

EXAMPLE 1 The types of objects that may be named relative to an organizational unit name are: user, host, service, file, and site. Here are some examples of names name objects relative to organizational unit names:

`org/accounts_payable.finance/site/videoconference.northwing`
names a conference room `videoconference` located in the north wing of the site associated with the organizational unit `accounts_payable.finance`.

`org/finance/user/mjones`
names a user `mjones` in the organizational unit `finance`.

`org/finance/host/inmail`
names a machine `inmail` belonging to the organizational unit `finance`.

`org/accounts_payable.finance/fs/pub/blue-and-whites/FY92-124`
names a file `pub/blue-and-whites/FY92-124` belonging to the organizational unit `accounts_payable.finance`.

`org/accounts_payable.finance/service/calendar`
names the `calendar` service of the organizational unit `accounts_payable.finance`. This might manage the meeting schedules of the organizational unit.

EXAMPLE 2 The types of objects that may be named relative to a site name are services and files. Here are some examples of names that name objects relative to sites:

`site/b5.mtv/service/printer/speedy`
names a printer `speedy` in the `b5.mtv` site.

`site/admin/fs/usr/dist`
names a file directory `usr/dist` available in the site `admin`.

EXAMPLE 3 The types of objects that may be named relative to a user name are services and files. Here are some examples of names that name objects relative to users:

`user/jsmith/service/calendar`
names the `calendar` service of the user `jsmith`.

fns_policies(5)

EXAMPLE 3 The types of objects that may be named relative to a user name are services and files. Here are some examples of names that name objects relative to users: *(Continued)*

```
user/jsmith/fs/bin/games/riddles
names the file bin/games/riddles of the user jsmith.
```

EXAMPLE 4 The types of objects that may be named relative to a host name are services and files. Here are some examples of names that name objects relative to hosts:

```
host/mailhop/service/mailbox
names the mailbox service associated with the machine mailhop.
```

```
host/mailhop/fs/pub/saf/archives.91
names the directory pub/saf/archives.91 found under the root directory of the
machine mailhop.
```

SEE ALSO `fncreate(1M)`, `nis+(1)`, `xfn(3XFN)`, `fns(5)`, `fns_initial_context(5)`,
`fns_references(5)`

NAME	fns_references – overview of FNS References																														
DESCRIPTION	<p>Every composite name in FNS is bound to a <i>reference</i>. A reference consists of a type and a list of addresses. The reference type is used to identify the type of object.</p> <p>An address is something that can be used with some communication mechanism to invoke operations on an object or service. Multiple addresses are intended to identify multiple communication endpoints for a single conceptual object or service. Each address in a reference consists of an address type and an opaque buffer. The address type determines the format and interpretation of the address data. Together, the address's type and data specify how to reach the object. Many communication mechanisms are possible; FNS does not place any restrictions on them.</p> <p>The following summarizes the reference and address types that are currently defined. New types should be registered with the Federated Naming Group at SunSoft.</p>																														
Reference Types	<p>All reference types use the FN_ID_STRING identifier format unless otherwise qualified.</p> <table border="0"> <tr> <td>onc_fn_enterprise</td> <td>Enterprise root context.</td> </tr> <tr> <td>onc_fn_organization</td> <td>A context for naming objects related to an organizational unit.</td> </tr> <tr> <td>onc_fn_hostname</td> <td>A context for naming hosts.</td> </tr> <tr> <td>onc_fn_username</td> <td>A context for naming users.</td> </tr> <tr> <td>onc_fn_user</td> <td>A context for naming objects related to a user.</td> </tr> <tr> <td>onc_fn_host</td> <td>A context for naming objects related to a computer.</td> </tr> <tr> <td>onc_fn_site</td> <td>A context for naming sites.</td> </tr> <tr> <td>onc_fn_service</td> <td>A context for naming services.</td> </tr> <tr> <td>onc_fn_nsid</td> <td>A context for naming namespace identifiers.</td> </tr> <tr> <td>onc_fn_generic</td> <td>A context for naming application-specific objects.</td> </tr> <tr> <td>onc_fn_fs</td> <td>A context for naming files, directories, and file systems.</td> </tr> <tr> <td>onc_fn_printername</td> <td>A context for naming printers.</td> </tr> <tr> <td>onc_printers</td> <td>A printer object. When implemented on top of NIS+, this could also be a context for naming printers.</td> </tr> <tr> <td>fn_link_ref</td> <td>An XFN link.</td> </tr> <tr> <td>inet_domain</td> <td>An Internet domain.</td> </tr> </table>	onc_fn_enterprise	Enterprise root context.	onc_fn_organization	A context for naming objects related to an organizational unit.	onc_fn_hostname	A context for naming hosts.	onc_fn_username	A context for naming users.	onc_fn_user	A context for naming objects related to a user.	onc_fn_host	A context for naming objects related to a computer.	onc_fn_site	A context for naming sites.	onc_fn_service	A context for naming services.	onc_fn_nsid	A context for naming namespace identifiers.	onc_fn_generic	A context for naming application-specific objects.	onc_fn_fs	A context for naming files, directories, and file systems.	onc_fn_printername	A context for naming printers.	onc_printers	A printer object. When implemented on top of NIS+, this could also be a context for naming printers.	fn_link_ref	An XFN link.	inet_domain	An Internet domain.
onc_fn_enterprise	Enterprise root context.																														
onc_fn_organization	A context for naming objects related to an organizational unit.																														
onc_fn_hostname	A context for naming hosts.																														
onc_fn_username	A context for naming users.																														
onc_fn_user	A context for naming objects related to a user.																														
onc_fn_host	A context for naming objects related to a computer.																														
onc_fn_site	A context for naming sites.																														
onc_fn_service	A context for naming services.																														
onc_fn_nsid	A context for naming namespace identifiers.																														
onc_fn_generic	A context for naming application-specific objects.																														
onc_fn_fs	A context for naming files, directories, and file systems.																														
onc_fn_printername	A context for naming printers.																														
onc_printers	A printer object. When implemented on top of NIS+, this could also be a context for naming printers.																														
fn_link_ref	An XFN link.																														
inet_domain	An Internet domain.																														
Address Types	<p>All address types use the FN_ID_STRING identifier format unless otherwise qualified. The format of address contents is determined by the corresponding address type.</p> <table border="0"> <tr> <td>onc_fn_nisplus</td> <td>For an FNS enterprise-level object implemented on top of NIS+. The address</td> </tr> </table>	onc_fn_nisplus	For an FNS enterprise-level object implemented on top of NIS+. The address																												
onc_fn_nisplus	For an FNS enterprise-level object implemented on top of NIS+. The address																														

fns_references(5)

	contains the context type, context representation type (either normal or merged), version number of the reference, and the NIS+ name of the object. The only intended use of this reference is that it be passed to <code>fn_ctx_handle_from_ref(3XFN)</code>
<code>onc_fn_nis</code>	For an FNS enterprise-level object implemented on top of NIS. The address contains the context type and version number of the reference, and the NIS name of the object. The only intended use of this reference is that it be passed to <code>fn_ctx_handle_from_ref(3XFN)</code> .
<code>onc_fn_files</code>	For an FNS enterprise-level object implemented on top of <code>/etc</code> files. The address contains the context type and version number of the reference, and the location of the object in the <code>/etc</code> file system. The only intended use of this reference is that it be passed to <code>fn_ctx_handle_from_ref(3XFN)</code> .
<code>onc_fn_fs_user</code>	For a user's home directory. The address contains the user's name and the name of the naming service password table where the user's home directory is stored.
<code>onc_fn_fs_user_nisplus</code>	For a user's home directory. The address contains the user's name and the name of the NIS+ password table where the user's home directory is stored.
<code>onc_fn_fs_host</code>	For all file systems exported by a host. The address contains the host's name.
<code>onc_fn_fs_mount</code>	For a single mount point. The address contains the mount options, the name of the servers and the exported path. See <code>mount(1M)</code> .
<code>onc_fn_printer_files</code>	For a printer's address in the files naming service.
<code>onc_fn_printer_nis</code>	For a printer's address in the NIS naming service.
<code>onc_fn_printer_nisplus</code>	For a printer's address in the NIS+ naming service.

fn_link_addr	For an XFN link address. The contents is the string form of the composite name.
inet_domain	For an Internet domain. The address contains the fully-qualified domain name (for example, "Wiz.COM.")
inet_ipaddr_string	For an object with an Internet address. The address contains an internet IP address in dotted string form (for example, "192.144.2.3").
x500	For an X.500 object. The address contains an X.500 Distinguished Name, in the syntax specified in the X/Open DCE: Directory Services.
osi_paddr	For an object with an OSI presentation address. The address contains the string encoding of an OSI Presentation Address as defined in <i>A string encoding of Presentation Address</i> (RFC 1278).
onc_printers_bsaddr	For a printer that understands the BSD print protocol. The address contains the machine name and printer name used by the protocol.
onc_printers_use	For a printer alias. The address contains a printer name.
onc_printers_all	For a list of printers that are enumerated using the "all" option. The address contains a list of printer names.
onc_printers_location	For a printer's location. The address format is unspecified.
onc_printers_type	For a printer's type. The address format is unspecified.
onc_printers_speed	For a printer's speed. The address format is unspecified.

SEE ALSO mount(1M), fn_ctx_handle_from_ref(3XFN), xfn(3XFN), fns(5), fns_policies(5)

Hardcastle-Kille, S.E., *A string encoding of Presentation Address*, RFC 1278, University College London, November 1991.

fns_x500(5)

NAME	fns_x500 – overview of FNS over X.500 implementation
DESCRIPTION	<p>Federated Naming Service (FNS) provides a method for federating multiple naming services under a single, simple interface for the basic naming operations. One of the naming services supported by FNS is the X.500 Directory Service (see ITU-T X.500 or ISO/IEC 9594). X.500 is a global directory service. Its components cooperate to manage information about a hierarchy of objects on a worldwide scope. Such objects include countries, organizations, people, services, and machines. FNS uses X.500 to name entities globally.</p> <p>FNS provides the XFN interface for retrieval and modification of information stored in X.500. In addition, enterprise namespaces such as those served by NIS+ and NIS can be federated with X.500 by adding reference information to X.500 describing how to reach the desired next naming service. To federate a NIS+ or NIS namespace under X.500, perform the following steps:</p> <ol style="list-style-type: none">1. Obtain the root reference for the NIS+ hierarchy or NIS domain.2. Enhance the X.500 schema to support the addition of XFN references.3. Create an X.500 entry to store the XFN reference.4. Add the XFN reference. <p>The root reference is referred to as the <i>next naming system reference</i> because it refers to the <i>next</i> naming system beneath X.500. This reference contains information about how to communicate with the NIS+ or NIS servers and has the following format:</p> <pre><domainname> <server name> [<server address>]</pre> <p>where <domainname> is the fully qualified domain name. Notice that NIS+ and NIS have slightly different syntaxes for domain names. For NIS+, the fully qualified domain name is case-insensitive and terminated by a dot character ('.'). For NIS, the fully qualified domain name is case-sensitive and <i>not</i> terminated by a dot character. For both NIS+ and NIS, <server address> is optional. If it is not supplied, a host name lookup will be performed to get the machine's address.</p> <p>For example, if the machine <code>wiz-nisplus-server</code> with address <code>133.33.33.33</code> serves the NIS+ domain <code>wiz.com.</code>, the reference would look like this:</p> <pre>wiz.com. wiz-nisplus-server 133.33.33.33</pre> <p>For another example, if the machine <code>woz-nis-server</code> serves the NIS domain <code>Woz.COM</code>, the reference would look like this:</p> <pre>Woz.COM woz-nis-server</pre> <p>Before the next naming system reference can be added to X.500, the X.500 schema must be altered to include the following object class and associated attributes (defined in ASN.1 notation).</p>


```

xFNSSupplement OBJECT-CLASS ::= {
    SUBCLASS OF { top }
    KIND          auxiliary
    MAY CONTAIN { objectReferenceString | nNSReference-
String }
    ID            id-oc-xFNSSupplement }

id-oc-xFNSSupplement OBJECT IDENTIFIER ::= {
    iso member-body(2) ansi(840) sun(113536) 25 }

objectReferenceString ATTRIBUTE ::= {
    WITH SYNTAX          OCTET STRING
    EQUALITY MATCHING RULE octetStringMatch
    SINGLE VALUE        TRUE
    ID                   id-at-objectReferenceString }

id-at-objectReferenceString OBJECT IDENTIFIER ::= {
    iso member-body(2) ansi(840) sun(113536) 30 }

nNSReferenceString ATTRIBUTE ::= {
    WITH SYNTAX          OCTET STRING
    EQUALITY MATCHING RULE octetStringMatch
    SINGLE VALUE        TRUE
    ID                   id-at-nNSReferenceString }

id-at-nNSReferenceString OBJECT IDENTIFIER ::= {
    iso member-body(2) ansi(840) sun(113536) 31 }

```

The procedures for altering the X.500 schema will vary from implementation to implementation. Consult *Solstice X.500* or the schema administration guide for your X.500 product.

Once X.500 supports XFN references, the next naming system reference can be added by first creating an X.500 object and then adding the new reference to it. For example, the following commands create entries for the *Wiz* and *Woz* organizations in the U.S.A. and add the reference information shown in the examples above to them.

For NIS+:

```

example% fnattr ../c=us/o=wiz -a objectclass \
    top organization xfnssupplement

example% fnbind -r ../c=us/o=wiz/ onc_fn_enterprise \
    onc_fn_nisplus_root "wiz.com. wiz-nisplus-server"

```

For NIS:

```

example% fnattr ../c=us/o=woz -a objectclass \
    top organization xfnssupplement

example% fnbind -r ../c=us/o=woz/ onc_fn_enterprise \
    onc_fn_nis_root "Woz.COM woz-nis-server"

```

Notice the mandatory trailing slash ('/') in the name argument to `fnbind(1)`.

fns_x500(5)

This modification effectively adds the next naming system reference to X.500. The reference may be retrieved using `fnlookup(1)` to see if the information has been added properly. For example, the following command looks up the next naming system reference of the `Wiz` organization:

```
example% fnlookup -v ../c=us/o=wiz/
```

Note the mandatory trailing slash.

After this administrative step has been taken, clients outside of the NIS+ hierarchy or NIS domain can access and perform operations on the contexts in the NIS+ hierarchy or NIS domain. Foreign NIS+ clients access the hierarchy as unauthenticated NIS+ clients. Continuing the example above, and assuming that NIS+ is federated underneath the `Wiz` organization, the root of the NIS+ enterprise may be listed using the command:

```
example% fnlist ../c=us/o=wiz/
```

Note the mandatory trailing slash.

The next naming system reference may be removed using the command:

```
example% fnunbind ../c=us/o=wiz/
```

Note the mandatory trailing slash.

SEE ALSO `fnattr(1)`, `fnbind(1)`, `fnlist(1)`, `fnlookup(1)`, `nis+(1)`, `ypserv(1M)`, `xfn(3XFN)`, `fns(5)`, `fns_dns(5)`, `fns_nis(5)`, `fns_nis+(5)`, `fns_references(5)`

Solstice X.500

NOTES In a 64-bit XFN application, retrieval and modification of information stored in the X.500 directory service is not supported.

NAME	formats – file format notation												
DESCRIPTION	Utility descriptions use a syntax to describe the data organization within files—stdin, stdout, stderr, input files, and output files—when that organization is not otherwise obvious. The syntax is similar to that used by the <code>printf(3C)</code> function. When used for stdin or input file descriptions, this syntax describes the format that could have been used to write the text to be read, not a format that could be used by the <code>scanf(3C)</code> function to read the input file.												
Format	<p>The description of an individual record is as follows:</p> <pre>"<format>", [<arg1>, <arg2>, . . . , <argn>]</pre> <p>The <code>format</code> is a character string that contains three types of objects defined below:</p> <p><i>characters</i> Characters that are not <i>escape sequences</i> or <i>conversion specifications</i>, as described below, are copied to the output.</p> <p><i>escape sequences</i> Represent non-graphic characters.</p> <p><i>conversion specifications</i> Specifies the output format of each argument. (See below.)</p> <p>The following characters have the following special meaning in the format string:</p> <p>” (An empty character position.) One or more blank characters.</p> <p>^ Exactly one space character.</p> <p>The notation for spaces allows some flexibility for application output. Note that an empty character position in <code>format</code> represents one or more blank characters on the output (not <i>white space</i>, which can include newline characters). Therefore, another utility that reads that output as its input must be prepared to parse the data using <code>scanf(3C)</code>, <code>awk(1)</code>, and so forth. The character is used when exactly one space character is output.</p>												
Escape Sequences	<p>The following table lists escape sequences and associated actions on display devices capable of the action.</p> <table border="1"> <thead> <tr> <th>Sequence</th> <th>Character</th> <th>Terminal Action</th> </tr> </thead> <tbody> <tr> <td>\\</td> <td>backslash</td> <td>None.</td> </tr> <tr> <td>\a</td> <td>alert</td> <td>Attempts to alert the user through audible or visible notification.</td> </tr> <tr> <td>\b</td> <td>backspace</td> <td>Moves the printing position to one column before the current position, unless the current position is the start of a line.</td> </tr> </tbody> </table>	Sequence	Character	Terminal Action	\\	backslash	None.	\a	alert	Attempts to alert the user through audible or visible notification.	\b	backspace	Moves the printing position to one column before the current position, unless the current position is the start of a line.
Sequence	Character	Terminal Action											
\\	backslash	None.											
\a	alert	Attempts to alert the user through audible or visible notification.											
\b	backspace	Moves the printing position to one column before the current position, unless the current position is the start of a line.											

formats(5)

Sequence	Character	Terminal Action
\f	form-feed	Moves the printing position to the initial printing position of the next logical page.
\n	newline	Moves the printing position to the start of the next line.
\r	carriage-return	Moves the printing position to the start of the current line.
\t	tab	Moves the printing position to the next tab position on the current line. If there are no more tab positions left on the line, the behavior is undefined.
\v	vertical-tab	Moves the printing position to the start of the next vertical tab position. If there are no more vertical tab positions left on the page, the behavior is undefined.

Conversion Specifications

Each conversion specification is introduced by the percent-sign character (%). After the character %, the following appear in sequence:

- flags* Zero or more *flags*, in any order, that modify the meaning of the conversion specification.
- field width* An optional string of decimal digits to specify a minimum *field width*. For an output field, if the converted value has fewer bytes than the field width, it is padded on the left (or right, if the left-adjustment flag (-), described below, has been given to the field width).
- precision* Gives the minimum number of digits to appear for the d, o, i, u, x or X conversions (the field is padded with leading zeros), the number of digits to appear after the radix character for the e and f conversions, the maximum number of significant digits for the g conversion; or the maximum number of bytes to be written from a string in s conversion. The precision takes the form of a period (.) followed by a decimal digit string; a null digit string is treated as zero.
- conversion characters* A conversion character (see below) that indicates the type of conversion to be applied.
- flags* The *flags* and their meanings are:
 - The result of the conversion is left-justified within the field.
 - + The result of a signed conversion always begins with a sign (+ or -).

<code><space></code>	If the first character of a signed conversion is not a sign, a space character is prefixed to the result. This means that if the space character and + flags both appear, the space character flag is ignored.
#	The value is to be converted to an alternative form. For c, d, i, u, and s conversions, the behaviour is undefined. For o conversion, it increases the precision to force the first digit of the result to be a zero. For x or X conversion, a non-zero result has 0x or 0X prefixed to it, respectively. For e, E, f, g, and G conversions, the result always contains a radix character, even if no digits follow the radix character. For g and G conversions, trailing zeros are not removed from the result as they usually are.
0	For d, i, o, u, x, X, e, E, f, g, and G conversions, leading zeros (following any indication of sign or base) are used to pad to the field width; no space padding is performed. If the 0 and - flags both appear, the 0 flag is ignored. For d, i, o, u, x and X conversions, if a precision is specified, the 0 flag is ignored. For other conversions, the behaviour is undefined.
Conversion Characters	<p>Each conversion character results in fetching zero or more arguments. The results are undefined if there are insufficient arguments for the format. If the format is exhausted while arguments remain, the excess arguments are ignored.</p> <p>The <i>conversion characters</i> and their meanings are:</p> <p><i>d,i,o,u,x,X</i> The integer argument is written as signed decimal (d or i), unsigned octal (o), unsigned decimal (u), or unsigned hexadecimal notation (x and X). The d and i specifiers convert to signed decimal in the style [-]ddd. The x conversion uses the numbers and letters 0123456789abcdef and the X conversion uses the numbers and letters 0123456789ABCDEF. The <i>precision</i> component of the argument specifies the minimum number of digits to appear. If the value being converted can be represented in fewer digits than the specified minimum, it is expanded with leading zeros. The default precision is 1. The result of converting a zero value with a precision of 0 is no characters. If both the field width and precision are omitted, the implementation may precede, follow or precede and follow numeric arguments of types d, i and u with blank characters; arguments of type o (octal) may be preceded with leading zeros.</p> <p>The treatment of integers and spaces is different from the printf(3C) function in that they can be surrounded with blank characters. This was done so that, given a format such as:</p> <pre>"%d\n", <foo></pre> <p>the implementation could use a printf() call such as:</p>

formats(5)

```
printf("%6d\n", foo);
```

and still conform. This notation is thus somewhat like `scanf()` in addition to `printf()`.

f The floating point number argument is written in decimal notation in the style `[-]ddd.ddd`, where the number of digits after the radix character (shown here as a decimal point) is equal to the *precision* specification. The `LC_NUMERIC` locale category determines the radix character to use in this format. If the *precision* is omitted from the argument, six digits are written after the radix character; if the *precision* is explicitly 0, no radix character appears.

e,E The floating point number argument is written in the style `[-]d.ddde±dd` (the symbol \pm indicates either a plus or minus sign), where there is one digit before the radix character (shown here as a decimal point) and the number of digits after it is equal to the precision. The `LC_NUMERIC` locale category determines the radix character to use in this format. When the precision is missing, six digits are written after the radix character; if the precision is 0, no radix character appears. The E conversion character produces a number with E instead of e introducing the exponent. The exponent always contains at least two digits. However, if the value to be written requires an exponent greater than two digits, additional exponent digits are written as necessary.

g,G The floating point number argument is written in style *f* or *e* (or in style *E* in the case of a *G* conversion character), with the precision specifying the number of significant digits. The style used depends on the value converted: style *g* is used only if the exponent resulting from the conversion is less than -4 or greater than or equal to the precision. Trailing zeros are removed from the result. A radix character appears only if it is followed by a digit.

c The integer argument is converted to an unsigned char and the resulting byte is written.

s The argument is taken to be a string and bytes from the string are written until the end of the string or the number of bytes indicated by the *precision* specification of the argument is reached. If the precision is omitted from the argument, it is taken to be infinite, so all bytes up to the end of the string are written.

% Write a % character; no argument is converted.

In no case does a non-existent or insufficient *field width* cause truncation of a field; if the result of a conversion is wider than the field width, the field is simply expanded to contain the conversion result. The term *field width* should not be confused with the term *precision* used in the description of *%s*.

One difference from the C function `printf()` is that the `l` and `h` conversion characters are not used. There is no differentiation between decimal values for type `int`, type `long`, or type `short`. The specifications `%d` or `%i` should be interpreted as an arbitrary length sequence of digits. Also, no distinction is made between single precision and double precision numbers (`float` or `double` in C). These are simply referred to as floating point numbers.

Many of the output descriptions use the term `line`, such as:

```
"%s", <input line>
```

Since the definition of `line` includes the trailing newline character already, there is no need to include a `\n` in the format; a double newline character would otherwise result.

EXAMPLES

EXAMPLE 1 To represent the output of a program that prints a date and time in the form Sunday, July 3, 10:02, where `<weekday>` and `<month>` are strings:

```
"%s, /\%s/\%d, /\%d:%.2d\n", <weekday>, <month>, <day>, <hour>, <min>
```

EXAMPLE 2 To show pi written to 5 decimal places:

```
"pi/\=/\%.5f\n", <value of pi>
```

EXAMPLE 3 To show an input file format consisting of five colon-separated fields:

```
"%s:%s:%s:%s:%s\n", <arg1>, <arg2>, <arg3>, <arg4>, <arg5>
```

SEE ALSO

`awk(1)`, `printf(1)`, `printf(3C)`, `scanf(3C)`

fsattr(5)

NAME	fsattr – extended file attributes						
DESCRIPTION	<p>Attributes are logically supported as files within the file system. The file system is therefore augmented with an orthogonal name space of file attributes. Any file (including attribute files) can have an arbitrarily deep attribute tree associated with it. Attribute values are accessed by file descriptors obtained through a special attribute interface. This logical view of "attributes as files" allows the leveraging of existing file system interface functionality to support the construction, deletion, and manipulation of attributes.</p> <p>The special files "." and ".." retain their accustomed semantics within the attribute hierarchy. The "." attribute file refers to the current directory and the ".." attribute file refers to the parent directory. The unnamed directory at the head of each attribute tree is considered the "child" of the file it is associated with and the ".." file refers to the associated file. For any non-directory file with attributes, the ".." entry in the unnamed directory refers to a file that is not a directory.</p> <p>Conceptually, the attribute model is fully general. Extended attributes can be any type of file (doors, links, directories, and so forth) and can even have their own attributes (fully recursive). As a result, the attributes associated with a file could be an arbitrarily deep directory hierarchy where each attribute could have an equally complex attribute tree associated with it. Not all implementations are able to, or want to, support the full model. Implementation are therefore permitted to reject operations that are not supported. For example, the implementation for the UFS file system allows only regular files as attributes (for example, no sub-directories) and rejects attempts to place attributes on attributes.</p> <p>The following list details the operations that are rejected in the current implementation:</p> <table><tr><td>link</td><td>Any attempt to create links between attribute and non-attribute space is rejected to prevent security-related or otherwise sensitive attributes from being exposed, and therefore manipulable, as regular files.</td></tr><tr><td>rename</td><td>Any attempt to rename between attribute and non-attribute space is rejected to prevent an already linked file from being renamed and thereby circumventing the link restriction above.</td></tr><tr><td>mkdir, symlink, mknod</td><td>Any attempt to create a "non-regular" file in attribute space is rejected to reduce the functionality, and therefore exposure and risk, of the initial implementation.</td></tr></table>	link	Any attempt to create links between attribute and non-attribute space is rejected to prevent security-related or otherwise sensitive attributes from being exposed, and therefore manipulable, as regular files.	rename	Any attempt to rename between attribute and non-attribute space is rejected to prevent an already linked file from being renamed and thereby circumventing the link restriction above.	mkdir, symlink, mknod	Any attempt to create a "non-regular" file in attribute space is rejected to reduce the functionality, and therefore exposure and risk, of the initial implementation.
link	Any attempt to create links between attribute and non-attribute space is rejected to prevent security-related or otherwise sensitive attributes from being exposed, and therefore manipulable, as regular files.						
rename	Any attempt to rename between attribute and non-attribute space is rejected to prevent an already linked file from being renamed and thereby circumventing the link restriction above.						
mkdir, symlink, mknod	Any attempt to create a "non-regular" file in attribute space is rejected to reduce the functionality, and therefore exposure and risk, of the initial implementation.						

The entire available name space has been allocated to "general use" to bring the implementation in line with the NFSv4 draft standard [NFSv4]. That standard defines "named attributes" (equivalent to Solaris Extended Attributes) with no naming restrictions. All Sun applications making use of opaque extended attributes will use the prefix "SUNW".

Shell-level API

The command interface for extended attributes is the set of applications provided by Solaris for the manipulation of attributes from the command line. This interface consists of a set of existing utilities that have been extended to be "attribute-aware", plus the `runat` utility designed to "expose" the extended attribute space so that extended attributes can be manipulated as regular files.

The `-@` option enable utilities to manipulate extended attributes. As a rule, this option enables the utility to enter into attribute space when the utility is performing a recursive traversal of file system space. This is a fully recursive concept. If the underlying file system supports recursive attributes and directory structures, the `-@` option opens these spaces to the file tree-walking algorithms.

The following utilities accommodate extended attributes (see the individual manual pages for details):

<code>cp</code>	By default, <code>cp</code> ignores attributes and copies only file data. This is intended to maintain the semantics implied by <code>cp</code> currently, where attributes (such as owner and mode) are not copied unless the <code>-p</code> option is specified. With the <code>-@</code> (or <code>-p</code>) option, <code>cp</code> attempts to copy all attributes along with the file data.
<code>cpio</code>	The <code>-@</code> option informs <code>cpio</code> to archive attributes, but by default <code>cpio</code> ignores extended attributes. See <i>Extended Archive Formats</i> below for a description of the new archive records.
<code>du</code>	File sizes computed include the space allocated for any extended attributes present.
<code>find</code>	By default, <code>find</code> ignores attributes. The <code>-xattr</code> expression provides support for searches involving attribute space. It returns true if extended attributes are present on the current file.
<code>fsck</code>	The <code>fsck</code> utility manages extended attribute data on the disk. A file system with extended attributes can be mounted on versions of Solaris that are not attribute-aware (versions prior to Solaris 9), but the attributes will not be accessible and <code>fsck</code> will strip them from the files and place them in <code>lost+found</code> . Once the attributes have been stripped the file system is completely stable on Solaris versions that are not attribute-aware, but would now be considered

	corrupted on attribute-aware versions of Solaris. The attribute-aware <code>fsck</code> utility should be run to stabilize the file system before using it in an attribute-aware environment.
<code>fsdb</code>	This <code>fsdb</code> utility is able to find the inode for the "hidden" extended attribute directory.
<code>ls</code>	The <code>ls -@</code> command displays an "@" following the mode information when extended attributes are present. More precisely, the output line for a given file contains an "@" character following the mode characters if the <code>pathconf(2)</code> variable <code>XATTR_EXISTS</code> is set to true. See the <code>pathconf()</code> section below. The <code>-@</code> option uses the same general output format as the <code>-l</code> option.
<code>mv</code>	When a file is moved, all attributes are carried along with the file rename. When a file is moved across a file system boundary, the copy command invoked is similar to the <code>cp -p</code> variant described above and extended attributes are "moved". If the extended file attributes cannot be replicated, the move operation fails and the source file is not removed.
<code>pax</code>	The <code>-@</code> option informs <code>pax</code> to archive attributes, but by default <code>pax</code> ignores extended attributes. The <code>pax(1)</code> utility is a generic replacement for both <code>tar(1)</code> and <code>cpio(1)</code> and is able to produce either output format in its archive. See <code>Extended Archive Formats</code> below for a description of the new archive records.
<code>tar</code>	In the default case, <code>tar</code> does not attempt to place attributes in the archive. If the <code>-@</code> option is specified, however, <code>tar</code> traverses into the attribute space of all files being placed in the archive and attempts to add the attributes to the archive. A new record type has been introduced for extended attribute entries in <code>tar</code> archive files (the same is true for <code>pax</code> and <code>cpio</code> archives) similar to the way <code>ACLs</code> records were defined. See <code>Extended Archive Formats</code> below for a description of the new archive records.

There is a class of utilities (`chmod`, `chown`, `chgrp`) that one might expect to be modified in a manner similar to those listed above. For example, one might expect that performing `chmod` on a file would not only affect the file itself but would also affect at least the extended attribute directory if not any existing extended attribute files. This is not the case. The model chosen for extended attributes implies that the attribute directory and the attributes themselves are all file objects in their own right, and can therefore have independent file status attributes associated with them (a given

implementation cannot support this, for example, for intrinsic attributes). The relationship is left undefined and a fine-grained control mechanism (`runat(1)`) is provided to allow manipulation of extended attribute status attributes as necessary.

The `runat` utility has the following syntax:

```
runat filename [command]
```

The `runat` utility executes the supplied command in the context of the "attribute space" associated with the indicated file. If no command argument is supplied, a shell is invoked. See `runat(1)` for details.

Application-level API

The primary interface required to access extended attributes at the programmatic level is the `openat(2)` function. Once a file descriptor has been obtained for an attribute file by an `openat()` call, all normal file system semantics apply. There is no attempt to place special semantics on `read(2)`, `write(2)`, `ftruncate(3C)`, or other functions when applied to attribute file descriptors relative to "normal" file descriptors.

The set of existing attributes can be browsed by calling `openat()` with "." as the file name and the `O_XATTR` flag set, resulting in a file descriptor for the attribute directory. The list of attributes is obtained by calls to `getdents(2)` on the returned file descriptor. If the target file did not previously have any attributes associated with it, an empty top-level attribute directory is created for the file and subsequent `getdents()` calls will return only "." and "..". While the owner of the parent file owns the extended attribute directory, it is not charged against its quota if the directory is empty. Attribute files themselves, however, are charged against the user quota as any other regular file.

Additional system calls have been provided as convenience functions. These include the `fchownat(2)`, `fstatat(2)`, `futimesat(2)`, `renameat(2)`, `unlinkat(2)`. These new functions, along with `openat()`, provide a mechanism to access files relative to an arbitrary point in the file system, rather than only the current working directory. This mechanism is particularly useful in situations when a file descriptor is available with no path. The `openat()` function, in particular, can be used in many contexts where `chdir()` or `fchdir()` is currently required. See `chdir(2)`.

Open a file relative to a file descriptor

```
int openat (int fd, const char *path, int oflag [, mode_t mode])
```

The `openat(2)` function behaves exactly as `open(2)` except when given a relative path. Where `open()` resolves a relative path from the current working directory, `openat()` resolves the path based on the vnode indicated by the supplied file descriptor. When `oflag` is `O_XATTR`, `openat()` interprets the `path` argument as an extended attribute reference. The following code fragment uses `openat()` to examine the attributes of some already opened file:

```
dfd = openat(fd, ".", O_RDONLY|O_XATTR);
(void) getdents(dfd, buf, nbytes);
```

If `openat()` is passed the special value `AT_FDCWD` as its first (*fd*) argument, its behavior is identical to `open()` and the relative path arguments are interpreted relative to the current working directory. If the `O_XATTR` flag is provided to `openat()` or to `open()`, the supplied path is interpreted as a reference to an extended attribute on the current working directory.

Unlink a file relative to a directory file descriptor

```
int unlinkat (int dirfd, const char *pathflag, int flagflag)
```

The `unlinkat(2)` function deletes an entry from a directory. The *path* argument indicates the name of the entry to remove. If *path* an absolute path, the *dirfd* argument is ignored. If it is a relative path, it is interpreted relative to the directory indicated by the *dirfd* argument. If *dirfd* does not refer to a valid directory, the function returns `ENOTDIR`. If the special value `AT_FDCWD` is specified for *dirfd*, a relative path argument is resolved relative to the current working directory. If the *flag* argument is 0, all other semantics of this function are equivalent to `unlink(2)`. If *flag* is set to `AT_REMOVEDIR`, all other semantics of this function are equivalent to `rmdir(2)`.

Rename a file relative to directories

```
int renameat (int fromfd, const char *old, int tofd, const char *new)
```

The `renameat(2)` function renames an entry in a directory, possibly moving the entry into a different directory. The *old* argument indicates the name of the entry to rename. If this argument is a relative path, it is interpreted relative to the directory indicated by the *fd* argument. If it is an absolute path, the *fromfd* argument is ignored. The *new* argument indicates the new name for the entry. If this argument is a relative path, it is interpreted relative to the directory indicated by the *tofd* argument. If it is an absolute path, the *tofd* argument is ignored.

In the relative path cases, if the directory file descriptor arguments do not refer to a valid directory, the function returns `ENOTDIR`. All other semantics of this function are equivalent to `rename(2)`.

If a special value `AT_FDCWD` is specified for either the *fromfd* or *tofd* arguments, their associated path arguments (*old* and *new*) are interpreted relative to the current working directory if they are not specified as absolute paths. Any attempt to use `renameat()` to move a file that is not an extended attribute into an extended attribute directory (so that it becomes an extended attribute) will fail. The same is true for an attempt to move a file that is an extended attribute into a directory that is not an extended attribute directory.

Obtain information about a file

```
int fstatat (int fd, const char *path, struct stat* buf, int flag)
```

The `fstatat(2)` function obtains information about a file. If the *path* argument is relative, it is resolved relative to the *fd* argument file descriptor, otherwise the *fd* argument is ignored. If the *fd* argument is a special value `AT_FDCWD` the path is resolved relative to the current working directory. If the *path* argument is a null

pointer, the function returns information about the file referenced by the *fd* argument. In all other relative path cases, if the *fd* argument does not refer to a valid directory, the function returns ENOTDIR. If the *flag* argument is set to AT_SYMLINK_NOFOLLOW, the function will not automatically traverse a symbolic link at the position of the path. The `fstatat()` function is a multi-purpose function that can be used in place of `stat()`, `lstat()`, or `fstat()`. See `stat(2)`.

The function call `stat(path, buf)` is identical to `fstatat(AT_FDCWD, path, buf, 0)`.

The function call `lstat(path, buf)` is identical to `fstatat(AT_FDCWD, path, buf, AT_SYMLINK_NOFOLLOW)`

The function call `fstat(fildes, buf)` is identical to `fstatat(fildes, NULL, buf, 0)`.

Set owner and group ID

```
int fchownat (int fd, const char *path, uid_t owner, gid_t group, int flag)
```

The `fchownat(2)` function sets the owner ID and group ID for a file. If the *path* argument is relative, it is resolved relative to the *fd* argument file descriptor, otherwise the *fd* argument is ignored. If the *fd* argument is a special value AT_FDCWD the path is resolved relative to the current working directory. If the *path* argument is a null pointer, the function sets the owner and group ID of the file referenced by the *fd* argument. In all other relative path cases, if the *fd* argument does not refer to a valid directory, the function returns ENOTDIR. If the *flag* argument is set to AT_SYMLINK_NOFOLLOW, the function will not automatically traverse a symbolic link at the position of the path. The `fchownat()` function is a multi-purpose function that can be used in place of `chown()`, `lchown()`, or `fchown()`. See `chown(2)`.

The function call `chown(path, owner, group)` is equivalent to `fchownat(AT_FDCWD, path, owner, group, 0)`.

The function call `lchown(path, owner, group)` is equivalent to `fchownat(AT_FDCWD, path, owner, group, AT_SYMLINK_NOFOLLOW)`.

Set file access and modification times

```
int futimesat (int fd, const char *path, const struct timeval times[2])
```

The `futimesat(2)` function sets the access and modification times for a file. If the *path* argument is relative, it is resolved relative to the *fd* argument file descriptor; otherwise the *fd* argument is ignored. If the *fd* argument is the special value AT_FDCWD, the path is resolved relative to the current working directory. If the *path* argument is a null pointer, the function sets the access and modification times of the file referenced by the *fd* argument. In all other relative path cases, if the *fd* argument does not refer to a valid directory, the function returns ENOTDIR. The `futimesat()` function can be used in place of `utimes(2)`.

The function call `utimes(path, times)` is equivalent to `futimesat(AT_FDCWD, path, times)`.

New pathconf() functionality

```
long int pathconf(const char *path, int name)
```

Two variables have been added to `pathconf(2)` to provide enhanced support for extended attribute manipulation. The `XATTR_ENABLED` variable allows an application to determine if attribute support is currently enabled for the file in question. The `XATTR_EXISTS` variable allows an application to determine whether there are any extended attributes associated with the supplied path.

Open/Create an attribute file

```
int attropen (const char *path, const char *attrpath, int oflag [, mode_t mode])
```

The `attropen(3C)` function returns a file descriptor for the named attribute, `attrpath`, of the file indicated by `path`. The `oflag` and `mode` arguments are identical to the `open(2)` arguments and are applied to the open operation on the attribute file (for example, using the `O_CREAT` flag creates a new attribute). Once opened, all normal file system operations can be used on the attribute file descriptor. The `attropen()` function is a convenience function and is equivalent to the following sequence of operations:

```
fd = open (path, O_RDONLY);
attrfd = openat (fd, attrpath, oflag|O_XATTR, mode);
close (fd);
```

The set of existing attributes can be browsed by calling `attropen()` with "." as the attribute name. The list of attributes is obtained by calling `getdents(2)` (or `fdopendir(3C)` followed by `readdir(3C)`, see below) on the returned file descriptor.

Convert an open file descriptor for a directory into a directory descriptor

```
DIR * fdopendir (const int fd)
```

The `fdopendir(3C)` function promotes a file descriptor for a directory to a directory pointer suitable for use with the `readdir(3C)` function. The originating file descriptor should not be used again following the call to `fdopendir()`. The directory pointer should be closed with a call to `closedir(3C)`. If the provided file descriptor does not reference a directory, the function returns `ENOTDIR`. This function is useful in circumstances where the only available handle on a directory is a file descriptor. See `attropen(3C)` and `openat(2)`.

Using the API

The following examples demonstrate how the API might be used to perform basic operations on extended attributes:

EXAMPLE 1 List extended attributes on a file.

```
attrdirfd = attropen("test", ".", O_RDONLY);
dirp = fdopendir(attrdirfd);
while (dp = readdir(dirp)) {
...
}
```

EXAMPLE 2 Open an extended attribute.

```
attrfd = attropen("test", dp->d_name, O_RDONLY);
```

or

```
attrfd = openat(attrdirfd, dp->d_name, O_RDONLY);
```

EXAMPLE 3 Read from an extended attribute.

```
while (read(attrfd, buf, 512) > 0) {
    ...
}
```

EXAMPLE 4 Create an extended attribute.

```
newfd = attropen("test", "attr", O_CREAT|O_RDWR);
```

or

```
newfd = openat(attrdirfd, "attr", O_CREAT|O_RDWR);
```

EXAMPLE 5 Write to an extended attribute.

```
count = write(newfd, buf, length);
```

EXAMPLE 6 Delete an extended attribute.

```
error = unlinkat(attrdirfd, "attr");
```

Applications intending to access the interfaces defined here as well as the POSIX and X/Open specification-conforming interfaces should define the macro `_ATFILE_SOURCE` to be 1 and set whichever feature test macros are appropriate to obtain the desired environment. See `standards(5)`.

Extended Archive Formats

As noted above in the description of command utilities modified to provide support for extended attributes, the archive formats for `tar(1)` and `cpio(1)` have been extended to provide support for archiving extended attributes. This section describes the specifics of the archive format extensions.

Extended tar format

The `tar` archive is made up of a series of 512 byte blocks. Each archived file is represented by a header block and zero or more data blocks containing the file contents. The header block is structured as shown in the following table.

Field Name	Length (in Octets)	Description
Name	100	File name string
Mode	8	12 file mode bits

fsattr(5)

Field Name	Length (in Octets)	Description
Uid	8	User ID of file owner
Gid	8	Group ID of file owner
Size	12	Size of file
Mtime	12	File modification time
Chksum	8	File contents checksum
Typeflag	1	File type flag
Linkname	100	Link target name if file linked
Magic	6	"ustar"
Version	2	"00"
Uname	32	User name of file owner
Gname	32	Group name of file owner
Devmajor	8	Major device ID if special file
Devminor	8	Minor device ID if special file
Prefix	155	Path prefix string for file

The extended attribute project extends the above header format by defining a new header type (for the `Typeflag` field). The type 'E' is defined to be used for all extended attribute files. Attribute files are stored in the `tar` archive as a sequence of two `<header , data>` pairs. The first file contains the data necessary to locate and name the extended attribute in the file system. The second file contains the actual attribute file data. Both files use an 'E' type header. The prefix and name fields in extended attribute headers are ignored, though they should be set to meaningful values for the benefit of archivers that do not process these headers. Solaris archivers set the prefix field to `"/dev/null"` to prevent archivers that do not understand the type 'E' header from trying to restore extended attribute files in inappropriate places.

Extended cpio format

The `cpio` archive format is octet-oriented rather than block-oriented. Each file entry in the archive includes a header that describes the file, followed by the file name, followed by the contents of the file. These data are arranged as described in the following table.

Field Name	Length (in Octets)	Description
<code>c_magic</code>	6	70707
<code>c_dev</code>	6	First half of unique file ID

Field Name	Length (in Octets)	Description
c_ino	6	Second half of unique file ID
c_mode	6	File mode bits
c_uid	6	User ID of file owner
c_gid	6	Group ID of file owner
c_nlink	6	Number of links referencing file
c_rdev	6	Information for special files
c_mtime	11	Modification time of file
c_namesize	6	Length of file pathname
c_filesize	11	Length of file content
c_name	c_namesize	File pathname
c_filedata	c_filesize	File content

The basic archive file structure is not changed for extended attributes. The file type bits stored in the `c_mode` field for an attribute file are set to `0xB000`. As with the `tar` archive format, extended attributes are stored in `cpio` archives as two consecutive file entries. The first file describes the location/name for the extended attribute. The second file contains the actual attribute file content. The `c_name` field in extended attribute headers is ignored, though it should be set to a meaningful value for the benefit of archivers that do not process these headers. Solaris archivers start the pathname with `"/dev/null/"` to prevent archivers that do not understand the type 'E' header from trying to restore extended attribute files in inappropriate places.

Attribute identification data format

Both the `tar` and `cpio` archive formats can contain the special files described above, always paired with the extended attribute data record, for identifying the precise location of the extended attribute. These special data files are necessary because there is no simple naming mechanism for extended attribute files. Extended attributes are not visible in the file system name space. The extended attribute name space must be "tunneled into" using the `openat()` function. The attribute identification data must support not only the flat naming structure for extended attributes, but also the possibility of future extensions allowing for attribute directory hierarchies and recursive attributes. The data file is therefore composed of a sequence of records. It begins with a fixed length header describing the content. The following table describes the format of this data file.

fsattr(5)

Field Name	Length (in Octets)	Description
h_version	7	Name file version
h_size	10	Length of data file
h_component_len	10	Total length of all path segments
h_link_comp_len	10	Total length of all link segments
path	h_component_len	Complex path
link_path	h_link_comp_len	Complex link path

As demonstrated above, the header is followed by a record describing the "path" to the attribute file. This path is composed of two or more path segments separated by a null character. Each segment describes a path rooted at the hidden extended attribute directory of the leaf file of the previous segment, making it possible to name attributes on attributes. The first segment is always the path to the parent file that roots the entire sequence in the normal name space. The following table describes the format of each segment.

Field Name	Length (in Octets)	Description
h_namesz	7	Length of segment path
h_typeflag	1	Actual file type of attribute file
h_names	h_namesz	Parent path + segment path

If the attribute file is linked to another file, the path record is followed by a second record describing the location of the referencing file. The structure of this record is identical to the record described above.

SEE ALSO cp(1), cpio(1), find(1), ls(1), mv(1), pax(1), runat(1), tar(1), du(1), fsck(1M), chown(2), link(2), open(2), pathconf(2), rename(2), stat(2), unlink(2), utimes(2), attropen(3C), standards(5)

NAME iconv_1250 – code set conversion tables for MS 1250 (Windows Latin 2)

DESCRIPTION The following code set conversions are supported:

Code Set Conversions Supported				
Code	Symbol	Target Code	Symbol	Target Output
MS 1250	win2	ISO 8859-2	iso2	ISO Latin 2
MS 1250	win2	MS 852	dos2	MS-DOS Latin 2
MS 1250	win2	Mazovia	maz	Mazovia
MS 1250	win2	DHN	dhn	Dom Handlowy Nauki

CONVERSIONS The conversions are performed according to the following tables. All values in the tables are given in octal.

MS 1250 to ISO 8859-2 For the conversion of MS 1250 to ISO 8859-2, all characters not in the following table are mapped unchanged.

Conversions Performed			
MS 1250	ISO 8859-2	MS 1250	ISO 8859-2
24-211	40	235	273
212	251	236	276
213	40	237	274
214	246	241	267
215	253	245	241
216	256	246-267	40
217	254	271	261
221-231	40	273	40
232	271	274	245
233	40	276	265
234	266	247	365

MS 1250 to MS 852 For the conversion of MS 1250 to MS 852, all characters not in the following table are mapped unchanged.

iconv_1250(5)

Conversions Performed			
MS 1250	MS 852	MS 1250	MS 852
200-211	40	311	220
212	346	312	250
213	40	313	323
214	227	314	267
215	233	315	326
216	246	316	327
217	215	317	322
220-231	40	320	321
232	347	321	343
233	40	322	325
234	230	323	340
235	234	324	342
236	247	325	212
237	253	326	231
240	377	327	236
241	363	330	374
242	364	331	336
243	235	332	351
244	317	333	353
245	244	334	232
246	40	335	355
247	365	336	335
250	371	337	341
251	40	340	352
252	270	341	240
253	256	342	203
254	252	343	307
255	360	344	204

Conversions Performed			
MS 1250	MS 852	MS 1250	MS 852
256	40	345	222
257	275	346	206
260	370	347	207
261	40	350	237
262	362	351	202
263	210	352	251
264	357	353	211
265-267	40	354	330
270	367	355	241
271	245	356	214
272	255	357	324
273	257	360	320
274	225	361	344
275	361	362	345
276	226	363	242
277	276	364	223
300	350	365	213
301	265	366	224
302	266	367	366
303	306	370	375
304	216	371	205
305	221	372	243
306	217	374	201
307	200	375	354
310	254	376	356

MS 1250 to Mazovia

For the conversion of MS 1250 to Mazovia, all characters not in the following table are mapped unchanged.

iconv_1250(5)

Conversions Performed			
MS 1250	Mazovia	MS 1250	Mazovia
200-213	40	310-311	40
214	230	312	220
215-216	40	313-320	40
217	240	321	245
220-233	40	322	40
234	236	323	243
235-236	40	324-325	40
237	246	326	231
240	377	327-333	40
241-242	40	334	232
243	234	335-336	40
244	40	337	341
245	217	340-341	40
246-252	40	342	203
253	256	343	40
254	252	344	204
255-256	40	345	40
257	241	346	215
260	370	347	207
261	361	350	40
262	40	351	202
263	222	352	221
264	40	353	211
265	346	354-355	40
266	40	356	214
267	372	357-360	40
270	40	361	244
271	206	362	40

Conversions Performed			
MS 1250	Mazovia	MS 1250	Mazovia
272	40	363	242
273	257	364	223
274-276	40	365	40
277	247	366	224
300-303	40	367	366
304	216	370-373	40
305	40	374	201
306	225	375-376	40
307	200		

MS 1250 to DHN For the conversion of MS 1250 to DHN, all characters not in the following table are mapped unchanged.

Conversions Performed			
MS 1250	DHN	MS 1250	DHN
200-213	40	306	201
214	206	307-311	40
215-216	40	312	202
217	207	313-320	40
220-233	40	321	204
234	217	322	40
235-236	40	323	205
237	220	324-325	40
240	377	326	231
241-242	40	327-333	40
243	203	334	232
244	40	335-336	40
245	200	337	341
246-252	40	340	40

iconv_1250(5)

Conversions Performed			
MS 1250	DHN	MS 1250	DHN
253	256	341	240
254	252	342-345	40
255-256	40	346	212
257	210	347-351	40
260	370	352	213
261	361	353-354	40
262	40	355	241
263	214	356-360	40
264	40	361	215
265	346	362	40
266	40	363	216
267	372	364	223
270	40	365	40
271	211	366	224
272	40	367	366
273	257	370-371	40
274-276	40	372	243
277	221	373-376	40
300-305	40		

FILES /usr/lib/iconv/*.so conversion modules
/usr/lib/iconv/*.t conversion tables
/usr/lib/iconv/iconv_data list of conversions supported by conversion tables

SEE ALSO iconv(1), iconv(3C), iconv(5)

NAME iconv_1251 – code set conversion tables for MS 1251 (Windows Cyrillic)

DESCRIPTION The following code set conversions are supported:

Code Set Conversions Supported				
Code	Symbol	Target Code	Symbol	Target Output
MS 1251	win5	ISO 8859-5	iso5	ISO 8859-5 Cyrillic
MS 1251	win5	KOI8-R	koi8	KOI8-R
MS 1251	win5	PC Cyrillic	alt	Alternative PC Cyrillic
MS 1251	win5	Mac Cyrillic	mac	Macintosh Cyrillic

CONVERSIONS The conversions are performed according to the following tables. All values in the tables are given in octal.

MS 1251 to ISO 8859-5 For the conversion of MS 1251 to ISO 8859-5, all characters not in the following table are mapped unchanged.

Conversions Performed			
MS 1251	ISO 8859-5	MS 1251	ISO 8859-5
24	4	310	270
200	242	311	271
201	243	312	272
202	40	313	273
203	363	314	274
204-207	40	315	275
210	255	316	276
211	40	317	277
212	251	320	300
213	40	321	301
214	252	322	302
215	254	323	303
216	253	324	304
217	257	325	305

iconv_1251(5)

Conversions Performed			
MS 1251	ISO 8859-5	MS 1251	ISO 8859-5
220	362	326	306
221-227	40	327	307
230	255	330	310
231	40	331	311
232	371	332	312
233	40	333	313
234	372	334	314
235	374	335	315
236	373	336	316
237	377	337	317
241	256	340	320
242	376	341	321
243	250	342	322
244-247	40	343	323
250	241	344	324
251	40	345	325
252	244	346	326
253-254	40	347	327
255	55	350	330
256	40	351	331
257	247	352	332
260-261	40	353	333
262	246	354	334
263	366	355	335
264-267	40	356	336
270	361	357	337
271	360	360	340
272	364	361	341

Conversions Performed			
MS 1251	ISO 8859-5	MS 1251	ISO 8859-5
273	40	362	342
274	370	363	343
275	245	364	344
276	365	365	345
277	367	366	346
300	260	367	347
301	261	370	350
302	262	371	351
303	263	372	352
304	264	373	353
305	265	374	354
306	266	375	355
307	267	376	356

**MS 1251 to
KOI8-R**

For the conversion of MS 1251 to KOI8-R, all characters not in the following table are mapped unchanged.

Conversions Performed			
MS 1251	KOI8-R	MS 1251	KOI8-R
24	4	310	351
200	261	311	352
201	262	312	353
202	40	313	354
203	242	314	355
204-207	40	315	356
210	255	316	357
211	40	317	360
212	271	320	362
213	40	321	363

iconv_1251(5)

Conversions Performed			
MS 1251	KO18-R	MS 1251	KO18-R
214	272	322	364
215	274	323	365
216	273	324	346
217	277	325	350
220	241	326	343
221-227	40	327	376
230	255	330	373
231	40	331	375
232	251	332	377
233	40	333	371
234	252	334	370
235	254	335	374
236	253	336	340
237	257	337	361
241	276	340	301
242	256	341	302
243	270	342	327
244-247	40	343	307
250	263	344	304
251	40	345	305
252	264	346	326
253-254	40	347	332
255	55	350	311
256	40	351	312
257	267	352	313
260-261	40	353	314
262	266	354	315
263	246	355	316

Conversions Performed			
MS 1251	KOI8-R	MS 1251	KOI8-R
264-267	40	356	317
270	243	357	320
271	260	360	322
272	244	361	323
273	40	362	324
274	250	363	325
275	265	364	306
276	245	365	310
277	247	366	303
300	341	367	336
301	342	370	333
302	367	371	335
303	347	372	337
304	344	373	331
305	345	374	330
306	366	375	334
307	372	376	300

MS 1251 to PC Cyrillic

For the conversion of MS 1251 to PC Cyrillic, all characters not in the following table are mapped unchanged.

Conversions Performed			
MS 1251	PC Cyrillic	MS 1251	PC Cyrillic
24	4	332	232
200-207	40	333	233
210	260	334	234
211-227	40	335	235
230	260	336	236
231-247	40	337	237

iconv_1251(5)

Conversions Performed			
MS 1251	PC Cyrillic	MS 1251	PC Cyrillic
250	360	340	240
251-254	40	341	241
255	55	342	242
256-267	40	343	243
270	361	344	244
271-277	40	345	245
300	200	346	246
301	201	347	247
302	202	350	250
303	203	351	251
304	204	352	252
305	205	353	253
306	206	354	254
307	207	355	255
310	210	356	256
311	211	357	257
312	212	360	340
313	213	361	341
314	214	362	342
315	215	363	343
316	216	364	344
317	217	365	345
320	220	366	346
321	221	367	347
322	222	370	350
323	223	371	351
324	224	372	352
325	225	373	353

Conversions Performed			
MS 1251	PC Cyrillic	MS 1251	PC Cyrillic
326	226	374	354
327	227	375	355
330	230	376	356
331	231		

**MS 1251 to Mac
Cyrillic**

For the conversion of MS 1251 to Mac Cyrillic, all characters not in the following table are mapped unchanged.

Conversions Performed			
MS 1251	Mac Cyrillic	MS 1251	Mac Cyrillic
24	4	260	241
200	253	262	247
201	256	263	264
202	40	264	266
203	257	266	246
204	327	267	245
205	311	270	336
206	240	271	334
207-211	40	272	271
212	274	273	310
213	40	274	300
214	276	275	301
215	315	276	317
216	40	277	273
217	332	300	200
220	254	301	201
221	324	302	202
222	325	303	203
223	322	304	204

iconv_1251(5)

Conversions Performed			
MS 1251	Mac Cyrillic	MS 1251	Mac Cyrillic
224	323	305	205
225	40	306	206
226	320	307	207
227	321	310	210
230	40	311	211
231	252	312	212
232	275	313	213
233	40	314	214
234	277	315	215
235	316	316	216
236	40	317	217
237	333	320	220
240	312	321	221
241	330	322	222
242	331	323	223
243	267	324	224
244	377	325	225
245	242	326	226
246	40	327	227
247	244	330	230
250	335	331	231
252	270	332	232
253	307	333	233
254	302	334	234
255	55	335	235
256	250	336	236
257	272	337	237
355	316		

iconv_1251(5)

FILES /usr/lib/iconv/*.so conversion modules
/usr/lib/iconv/*.t conversion tables
/usr/lib/iconv/iconv_data list of conversions supported by conversion tables

SEE ALSO iconv(1), iconv(3C), iconv(5)

iconv(5)

NAME	iconv – code set conversion tables				
DESCRIPTION	The following code set conversions are supported:				
	Code Set Conversions Supported				
Code	Symbol	Target Code	Symbol	Target Output	
ISO 646	646	ISO 8859-1	8859	US ASCII	
ISO 646de	646de	ISO 8859-1	8859	German	
ISO 646da	646da	ISO 8859-1	8859	Danish	
ISO 646en	646en	ISO 8859-1	8859	English ASCII	
ISO 646es	646es	ISO 8859-1	8859	Spanish	
ISO 646fr	646fr	ISO 8859-1	8859	French	
ISO 646it	646it	ISO 8859-1	8859	Italian	
ISO 646sv	646sv	ISO 8859-1	8859	Swedish	
ISO 8859-1	8859	ISO 646	646	7 bit ASCII	
ISO 8859-1	8859	ISO 646de	646de	German	
ISO 8859-1	8859	ISO 646da	646da	Danish	
ISO 8859-1	8859	ISO 646en	646en	English ASCII	
ISO 8859-1	8859	ISO 646es	646es	Spanish	
ISO 8859-1	8859	ISO 646fr	646fr	French	
ISO 8859-1	8859	ISO 646it	646it	Italian	
ISO 8859-1	8859	ISO 646sv	646sv	Swedish	
ISO 8859-16	iso16	ISO 8859-2	iso2	ISO Latin 2	
ISO 8859-2	iso2	ISO 8859-16	iso16	ISO Latin 10	
ISO 8859-16	iso16	IBM 850	ibm850	IBM 850 code page	
ISO 8859-16	iso16	IBM 870	ibm870	IBM 870 code page	
ISO 8859-2	iso2	MS 1250	win2	Windows Latin 2	
ISO 8859-2	iso2	MS 852	dos2	MS-DOS Latin 2	
ISO 8859-2	iso2	Mazovia	maz	Mazovia	
IBM 850	ibm850	ISO 8859-16	iso16	ISO Latin 10	
IBM 870	ibm870	ISO 8859-16	iso16	ISO Latin 10	
MS 1250	win2	DHN	dhn	Dom Handlowy Nauki	
MS 852	dos2	ISO 8859-2	iso2	ISO Latin 2	
MS 852	dos2	MS 1250	win2	Windows Latin 2	
MS 852	dos2	Mazovia	maz	Mazovia	
MS 852	dos2	DHN	dhn	Dom Handlowy Nauki	
Mazovia	maz	ISO 8859-2	iso2	ISO Latin 2	
Mazovia	maz	MS 1250	win2	Windows Latin 2	
Mazovia	maz	MS 852	dos2	MS-DOS Latin 2	
Mazovia	maz	DHN	dhn	Dom Handlowy Nauki	
DHN	dhn	ISO 8859-2	iso2	ISO Latin 2	
DHN	dhn	MS 1250	win2	Windows Latin 2	
DHN	dhn	MS 852	dos2	MS-DOS Latin 2	
DHN	dhn	Mazovia	maz	Mazovia	
ISO 8859-5	iso5	KOI8-R	koi8	KOI8-R	
ISO 8859-5	iso5	PC Cyrillic	alt	Alternative PC Cyrillic	
ISO 8859-5	iso5	MS 1251	win5	Windows Cyrillic	
ISO 8859-5	iso5	Mac Cyrillic	mac	Macintosh Cyrillic	
KOI8-R	koi8	ISO 8859-5	iso5	ISO 8859-5 Cyrillic	
KOI8-R	koi8	PC Cyrillic	alt	Alternative PC Cyrillic	
KOI8-R	koi8	MS 1251	win5	Windows Cyrillic	
KOI8-R	koi8	Mac Cyrillic	mac	Macintosh Cyrillic	
PC Cyrillic	alt	ISO 8859-5	iso5	ISO 8859-5 Cyrillic	
PC Cyrillic	alt	KOI8-R	koi8	KOI8-R	
PC Cyrillic	alt	MS 1251	win5	Windows Cyrillic	

PC Cyrillic	alt	Mac Cyrillic	mac	Macintosh Cyrillic
MS 1251	win5	ISO 8859-5	iso5	ISO 8859-5 Cyrillic
MS 1251	win5	KOI8-R	koi8	KOI8-R
MS 1251	win5	PC Cyrillic	alt	Alternative PC Cyrillic
MS 1251	win5	Mac Cyrillic	mac	Macintosh Cyrillic
Mac Cyrillic	mac	ISO 8859-5	iso5	ISO 8859-5 Cyrillic
Mac Cyrillic	mac	KOI8-R	koi8	KOI8-R
Mac Cyrillic	mac	PC Cyrillic	alt	Alternative PC Cyrillic
Mac Cyrillic	mac	MS 1251	win5	Windows Cyrillic

CONVERSIONS

The conversions are performed according to the tables contained in the manual pages cross-referenced in the Index of Conversion Code Tables below.

Index of Conversion Code Tables		
Code	Target Code	See Manual Page
ISO 646	ISO 8859-1	iconv_646 (5)
ISO 646de	ISO 8859-1	
ISO 646da	ISO 8859-1	
ISO 646en	ISO 8859-1	
ISO 646es	ISO 8859-1	
ISO 646fr	ISO 8859-1	
ISO 646it	ISO 8859-1	
ISO 646sv	ISO 8859-1	
ISO 8859-1	ISO 646	iconv_8859-1 (5)
ISO 8859-1	ISO 646de	
ISO 8859-1	ISO 646da	
ISO 8859-1	ISO 646en	
ISO 8859-1	ISO 646es	
ISO 8859-1	ISO 646fr	
ISO 8859-1	ISO 646it	
ISO 8859-1	ISO 646sv	
ISO 8859-2	MS 1250	iconv_8859-2 (5)
ISO 8859-2	MS 852	
ISO 8859-2	Mazovia	
ISO 8859-2	DHN	

iconv(5)

Index of Conversion Code Tables		
MS 1250	ISO 8859-2	iconv_1250 (5)
MS 1250	MS 852	
MS 1250	Mazovia	
MS 1250	DHN	
MS 852	ISO 8859-2	iconv_852 (5)
MS 852	MS 1250	
MS 852	Mazovia	
MS 852	DHN	
Mazovia	ISO 8859-2	iconv_maz (5)
Mazovia	MS 1250	
Mazovia	MS 852	
Mazovia	DHN	

Index of Conversion Code Tables		
Code	Target Code	See Manual Page
DHN	ISO 8859-2	iconv_dhn (5)
DHN	MS 1250	
DHN	MS 852	
DHN	Mazovia	
ISO 8859-5	KOI8-R	iconv_8859-5 (5)
ISO 8859-5	PC Cyrillic	
ISO 8859-5	MS 1251	
ISO 8859-5	Mac Cyrillic	
KOI8-R	ISO 8859-5	iconv_koi8-r (5)
KOI8-R	PC Cyrillic	
KOI8-R	MS 1251	
KOI8-R	Mac Cyrillic	
PC Cyrillic	ISO 8859-5	iconv_pc_cyr (5)
PC Cyrillic	KOI8-R	

Index of Conversion Code Tables		
PC Cyrillic	MS 1251	
PC Cyrillic	Mac Cyrillic	
MS 1251	ISO 8859-5	iconv_1251 (5)
MS 1251	KOI8-R	
MS 1251	PC Cyrillic	
MS 1251	Mac Cyrillic	
Mac Cyrillic	ISO 8859-5	iconv_mac_cyr (5)
Mac Cyrillic	KOI8-R	
Mac Cyrillic	PC Cyrillic	
Mac Cyrillic	MS 1251	

FILES /usr/lib/iconv/*.so
conversion modules

/usr/lib/iconv/*.t
Conversion tables.

/usr/lib/iconv/geniconvtbl/binarytables/*.bt
Conversion binary tables.

/usr/lib/iconv/iconv_data
List of conversions supported by conversion tables.

SEE ALSO iconv(1), iconv(3C), iconv_1250(5), iconv_1251(5), iconv_646(5),
iconv_852(5), iconv_8859-1(5), iconv_8859-2(5), iconv_8859-5(5),
iconv_dhn(5), iconv_koi8-r(5), iconv_mac_cyr(5), iconv_maz(5),
iconv_pc_cyr(5), iconv_unicode(5)

iconv_646(5)

NAME iconv_646 – code set conversion tables for ISO 646
DESCRIPTION The following code set conversions are supported:

Code Set Conversions Supported				
Code	Symbol	Target Code	Symbol	Target Output
ISO 646	646	ISO 8859-1	8859	US ASCII
ISO 646de	646de	ISO 8859-1	8859	German
ISO 646da	646da	ISO 8859-1	8859	Danish
ISO 646en	646en	ISO 8859-1	8859	English ASCII
ISO 646es	646es	ISO 8859-1	8859	Spanish
ISO 646fr	646fr	ISO 8859-1	8859	French
ISO 646it	646it	ISO 8859-1	8859	Italian
ISO 646sv	646sv	ISO 8859-1	8859	Swedish

CONVERSIONS The conversions are performed according to the following tables. All values in the tables are given in octal.

ISO 646 (US ASCII) to ISO 8859-1 For the conversion of ISO 646 to ISO 8859-1, all characters in ISO 646 can be mapped unchanged to ISO 8859-1

ISO 646de (GERMAN) to ISO 8859-1 For the conversion of ISO 646de to ISO 8859-1, all characters not in the following table are mapped unchanged.

Conversions Performed			
ISO 646de	ISO 8859-1	ISO 646de	ISO 8859-1
100	247	173	344
133	304	174	366
134	326	175	374
135	334	176	337

ISO 646da (DANISH) to ISO 8859-1 For the conversion of ISO 646da to ISO 8859-1, all characters not in the following table are mapped unchanged.

Conversions Performed			
ISO 646da	ISO 8859-1	ISO 646da	ISO 8859-1
133	306	173	346
134	330	174	370
135	305	175	345

**ISO 646en
(ENGLISH ASCII)
to ISO 8859-1**

For the conversion of ISO 646en to ISO 8859-1, all characters not in the following table are mapped unchanged.

Conversions Performed	
ISO 646en	ISO 8859-1
043	243

**ISO 646es
(SPANISH) to ISO
8859-1**

For the conversion of ISO 646es to ISO 8859-1, all characters not in the following table are mapped unchanged.

Conversions Performed			
ISO 646es	ISO 8859-1	ISO 646es	ISO 8859-1
100	247	173	260
133	241	174	361
134	321	175	347
135	277		

**ISO 646fr
(FRENCH) to ISO
8859-1**

For the conversion of ISO 646fr to ISO 8859-1, all characters not in the following table are mapped unchanged.

Conversions Performed			
ISO 646fr	ISO 8859-1	ISO 646fr	ISO 8859-1
043	243	173	351
100	340	174	371
133	260	175	350
134	347	176	250
135	247		

iconv_646(5)

**ISO 646it
(ITALIAN) to ISO
8859-1**

For the conversion of ISO 646it to ISO 8859-1, all characters not in the following table are mapped unchanged.

Conversions Performed			
ISO 646it	ISO 8859-1	ISO 646it	ISO 8859-1
043	243	140	371
100	247	173	340
133	260	174	362
134	347	175	350
135	351	176	354

**ISO 646sv
(SWEDISH) to ISO
8859-1**

For the conversion of ISO 646sv to ISO 8859-1, all characters not in the following table are mapped unchanged.

Conversions Performed			
ISO 646sv	ISO 8859-1	ISO 646sv	ISO 8859-1
100	311	140	351
133	304	173	344
134	326	174	366
135	305	175	345
136	334	176	374

FILES /usr/lib/iconv/*.so conversion modules
/usr/lib/iconv/*.t conversion tables
/usr/lib/iconv/iconv_data list of conversions supported by conversion tables

SEE ALSO iconv(1), iconv(3C), iconv(5)

NAME iconv_852 – code set conversion tables for MS 852 (MS-DOS Latin 2)

DESCRIPTION The following code set conversions are supported:

Code Set Conversions Supported				
Code	Symbol	Target Code	Symbol	Target Output
MS 852	dos2	ISO 8859-2	iso2	ISO Latin 2
MS 852	dos2	MS 1250	win2	Windows Latin 2
MS 852	dos2	Mazovia	maz	Mazovia
MS 852	dos2	DHN	dhn	Dom Handlowy Nauki

CONVERSIONS The conversions are performed according to the following tables. All values in the tables are given in octal.

MS 852 to ISO 8859-2 For the conversion of MS 852 to ISO 8859-2, all characters not in the following table are mapped unchanged.

Conversions Performed			
MS 852	ISO 8859-2	MS 852	ISO 8859-2
24-177	40	271-274	40
200	307	275	257
201	374	276	277
202	351	277-305	40
203	342	306	303
204	344	307	343
205	371	310-316	40
206	346	317	244
207	347	320	360
210	263	321	320
211	353	322	317
212	325	323	313
213	365	324	357
214	356	325	322

iconv_852(5)

Conversions Performed			
MS 852	ISO 8859-2	MS 852	ISO 8859-2
215	254	326	315
216	304	327	316
217	306	330	354
220	311	331-334	40
221	305	335	336
222	345	336	331
223	364	337	40
224	366	340	323
225	245	341	337
226	265	342	324
227	246	343	321
230	266	344	361
231	326	345	362
232	334	346	251
233	253	347	271
234	273	350	300
235	243	351	332
236	327	352	340
237	350	353	333
240	341	354	375
241	355	355	335
242	363	356	376
243	372	357	264
244	241	360	255
245	261	361	275
246	256	362	262
247	276	363	267
250	312	364	242

Conversions Performed			
MS 852	ISO 8859-2	MS 852	ISO 8859-2
251	352	365	247
252	40	366	367
253	274	367	270
254	310	370	260
255	272	371	250
256-264	40	372	377
265	301	374	330
266	302	375	370
267	314	376	40
270	252		

MS 852 to MS 1250

For the conversion of MS 852 to MS 1250, all characters not in the following table are mapped unchanged.

Conversions Performed			
MS 852	MS 1250	MS 852	MS 1250
200	307	270	252
201	374	271-274	40
202	351	275	257
203	342	276	277
204	344	277-305	40
205	371	306	303
206	346	307	343
207	347	310-316	40
210	263	317	244
211	353	320	360
212	325	321	320
213	365	322	317
214	356	323	313

iconv_852(5)

Conversions Performed			
MS 852	MS 1250	MS 852	MS 1250
215	217	324	357
216	304	325	322
217	306	326	315
220	311	327	316
221	305	330	354
222	345	331-334	40
223	364	335	336
224	366	336	331
225	274	337	40
226	276	340	323
227	214	341	337
230	234	342	324
231	326	343	321
232	334	344	361
233	215	345	362
234	235	346	212
235	243	347	232
236	327	350	300
237	350	351	332
240	341	352	340
241	355	353	333
242	363	354	375
243	372	355	335
244	245	356	376
245	271	357	264
246	216	360	255
247	236	361	275
250	312	362	262

Conversions Performed			
MS 852	MS 1250	MS 852	MS 1250
251	352	363	241
252	254	364	242
253	237	365	247
254	310	366	367
255	272	367	270
256	253	370	260
257	273	371	250
260-264	40	372	377
265	301	374	330
266	302	375	370
267	314	376	40

MS 852 to Mazovia

For the conversion of MS 852 to Mazovia, all characters not in the following table are mapped unchanged.

Conversions Performed			
MS 852	Mazovia	MS 852	Mazovia
205	40	246-247	40
206	215	250	220
210	222	251	221
212-213	40	253	246
215	240	254-270	40
217	225	275	241
220-226	40	276	247
227	230	306-336	40
230	236	340	243
233-234	40	342	40
235	234	343	245
236-243	40	344	244

iconv_852(5)

Conversions Performed			
MS 852	Mazovia	MS 852	Mazovia
244	217	345-375	40
245	206		

MS 852 to DHN

For the conversion of MS 852 to DHN, all characters not in the following table are mapped unchanged.

Conversions Performed			
MS 852	DHN	MS 852	DHN
200-205	40	244	200
206	212	245	211
207	40	246-247	40
210	214	250	202
211-214	40	251	213
215	207	253	220
216	40	254-270	40
217	201	275	210
220-226	40	276	221
227	206	306-336	40
230	217	340	205
233-234	40	342	40
235	203	343	204
236-237	40	344	215
242	216	345-375	40
252	254		

FILES /usr/lib/iconv/*.so conversion modules
 /usr/lib/iconv/*.t conversion tables
 /usr/lib/iconv/iconv_data list of conversions supported by conversion tables

SEE ALSO iconv(1), iconv(3C), iconv(5)

NAME	iconv_8859-1 – code set conversion tables for ISO 8859-1 (Latin 1)																																																		
DESCRIPTION	The following code set conversions are supported:																																																		
	<table border="1"> <thead> <tr> <th colspan="5">Code Set Conversions Supported</th> </tr> <tr> <th>Code</th> <th>Symbol</th> <th>Target Code</th> <th>Symbol</th> <th>Target Output</th> </tr> </thead> <tbody> <tr> <td>ISO 8859-1</td> <td>8859</td> <td>ISO 646</td> <td>646</td> <td>7 bit ASCII</td> </tr> <tr> <td>ISO 8859-1</td> <td>8859</td> <td>ISO 646de</td> <td>646de</td> <td>German</td> </tr> <tr> <td>ISO 8859-1</td> <td>8859</td> <td>ISO 646da</td> <td>646da</td> <td>Danish</td> </tr> <tr> <td>ISO 8859-1</td> <td>8859</td> <td>ISO 646en</td> <td>646en</td> <td>English ASCII</td> </tr> <tr> <td>ISO 8859-1</td> <td>8859</td> <td>ISO 646es</td> <td>646es</td> <td>Spanish</td> </tr> <tr> <td>ISO 8859-1</td> <td>8859</td> <td>ISO 646fr</td> <td>646fr</td> <td>French</td> </tr> <tr> <td>ISO 8859-1</td> <td>8859</td> <td>ISO 646it</td> <td>646it</td> <td>Italian</td> </tr> <tr> <td>ISO 8859-1</td> <td>8859</td> <td>ISO 646sv</td> <td>646sv</td> <td>Swedish</td> </tr> </tbody> </table>	Code Set Conversions Supported					Code	Symbol	Target Code	Symbol	Target Output	ISO 8859-1	8859	ISO 646	646	7 bit ASCII	ISO 8859-1	8859	ISO 646de	646de	German	ISO 8859-1	8859	ISO 646da	646da	Danish	ISO 8859-1	8859	ISO 646en	646en	English ASCII	ISO 8859-1	8859	ISO 646es	646es	Spanish	ISO 8859-1	8859	ISO 646fr	646fr	French	ISO 8859-1	8859	ISO 646it	646it	Italian	ISO 8859-1	8859	ISO 646sv	646sv	Swedish
Code Set Conversions Supported																																																			
Code	Symbol	Target Code	Symbol	Target Output																																															
ISO 8859-1	8859	ISO 646	646	7 bit ASCII																																															
ISO 8859-1	8859	ISO 646de	646de	German																																															
ISO 8859-1	8859	ISO 646da	646da	Danish																																															
ISO 8859-1	8859	ISO 646en	646en	English ASCII																																															
ISO 8859-1	8859	ISO 646es	646es	Spanish																																															
ISO 8859-1	8859	ISO 646fr	646fr	French																																															
ISO 8859-1	8859	ISO 646it	646it	Italian																																															
ISO 8859-1	8859	ISO 646sv	646sv	Swedish																																															
CONVERSIONS	The conversions are performed according to the following tables. All values in the tables are given in octal.																																																		
ISO 8859-1 to ISO 646 (7-bit ASCII)	<p>For the conversion of ISO 8859-1 to ISO 646, all characters not in the following table are mapped unchanged.</p> <p>Converted to Underscore '_' (137)</p> <p>-----</p> <pre> 200 201 202 203 204 205 206 207 210 211 212 213 214 215 216 217 220 221 222 223 224 225 226 227 230 231 232 233 234 235 236 237 240 241 242 243 244 245 246 247 250 251 252 253 254 255 256 257 260 261 262 263 264 265 266 267 270 271 272 273 274 275 276 277 300 301 302 303 304 305 306 307 310 311 312 313 314 315 316 317 320 321 322 323 324 325 326 327 330 331 332 333 334 335 336 337 340 341 342 343 344 345 346 347 350 351 352 353 354 355 356 357 360 361 362 363 364 365 366 367 370 371 372 373 374 375 376 377 </pre>																																																		
ISO 8859-1 to ISO 646de (GERMAN)	For the conversion of ISO 8859-1 to ISO 646de, all characters not in the following tables are mapped unchanged.																																																		

iconv_8859-1(5)

Conversions Performed			
ISO 8859-1	ISO 646de	ISO 8859-1	ISO 646de
247	100	337	176
304	133	344	173
326	134	366	174
334	135	374	175

Converted to Underscore '_' (137)

```
-----
100 133 134 135 173 174 175 176
200 201 202 203 204 205 206 207
210 211 212 213 214 215 216 217
220 221 222 223 224 225 226 227
230 231 232 233 234 235 236 237
240 241 242 243 244 245 246
250 251 252 253 254 255 256 257
260 261 262 263 264 265 266 267
270 271 272 273 274 275 276 277
300 301 302 303      305 306 307
310 311 312 313 314 315 316 317
320 321 322 323 324 325      327
330 331 332 333      335 336 337
340 341 342 343      345 346 347
350 351 352 353 354 355 356 357
360 361 362 363 364 365      367
370 371 372 373      375 376 377
```

ISO 8859-1 to ISO 646da (DANISH)

For the conversion of ISO 8859-1 to ISO 646da, all characters not in the following tables are mapped unchanged.

Conversions Performed			
ISO 8859-1	ISO 646da	ISO 8859-1	ISO 646da
305	135	345	175
306	133	346	173
330	134	370	174

Converted to Underscore '_' (137)

```
-----
133 134 135 173 174 175
200 201 202 203 204 205 206 207
210 211 212 213 214 215 216 217
220 221 222 223 224 225 226 227
230 231 232 233 234 235 236 237
240 241 242 243 244 245 246 247
250 251 252 253 254 255 256 257
260 261 262 263 264 265 266 267
270 271 272 273 274 275 276 277
300 301 302 303 304      307
```



```

310 311 312 313 314 315 316 317
320 321 322 323 324 325 326 327
    331 332 333 334 335 336 337
340 341 342 343 344      347
350 351 352 353 354 355 356 357
360 361 362 363 364 365 366 367
371 372 373 374      376 377
    
```

ISO 8859-1 to ISO 646en (ENGLISH ASCII)

For the conversion of ISO 8859-1 to ISO 646en, all characters not in the following tables are mapped unchanged.

Conversions Performed	
ISO 8859-1	ISO 646en
243	043

Converted to Underscore '_' (137)

```

-----
043
200 201 202 203 204 205 206 207
210 211 212 213 214 215 216 217
220 221 222 223 224 225 226 227
230 231 232 233 234 235 236 237
240 241 242      244 245 246 247
250 251 252 253 254 255 256 257
260 261 262 263 264 265 266 267
270 271 272 273 274 275 276 277
300 301 302 303 304 305 306 307
310 311 312 313 314 315 316 317
320 321 322 323 324 325 326 327
330 331 332 333 334 335 336 337
340 341 342 343 344 345 346 347
350 351 352 353 354 355 356 357
360 361 362 363 364 365 366 367
370 371 372 373 374 375 376 377
    
```

ISO 8859-1 to ISO 646fr (FRENCH)

For the conversion of ISO 8859-1 to ISO 646fr, all characters not in the following tables are mapped unchanged.

Conversions Performed			
ISO 8859-1	ISO 646fr	ISO 8859-1	ISO 646fr
243	043	347	134
247	135	350	175
250	176	351	173
260	133	371	174
340	100		

Converted to Underscore '_' (137)

iconv_8859-1(5)

```
043
100 133 134 135 173 174 175 176
200 201 202 203 204 205 206 207
210 211 212 213 214 215 216 217
220 221 222 223 224 225 226 227
230 231 232 233 234 235 236 237
240 241 242      244 245 246
      251 252 253 254 255 256 257
      261 262 263 264 265 266 267
270 271 272 273 274 275 276 277
300 301 302 303 304 305 306 307
310 311 312 313 314 315 316 317
320 321 322 323 324 325 326 327
330 331 332 333 334 335 336 337
      341 342 343 344 345 346
      352 353 354 355 356 357
360 361 362 363 364 365 366 367
370      372 373 374 375 376 377
```

**ISO 8859-1 to ISO
646it (ITALIAN)**

For the conversion of ISO 8859-1 to ISO 646it, all characters not in the following tables are mapped unchanged.

Conversions Performed			
ISO 8859-1	ISO 646it	ISO 8859-1	ISO 646it
243	043	350	175
247	100	351	135
260	133	354	176
340	173	362	174
347	134	371	140

Converted to Underscore '_' (137)

```
-----
043
100 133 134 135 173 174 175 176
200 201 202 203 204 205 206 207
210 211 212 213 214 215 216 217
220 221 222 223 224 225 226 227
230 231 232 233 234 235 236 237
240 241 242      244 245 246
250 251 252 253 254 255 256 257
      261 262 263 264 265 266 267
270 271 272 273 274 275 276 277
300 301 302 303 304 305 306 307
310 311 312 313 314 315 316 317
320 321 322 323 324 325 326 327
330 331 332 333 334 335 336 337
      341 342 343 344 345 346
      352 353 354 355 356 357
360 361 362 363 364 365 366 367
370      372 373 374 375 376 377
```

ISO 8859-1 to ISO 646es (SPANISH)

For the conversion of ISO 8859-1 to ISO 646es, all characters not in the following tables are mapped unchanged.

Conversions Performed			
ISO 8859-1	ISO 646es	ISO 8859-1	ISO 646es
241	133	321	134
247	100	347	175
260	173	361	174
277	135		

Converted to Underscore '_' (137)

```
-----
100 133 134 135 173 174 175
200 201 202 203 204 205 206 207
210 211 212 213 214 215 216 217
220 221 222 223 224 225 226 227
230 231 232 233 234 235 236 237
240   242 243 244 245 246
250 251 252 253 254 255 256 257
    261 262 263 264 265 266 267
270 271 272 273 274 275 276
300 301 302 303 304 305 306 307
310 311 312 313 314 315 316 317
320   322 323 324 325 326 327
330 331 332 333 334 335 336 337
340 341 342 343 344 345 346
350 351 352 353 354 355 356 357
360   362 363 364 365 366 367
370 371 372 373 374 375 376 377
```

ISO 8859-1 to ISO 646sv (SWEDISH)

For the conversion of ISO 8859-1 to ISO 646sv, all characters not in the following tables are mapped unchanged.

Conversions Performed			
ISO 8859-1	ISO 646sv	ISO 8859-1	ISO 646sv
304	133	344	173
305	135	345	175
311	100	351	140
326	134	366	174
334	136	374	176

Converted to Underscore '_' (137)

```
-----
100 133 134 135 136 140
173 174 175 176
```

iconv_8859-1(5)

```
200 201 202 203 204 205 206 207
210 211 212 213 214 215 216 217
220 221 222 223 224 225 226 227
230 231 232 233 234 235 236 237
240 241 242 243 244 245 246 247
250 251 252 253 254 255 256 257
260 261 262 263 264 265 266 267
270 271 272 273 274 275 276 277
300 301 302 303           306 307
310       312 313 314 315 316 317
320 321 322 323 324 325       327
330 331 332 333           335 336 337
340 341 342 343           346 347
350       352 353 354 355 356 357
360 361 362 363 364 365       367
370 371 372 373           375 376 377
```

FILES

/usr/lib/iconv/*.so
/usr/lib/iconv/*.t
/usr/lib/iconv/iconv_data

conversion modules

conversion tables

list of conversions supported by conversion tables

SEE ALSO

iconv(1), iconv(3C), iconv(5)

NAME iconv_8859-2 – code set conversion tables for ISO 8859-2 (Latin 2)

DESCRIPTION The following code set conversions are supported:

Code Set Conversions Supported				
Code	Symbol	Target Code	Symbol	Target Output
ISO 8859-2	iso2	MS 1250	win2	Windows Latin 2
ISO 8859-2	iso2	MS 852	dos2	MS-DOS Latin 2
ISO 8859-2	iso2	Mazovia	maz	Mazovia
ISO 8859-2	iso2	DHN	dhn	Dom Handlowy Nauki

CONVERSIONS The conversions are performed according to the following tables. All values in the tables are given in octal.

ISO 8859-2 to MS 1250 For the conversion of ISO 8859-2 to MS 1250, all characters not in the following table are mapped unchanged.

Conversions Performed			
ISO 8859-2	MS 1250	ISO 8859-2	MS 1250
24	4	261	271
177-237	40	265	276
241	245	266	234
245	274	267	241
246	214	271	232
251	212	273	235
253	215	274	237
254	217	276	236
256	216	266	236

ISO 8859-2 to MS 852 For the conversion of ISO 8859-2 to MS 852, all characters not in the following table are mapped unchanged.

iconv_8859-2(5)

Conversions Performed			
ISO 8859-2	MS 852	ISO 8859-2	MS 852
24	4	316	327
177-237	40	317	322
240	377	320	321
241	244	321	343
242	364	322	325
243	235	323	340
244	317	324	342
245	225	325	212
246	227	326	231
247	365	327	236
250	371	330	374
251	346	331	336
252	270	332	351
253	233	333	353
254	215	334	232
255	360	335	355
256	246	336	335
257	275	337	341
260	370	340	352
261	245	341	240
262	362	342	203
263	210	343	307
264	357	344	204
265	226	345	222
266	230	346	206
267	363	347	207
270	367	350	237
271	347	351	202

Conversions Performed			
ISO 8859-2	MS 852	ISO 8859-2	MS 852
272	255	352	251
273	234	353	211
274	253	354	330
275	361	355	241
276	247	356	214
277	276	357	324
300	350	360	320
301	265	361	344
302	266	362	345
303	306	363	242
304	216	364	223
305	221	365	213
306	217	366	224
307	200	367	366
310	254	370	375
311	220	371	205
312	250	372	243
313	323	374	201
314	267	375	354
315	326	376	356
366	367		

ISO 8859-2 to Mazovia

For the conversion of ISO 8859-2 to Mazovia, all characters not in the following table are mapped unchanged.

Conversions Performed			
ISO 8859-2	Mazovia	ISO 8859-2	Mazovia
24	4	323	243
177-237	40	324-325	40

iconv_8859-2(5)

Conversions Performed			
ISO 8859-2	Mazovia	ISO 8859-2	Mazovia
240	377	326	231
241	217	327-333	40
242	40	334	232
243	234	335-336	40
244-245	40	337	341
246	230	340-341	40
247-253	40	342	203
254	240	343	40
255-256	40	344	204
257	241	345	40
260	370	346	215
261	206	347	207
262	40	350	40
263	222	351	202
264-265	40	352	221
266	236	353	211
267-273	40	354-355	40
274	246	356	214
275-276	40	357-360	40
277	247	361	244
300-303	40	362	40
304	216	363	242
305	40	364	223
306	225	365	40
307	200	366	224
310-311	40	367	366
312	220	370-373	40
313-320	40	374	201

Conversions Performed			
ISO 8859-2	Mazovia	ISO 8859-2	Mazovia
321	245	375-376	40
322	40		

**ISO 8859-2 to
DHN**

For the conversion of ISO 8859-2 to DHN, all characters not in the following table are mapped unchanged.

Conversions Performed			
ISO 8859-2	DHN	ISO 8859-2	DHN
24	4	322	40
177-237	40	323	205
240	377	324-325	40
241	200	326	231
242	40	327-333	40
243	203	334	232
244-245	40	335-336	40
246	206	337	341
247-253	40	340	40
254	207	341	240
255-256	40	342-345	40
257	210	346	212
260	370	347-351	40
261	211	352	213
262	40	353-354	40
263	214	355	241
264-265	40	356-360	40
266	217	361	215
267-273	40	362	40
274	220	363	216
275-276	40	364	223

iconv_8859-2(5)

Conversions Performed			
ISO 8859-2	DHN	ISO 8859-2	DHN
277	221	365	40
300-305	40	366	224
306	201	367	366
307-311	40	370-371	40
312	202	372	243
313-320	40	373-376	40
321	204		

FILES /usr/lib/iconv/*.so conversion modules
/usr/lib/iconv/*.t conversion tables
/usr/lib/iconv/iconv_data list of conversions supported by conversion tables

SEE ALSO iconv(1), iconv(3C), iconv(5)

NAME iconv_8859-5 – code set conversion tables for ISO 8859-5 (Cyrillic)

DESCRIPTION The following code set conversions are supported:

Code Set Conversions Supported				
Code	Symbol	Target Code	Symbol	Target Output
ISO 8859-5	iso5	KOI8-R	koi8	KOI8-R
ISO 8859-5	iso5	PC Cyrillic	alt	Alternative PC Cyrillic
ISO 8859-5	iso5	MS 1251	win5	Windows Cyrillic
ISO 8859-5	iso5	Mac Cyrillic	mac	Macintosh Cyrillic

CONVERSIONS The conversions are performed according to the following tables. All values in the tables are given in octal.

ISO 8859-5 to KOI8-R For the conversion of ISO 8859-5 to KOI8-R, all characters not in the following table are mapped unchanged.

Conversions Performed			
ISO 8859-5	KOI8-R	ISO 8859-5	KOI8-R
24	4	320	301
241	263	321	302
242	261	322	327
243	262	323	307
244	264	324	304
245	265	325	305
246	266	327	332
247	267	330	311
250	270	331	312
251	271	332	313
252	272	333	314
253	273	334	315
254	274	335	316
256	276	336	317

iconv_8859-5(5)

Conversions Performed			
ISO 8859-5	KOI8-R	ISO 8859-5	KOI8-R
257	277	337	320
260	341	340	322
261	342	341	323
262	367	342	324
263	347	343	325
264	344	344	306
265	345	345	310
266	366	346	303
267	372	347	336
270	351	350	333
271	352	351	335
272	353	352	337
273	354	353	331
274	355	354	330
275	356	355	334
276	357	356	300
277	360	357	321
300	362	360	260
301	363	361	243
302	364	362	241
303	365	363	242
304	346	364	244
305	350	365	245
306	343	366	246
307	376	367	247
310	373	370	250
311	375	371	251
312	377	372	252

Conversions Performed			
ISO 8859-5	KOI8-R	ISO 8859-5	KOI8-R
313	371	373	253
314	370	374	254
315	374	375	255
316	340	376	256
317	361		

**ISO 8859-5 to PC
Cyrillic**

For the conversion of ISO 8859-5 to PC Cyrillic, all characters not in the following table are mapped unchanged.

Conversions Performed			
ISO 8859-5	PC Cyrillic	ISO 8859-5	PC Cyrillic
24	4	307	227
200-240	40	310	230
241	360	311	231
242-254	40	312	232
255	260	313	233
256-257	40	314	234
260	200	315	235
261	201	316	236
262	202	317	237
263	203	320	240
264	204	321	241
265	205	322	242
266	206	323	243
267	207	324	244
270	210	325	245
271	211	326	246
272	212	327	247
273	213	330	250

iconv_8859-5(5)

Conversions Performed			
ISO 8859-5	PC Cyrillic	ISO 8859-5	PC Cyrillic
274	214	331	251
275	215	332	252
276	216	333	253
277	217	334	254
300	220	335	255
301	221	336	256
302	222	337	257
303	223	360-374	40
304	224	375	260
305	225	376	40
306	226	365	40

ISO 8859-5 to MS 1251

For the conversion of ISO 8859-5 to MS 1251, all characters not in the following table are mapped unchanged.

Conversions Performed			
ISO 8859-5	MS 1251	ISO 8859-5	MS 1251
24	4	317	337
200-237	40	320	340
241	250	321	341
242	200	322	342
243	201	323	343
244	252	324	344
245	275	325	345
246	262	326	346
247	257	327	347
250	243	330	350
251	212	331	351
252	214	332	352

Conversions Performed			
ISO 8859-5	MS 1251	ISO 8859-5	MS 1251
253	216	333	353
254	215	334	354
255	210	335	355
256	241	336	356
257	217	337	357
260	300	340	360
261	301	341	361
262	302	342	362
263	303	343	363
264	304	344	364
265	305	345	365
266	306	346	366
267	307	347	367
270	310	350	370
271	311	351	371
272	312	352	372
273	313	353	373
274	314	354	374
275	315	355	375
276	316	356	376
277	317	357	377
300	320	360	271
301	321	361	270
302	322	362	220
303	323	363	203
304	324	364	272
305	325	365	276
306	326	366	263

iconv_8859-5(5)

Conversions Performed			
ISO 8859-5	MS 1251	ISO 8859-5	MS 1251
307	327	367	277
310	330	370	274
311	331	371	232
312	332	372	234
313	333	373	236
314	334	374	235
315	335	375	210
316	336	376	242
376	331		

ISO 8859-5 to Mac Cyrillic

For the conversion of ISO 8859-5 to Mac Cyrillic, all characters not in the following table are mapped unchanged.

Conversions Performed			
ISO 8859-5	Mac Cyrillic	ISO 8859-5	Mac Cyrillic
24	4	317	237
200-237	40	320	340
240	312	321	341
241	335	322	342
242	253	323	343
243	256	324	344
244	270	325	345
245	301	326	346
246	247	327	347
247	272	330	350
250	267	331	351
251	274	332	352
252	276	333	353
253	40	334	354

Conversions Performed			
ISO 8859-5	Mac Cyrillic	ISO 8859-5	Mac Cyrillic
254	315	335	355
255	40	336	356
256	330	337	357
257	332	340	360
260	200	341	361
261	201	342	362
262	202	343	363
263	203	344	364
264	204	345	365
265	205	346	366
266	206	347	367
267	207	350	370
270	210	351	371
271	211	352	372
272	212	353	373
273	213	354	374
274	214	355	375
275	215	356	376
276	216	357	337
277	217	360	334
300	220	361	336
301	221	362	254
302	222	363	257
303	223	364	271
304	224	365	317
305	225	366	264
306	226	367	273
307	227	370	300

iconv_8859-5(5)

Conversions Performed			
ISO 8859-5	Mac Cyrillic	ISO 8859-5	Mac Cyrillic
310	230	371	275
311	231	372	277
312	232	373	40
313	233	374	316
314	234	375	40
315	235	376	331
316	236		

FILES /usr/lib/iconv/*.so conversion modules
/usr/lib/iconv/*.t conversion tables
/usr/lib/iconv/iconv_data list of conversions supported by conversion tables

SEE ALSO iconv(1), iconv(3C), iconv(5)

NAME iconv_dhn – code set conversion tables for DHN (Dom Handlowy Nauki)

DESCRIPTION The following code set conversions are supported:

Code Set Conversions Supported				
Code	Symbol	Target Code	Symbol	Target Output
DHN	dhn	ISO 8859-2	iso2	ISO Latin 2
DHN	dhn	MS 1250	win2	Windows Latin 2
DHN	dhn	MS 852	dos2	MS-DOS Latin 2
DHN	dhn	Mazovia	maz	Mazovia

CONVERSIONS The conversions are performed according to the following tables. All values in the tables are given in octal.

DHN to ISO 8859-2 For the conversion of DHN to ISO 8859-2, all characters not in the following table are mapped unchanged.

Conversions Performed			
DHN	ISO 8859-2	DHN	ISO 8859-2
24-177	40	222	40
200	241	223	364
201	306	224	366
202	312	225-230	40
203	243	231	326
204	321	232	334
205	323	233-237	40
206	246	240	341
207	254	241	355
210	257	242	363
211	261	243	372
212	346	244-340	40
213	352	341	337
214	263	342-365	40

iconv_dhn(5)

Conversions Performed			
DHN	ISO 8859-2	DHN	ISO 8859-2
215	361	366	367
216	363	367	40
217	266	370	260
220	274	371-376	40
221	277		

DHN to MS 1250

For the conversion of DHN to MS 1250, all characters not in the following table are mapped unchanged.

Conversions Performed			
DHN	MS 1250	DHN	MS 1250
200	245	233-237	40
201	306	240	341
202	312	241	355
203	243	242	363
204	321	243	372
205	323	244-251	40
206	214	252	254
207	217	253-255	40
210	257	256	253
211	271	257	273
212	346	260-340	40
213	352	341	337
214	263	342-345	40
215	361	346	265
216	363	347-360	40
217	234	361	261
220	237	362-365	40
221	277	366	367

Conversions Performed			
DHN	MS 1250	DHN	MS 1250
222	40	367	40
223	364	370	260
224	366	371	40
225-230	40	372	267
231	326	373-376	40
232	334		

DHN to MS 852 For the conversion of DHN to MS 852, all characters not in the following table are mapped unchanged.

Conversions Performed			
DHN	MS 852	DHN	MS 852
200	244	212	206
201	217	213	251
202	250	214	210
203	235	215	344
204	343	216	242
205	340	217	230
206	227	220	253
207	215	221	276
210	275	222-375	40
211	245		

DHN to Mazovia For the conversion of DHN to Mazovia, all characters not in the following table are mapped unchanged.

Conversions Performed			
DHN	Mazovia	DHN	Mazovia
200	217	212	215
201	225	213	221

iconv_dhn(5)

Conversions Performed			
DHN	Mazovia	DHN	Mazovia
202	220	214	222
203	234	215	244
204	245	216	242
205	243	217	236
206	230	220	246
207	240	221	247
210	241	222-247	40
211	206		

FILES /usr/lib/iconv/*.so conversion modules
/usr/lib/iconv/*.t conversion tables
/usr/lib/iconv/iconv_data list of conversions supported by conversion tables

SEE ALSO iconv(1), iconv(3C), iconv(5)

NAME iconv_koi8-r – code set conversion tables for KOI8-R

DESCRIPTION The following code set conversions are supported:

Code Set Conversions Supported				
Code	Symbol	Target Code	Symbol	Target Output
KOI8-R	koi8	ISO 8859-5	iso5	ISO 8859-5 Cyrillic
KOI8-R	koi8	PC Cyrillic	alt	Alternative PC Cyrillic
KOI8-R	koi8	MS 1251	win5	Windows Cyrillic
KOI8-R	koi8	Mac Cyrillic	mac	Macintosh Cyrillic

CONVERSIONS The conversions are performed according to the following tables. All values in the tables are given in octal.

KOI8-R to ISO 8859-5 For the conversion of KOI8-R to ISO 8859-5, all characters not in the following table are mapped unchanged.

Conversions Performed			
KOI8-R	ISO 8859-5	KOI8-R	ISO 8859-5
24	4	320	337
241	362	321	357
242	363	322	340
243	361	323	341
244	364	324	342
245	365	325	343
246	366	327	322
247	367	330	354
250	370	331	353
251	371	332	327
252	372	333	350
253	373	334	355
254	374	335	351
256	376	336	347

iconv_koi8-r(5)

Conversions Performed			
KOI8-R	ISO 8859-5	KOI8-R	ISO 8859-5
257	377	337	352
260	360	340	316
261	242	341	260
262	243	342	261
263	241	343	306
264	244	344	264
265	245	345	265
266	246	346	304
267	247	347	263
270	250	350	305
271	251	351	270
272	252	352	271
273	253	353	272
274	254	354	273
275	255	355	274
276	256	356	275
277	257	357	276
300	356	360	277
301	320	361	317
302	321	362	300
303	346	363	301
304	324	364	302
305	325	365	303
306	344	366	266
307	323	367	262
310	345	370	314
311	330	371	313
312	331	372	267

Conversions Performed			
KOI8-R	ISO 8859-5	KOI8-R	ISO 8859-5
313	332	373	310
314	333	374	315
315	334	375	311
316	335	376	307
317	336		

KOI8-R to PC Cyrillic

For the conversion of KOI8-R to PC Cyrillic, all characters not in the following table are mapped unchanged.

Conversions Performed			
KOI8-R	PC Cyrillic	KOI8-R	PC Cyrillic
24	4	333	350
200-242	40	334	355
243	361	335	351
244-254	40	336	347
255	260	337	352
256-262	40	340	236
263	360	341	200
264-274	40	342	201
275	260	343	226
276-277	40	344	204
300	356	345	205
301	240	346	224
302	241	347	203
303	346	350	225
304	244	351	210
305	245	352	211
306	344	353	212
307	243	354	213

iconv_koi8-r(5)

Conversions Performed			
KOI8-R	PC Cyrillic	KOI8-R	PC Cyrillic
310	345	355	214
311	250	356	215
312	251	357	216
313	252	360	217
314	253	361	237
315	254	362	220
316	255	363	221
317	256	364	222
320	257	365	223
321	357	366	206
322	340	367	202
323	341	370	234
324	342	371	233
325	343	372	207
326	246	373	230
327	242	374	235
330	354	375	231
331	353	376	227
332	247		

KOI8-R to MS 1251

For the conversion of KOI8-R to MS 1251, all characters not in the following table are mapped unchanged.

Conversions Performed			
KOI8-R	MS 1251	KOI8-R	MS 1251
24	4	317	356
200-237	40	320	357
241	220	321	377
242	203	322	360

Conversions Performed			
KOI8-R	MS 1251	KOI8-R	MS 1251
243	270	323	361
244	272	324	362
245	276	325	363
246	263	326	346
247	277	327	342
250	274	330	374
251	232	331	373
252	234	332	347
253	236	333	370
254	235	334	375
255	210	335	371
256	242	336	367
257	237	337	372
260	271	340	336
261	200	341	300
262	201	342	301
263	250	343	326
264	252	344	304
265	275	345	305
266	262	346	324
267	257	347	303
270	243	350	325
271	212	351	310
272	214	352	311
273	216	353	312
274	215	354	313
275	210	355	314
276	241	356	315

iconv_koi8-r(5)

Conversions Performed			
KOI8-R	MS 1251	KOI8-R	MS 1251
277	217	357	316
300	376	360	317
301	340	361	337
302	341	362	320
303	366	363	321
304	344	364	322
305	345	365	323
306	364	366	306
307	343	367	302
310	365	370	334
311	350	371	333
312	351	372	307
313	352	373	330
314	353	374	335
315	354	375	331
316	355	376	327
376	227		

KOI8-R to Mac Cyrillic

For the conversion of KOI8-R to Mac Cyrillic, all characters not in the following table are mapped unchanged.

Conversions Performed			
KOI8-R	Mac Cyrillic	KOI8-R	Mac Cyrillic
24	4	317	356
200-237	40	320	357
240	312	321	337
241	254	322	360
242	257	323	361
243	336	324	362

Conversions Performed			
KOI8-R	Mac Cyrillic	KOI8-R	Mac Cyrillic
244	271	325	363
245	317	326	346
246	264	327	342
247	273	330	374
250	300	331	373
251	275	332	347
252	277	333	370
253	40	334	375
254	316	335	371
255	40	336	367
256	331	337	372
257	333	340	236
260	334	341	200
261	253	342	201
262	256	343	226
263	335	344	204
264	270	345	205
265	301	346	224
266	247	347	203
267	272	350	225
270	267	351	210
271	274	352	211
272	276	353	212
273	40	354	213
274	315	355	214
275	40	356	215
276	330	357	216
277	332	360	217

iconv_koi8-r(5)

Conversions Performed			
KOI8-R	Mac Cyrillic	KOI8-R	Mac Cyrillic
300	376	361	237
301	340	362	220
302	341	363	221
303	366	364	222
304	344	365	223
305	345	366	206
306	364	367	202
307	343	370	234
310	365	371	233
311	350	372	207
312	351	373	230
313	352	374	235
314	353	375	231
315	354	376	227
316	355		

FILES /usr/lib/iconv/*.so conversion modules
 /usr/lib/iconv/*.t conversion tables
 /usr/lib/iconv/iconv_data list of conversions supported by conversion tables

SEE ALSO iconv(1), iconv(3C), iconv(5)

NAME iconv_mac_cyr – code set conversion tables for Macintosh Cyrillic

DESCRIPTION The following code set conversions are supported:

Code Set Conversions Supported				
Code	Symbol	Target Code	Symbol	Target Output
Mac Cyrillic	mac	ISO 8859-5	iso5	ISO 8859-5 Cyrillic
Mac Cyrillic	mac	KOI8-R	koi8	KOI8-R
Mac Cyrillic	mac	PC Cyrillic	alt	Alternative PC Cyrillic
Mac Cyrillic	mac	MS 1251	win5	Windows Cyrillic

CONVERSIONS The conversions are performed according to the following tables. All values in the tables are given in octal.

Mac Cyrillic to ISO 8859-5 For the conversion of Mac Cyrillic to ISO 8859-5, all characters not in the following table are mapped unchanged.

Conversions Performed			
Mac Cyrillic	ISO 8859-5	Mac Cyrillic	ISO 8859-5
24	4	276	252
200	260	277	372
201	261	300	370
202	262	301	245
203	263	302-311	40
204	264	312	240
205	265	313	242
206	266	314	362
207	267	315	254
210	270	316	374
211	271	317	365
212	272	320-327	40
213	273	330	256
214	274	331	376

iconv_mac_cyr(5)

Conversions Performed			
Mac Cyrillic	ISO 8859-5	Mac Cyrillic	ISO 8859-5
215	275	332	257
216	276	333	377
217	277	334	360
220	300	335	241
221	301	336	361
222	302	337	357
223	303	340	320
224	304	341	321
225	305	342	322
226	306	343	323
227	307	344	324
230	310	345	325
231	311	346	326
232	312	347	327
233	313	350	330
234	314	351	331
235	315	352	332
236	316	353	333
237	317	354	334
240-246	40	355	335
247	246	356	336
250-252	40	357	337
253	242	360	340
254	362	361	341
255	40	362	342
256	243	363	343
257	363	364	344
260-263	40	365	345

Conversions Performed			
Mac Cyrillic	ISO 8859-5	Mac Cyrillic	ISO 8859-5
264	366	366	346
265-266	40	367	347
267	250	370	350
270	244	371	351
271	364	372	352
272	247	373	353
273	367	374	354
274	251	375	355
275	371	376	356
375	370		

**Mac Cyrillic to
KOI8-R**

For the conversion of Mac Cyrillic to KOI8-R, all characters not in the following table are mapped unchanged.

Conversions Performed			
Mac Cyrillic	KOI8-R	Mac Cyrillic	KOI8-R
24	4	276	272
200	341	277	252
201	342	300	250
202	367	301	265
203	347	302-311	40
204	344	312	240
205	345	313	261
206	366	314	241
207	372	315	274
210	351	316	254
211	352	317	245
212	353	320-327	40
213	354	330	276

iconv_mac_cyr(5)

Conversions Performed			
Mac Cyrillic	KOI8-R	Mac Cyrillic	KOI8-R
214	355	331	256
215	356	332	277
216	357	333	257
217	360	334	260
220	362	335	263
221	363	336	243
222	364	337	321
223	365	340	301
224	346	341	302
225	350	342	327
226	343	343	307
227	376	344	304
230	373	345	305
231	375	346	326
232	377	347	332
233	371	350	311
234	370	351	312
235	374	352	313
236	340	353	314
237	361	354	315
240-246	40	355	316
247	266	356	317
250-252	40	357	320
253	261	360	322
254	241	361	323
255	40	362	324
256	262	363	325
257	242	364	306

Conversions Performed			
Mac Cyrillic	KOI8-R	Mac Cyrillic	KOI8-R
260-263	40	365	310
264	246	366	303
265-266	40	367	336
267	270	370	333
270	264	371	335
271	244	372	337
272	267	373	331
273	247	374	330
274	271	375	334
275	251	376	300
375	370		

Mac Cyrillic to PC Cyrillic

For the conversion of Mac Cyrillic to PC Cyrillic, all characters not in the following table are mapped unchanged.

Conversions Performed			
Mac Cyrillic	PC Cyrillic	Mac Cyrillic	PC Cyrillic
24	4	355	255
240-334	40	356	256
335	360	357	257
336	361	360	340
337	357	361	341
340	240	362	342
341	241	363	343
342	242	364	344
343	243	365	345
344	244	366	346
345	245	367	347
346	246	370	350

iconv_mac_cyr(5)

Conversions Performed			
Mac Cyrillic	PC Cyrillic	Mac Cyrillic	PC Cyrillic
347	247	371	351
350	250	372	352
351	251	373	353
352	252	374	354
353	253	375	355
354	254	376	356
303	366		

Mac Cyrillic to MS 1251

For the conversion of Mac Cyrillic to MS 1251, all characters not in the following table are mapped unchanged.

Conversions Performed			
Mac Cyrillic	MS 1251	Mac Cyrillic	MS 1251
24	4	255	40
200	300	256	201
201	301	257	203
202	302	260-263	40
203	303	264	263
204	304	266	264
205	305	267	243
206	306	270	252
207	307	271	272
210	310	272	257
211	311	273	277
212	312	274	212
213	313	275	232
214	314	276	214
215	315	277	234
216	316	300	274

Conversions Performed			
Mac Cyrillic	MS 1251	Mac Cyrillic	MS 1251
217	317	301	275
220	320	302	254
221	321	303-306	40
222	322	307	253
223	323	310	273
224	324	311	205
225	325	312	240
226	326	313	200
227	327	314	220
230	330	315	215
231	331	316	235
232	332	317	276
233	333	320	226
234	334	321	227
235	335	322	223
236	336	323	224
237	337	324	221
240	206	325	222
241	260	326	40
242	245	327	204
243	40	330	241
244	247	331	242
245	267	332	217
246	266	333	237
247	262	334	271
250	256	335	250
252	231	336	270
253	200	337	377

iconv_mac_cyr(5)

Conversions Performed			
Mac Cyrillic	MS 1251	Mac Cyrillic	MS 1251
254	220	362	324

FILES /usr/lib/iconv/*.so conversion modules
/usr/lib/iconv/*.t conversion tables
/usr/lib/iconv/iconv_data list of conversions supported by conversion tables

SEE ALSO iconv(1), iconv(3C), iconv(5)

NAME iconv_maz – code set conversion tables for Mazovia

DESCRIPTION The following code set conversions are supported:

Code Set Conversions Supported				
Code	Symbol	Target Code	Symbol	Target Output
Mazovia	maz	ISO 8859-2	iso2	ISO Latin 2
Mazovia	maz	MS 1250	win2	Windows Latin 2
Mazovia	maz	MS 852	dos2	MS-DOS Latin 2
Mazovia	maz	DHN	dhn	Dom Hanlowy Nauki

CONVERSIONS The conversions are performed according to the following tables. All values in the tables are given in octal.

Mazovia to ISO 8859-2 For the conversion of Mazovia to ISO 8859-2, all characters not in the following table are mapped unchanged.

Conversions Performed			
Mazovia	ISO 8859-2	Mazovia	ISO 8859-2
24–177	40	230	246
200	307	231	326
201	374	232	334
202	351	233	40
203	342	234	243
204	344	235	40
205	40	236	266
206	261	237	40
207	347	240	254
210	40	241	257
211	353	242	363
212-213	40	243	323
214	356	244	361
215	346	245	321

iconv_maz(5)

Conversions Performed			
Mazovia	ISO 8859-2	Mazovia	ISO 8859-2
216	304	246	274
217	241	247	277
220	312	250-340	40
221	352	341	337
222	263	342-365	40
223	364	366	367
224	366	367	40
225	306	370	260
226-227	40	371-376	40
256	201		

Mazovia to MS 1250

For the conversion of Mazovia to MS 1250, all characters not in the following table are mapped unchanged.

Mazovia	MS 1250	Mazovia	MS 1250
200	307	236	234
201	374	237	40
202	351	240	217
203	342	241	257
204	344	242	363
205	40	243	323
206	271	244	361
207	347	245	321
210	40	246	237
211	353	247	277
212-213	40	250-251	40
214	356	252	254
215	346	253-255	40
216	304	256	253

Mazovia	MS 1250	Mazovia	MS 1250
217	245	257	273
220	312	260-340	40
221	352	341	337
222	263	342-345	40
223	364	346	265
224	366	347-360	40
225	306	361	261
226-227	40	362-365	0
230	214	366	367
231	326	367	40
232	334	370	260
233	40	371	40
234	243	372	267
235	40	373-376	40
274	212		

Mazovia to MS 852

For the conversion of Mazovia to MS 852, all characters not in the following table are mapped unchanged.

Conversions Performed			
Mazovia	MS 852	Mazovia	MS 852
205	40	234	235
206	245	235	40
210-213	40	236	230
215	206	237	40
217	244	240	215
220	250	241	275
221	251	243	340
222	210	244	344
225	217	245	343

iconv_maz(5)

Conversions Performed			
Mazovia	MS 852	Mazovia	MS 852
226-227	40	246	253
230	227	247	276
233	40	250-375	40
227	327		

Mazovia to DHN

For the conversion of Mazovia to DHN, all characters not in the following table are mapped unchanged.

Conversions Performed			
Mazovia	DHN	Mazovia	DHN
200-205	40	234	203
206	211	236	217
207-214	40	240	207
215	212	241	210
216	40	242	216
217	200	243	205
220	202	244	215
221	214	246	220
225	201	247	221
230	206		

FILES /usr/lib/iconv/*.so conversion modules
 /usr/lib/iconv/*.t conversion tables
 /usr/lib/iconv/iconv_data list of conversions supported by conversion tables

SEE ALSO iconv(1), iconv(3C), iconv(5)

NAME iconv_pc_cyr – code set conversion tables for Alternative PC Cyrillic

DESCRIPTION The following code set conversions are supported:

Code Set Conversions Supported				
Code	Symbol	Target Code	Symbol	Target Output
PC Cyrillic	alt	ISO 8859-5	iso5	ISO 8859-5 Cyrillic
PC Cyrillic	alt	KOI8-R	koi8	KOI8-R
PC Cyrillic	alt	MS 1251	win5	Windows Cyrillic
PC Cyrillic	alt	Mac Cyrillic	mac	Macintosh Cyrillic

CONVERSIONS The conversions are performed according to the following tables. All values in the tables are given in octal.

PC Cyrillic to ISO 8859-5 For the conversion of PC Cyrillic to ISO 8859-5, all characters not in the following table are mapped unchanged.

Conversions Performed			
PC Cyrillic	ISO 8859-5	PC Cyrillic	ISO 8859-5
24	4	231	311
200	260	232	312
201	261	233	313
202	262	234	314
203	263	235	315
204	264	236	316
205	265	237	317
206	266	240	320
207	267	241	321
210	270	242	322
211	271	243	323
212	272	244	324
213	273	245	325
214	274	246	326

iconv_pc_cyr(5)

Conversions Performed			
PC Cyrillic	ISO 8859-5	PC Cyrillic	ISO 8859-5
215	275	247	327
216	276	250	330
217	277	251	331
220	300	252	332
221	301	253	333
222	302	254	334
223	303	255	335
224	304	256	336
225	305	257	337
226	306	260-337	255
227	307	360	241
230	310	362-376	255

**PC Cyrillic to
KOI8-R**

For the conversion of PC Cyrillic to KOI8-R, all characters not in the following table are mapped unchanged.

Conversions Performed			
PC Cyrillic	KOI8-R	PC Cyrillic	KOI8-R
24	4	242	327
200	341	243	307
201	342	244	304
202	367	245	305
203	347	246	326
204	344	247	332
205	345	250	311
206	366	251	312
207	372	252	313
210	351	253	314
211	352	254	315

Conversions Performed			
PC Cyrillic	KOI8-R	PC Cyrillic	KOI8-R
212	353	255	316
213	354	256	317
214	355	257	320
215	356	260-337	255
216	357	340	322
217	360	341	323
220	362	342	324
221	363	343	325
222	364	344	306
223	365	345	310
224	346	346	303
225	350	347	336
226	343	350	333
227	376	351	335
230	373	352	337
231	375	353	331
232	377	354	330
233	371	355	334
234	370	356	300
235	374	357	321
236	340	360	263
237	361	361	243
240	301	362-376	255
241	302		

**PC Cyrillic to MS
1251**

For the conversion of PC Cyrillic to MS 1251, all characters not in the following table are mapped unchanged.

iconv_pc_cyr(5)

Conversions Performed			
PC Cyrillic	MS 1251	PC Cyrillic	MS 1251
24	4	242	342
200	300	243	343
201	301	244	344
202	302	245	345
203	303	246	346
204	304	247	347
205	305	250	350
206	306	251	351
207	307	252	352
210	310	253	353
211	311	254	354
212	312	255	355
213	313	256	356
214	314	257	357
215	315	260-337	210
216	316	340	360
217	317	341	361
220	320	342	362
221	321	343	363
222	322	344	364
223	323	345	365
224	324	346	366
225	325	347	367
226	326	350	370
227	327	351	371
230	330	352	372
231	331	353	373
232	332	354	374

Conversions Performed			
PC Cyrillic	MS 1251	PC Cyrillic	MS 1251
233	333	355	375
234	334	356	376
235	335	357	377
236	336	360	250
237	337	361	270
240	340	362-376	210
241	341		

PC Cyrillic to Mac Cyrillic

For the conversion of PC Cyrillic to Mac Cyrillic, all characters not in the following table are mapped unchanged.

Conversions Performed			
PC Cyrillic	Mac Cyrillic	PC Cyrillic	Mac Cyrillic
24	4	341	361
240	340	342	362
241	341	343	363
242	342	344	364
243	343	345	365
244	344	346	366
245	345	347	367
246	346	350	370
247	347	351	371
250	350	352	372
251	351	353	373
252	352	354	374
253	353	355	375
254	354	356	376
255	355	357	337
256	356	360	335

iconv_pc_cyr(5)

Conversions Performed			
PC Cyrillic	Mac Cyrillic	PC Cyrillic	Mac Cyrillic
257	357	361	336
260-337	40	362-376	40
340	360		

FILES /usr/lib/iconv/*.so conversion modules
/usr/lib/iconv/*.t conversion tables
/usr/lib/iconv/iconv_data list of conversions supported by conversion tables

SEE ALSO iconv(1), iconv(3C), iconv(5)

NAME	iconv_unicode – code set conversion tables for Unicode			
DESCRIPTION	The following code set conversions are supported:			
	CODE SET CONVERSIONS SUPPORTED -----			
	FROM Code Set Code	FROM Filename Element	TO Code Set Target Code	TO Filename Element
	ISO 8859-1 (Latin 1)	8859-1	UTF-8	UTF-8
	ISO 8859-2 (Latin 2)	8859-2	UTF-8	UTF-8
	ISO 8859-3 (Latin 3)	8859-3	UTF-8	UTF-8
	ISO 8859-4 (Latin 4)	8859-4	UTF-8	UTF-8
	ISO 8859-5 (Cyrillic)	8859-5	UTF-8	UTF-8
	ISO 8859-6 (Arabic)	8859-6	UTF-8	UTF-8
	ISO 8859-7 (Greek)	8859-7	UTF-8	UTF-8
	ISO 8859-8 (Hebrew)	8859-8	UTF-8	UTF-8
	ISO 8859-9 (Latin 5)	8859-9	UTF-8	UTF-8
	ISO 8859-10 (Latin 6)	8859-10	UTF-8	UTF-8
	Japanese EUC	eucJP	UTF-8	UTF-8
	Chinese/PRC EUC (GB 2312-1980)	gb2312	UTF-8	UTF-8
	ISO-2022	iso2022	UTF-8	UTF-8
	Korean EUC	ko_KR-euc	Korean UTF-8	ko_KR-UTF-8
	ISO-2022-KR	ko_KR-iso2022-7	Korean UTF-8	ko_KR-UTF-8
	Korean Johap (KS C 5601-1987)	ko_KR-johap	Korean UTF-8	ko_KR-UTF-8
	Korean Johap (KS C 5601-1992)	ko_KR-johap92	Korean UTF-8	ko_KR-UTF-8
	Korean UTF-8	ko_KR-UTF-8	Korean EUC	ko_KR-euc
	Korean UTF-8	ko_KR-UTF-8	Korean Johap (KS C 5601-1987)	ko_KR-johap
	Korean UTF-8	ko_KR-UTF-8	Korean Johap (KS C 5601-1992)	ko_KR-johap92
	KOI8-R (Cyrillic)	KOI8-R	UCS-2	UCS-2
	KOI8-R (Cyrillic)	KOI8-R	UTF-8	UTF-8
	PC Kanji (SJIS)	PCK	UTF-8	UTF-8
	PC Kanji (SJIS)	SJIS	UTF-8	UTF-8
	UCS-2	UCS-2	KOI8-R (Cyrillic)	KOI8-R
	UCS-2	UCS-2	UCS-4	UCS-4
	CODE SET CONVERSIONS SUPPORTED -----			
	FROM Code Set Code	FROM Filename Element	TO Code Set Target Code	TO Filename Element
	UCS-2	UCS-2	UTF-7	UTF-7
	UCS-2	UCS-2	UTF-8	UTF-8
	UCS-4	UCS-4	UCS-2	UCS-2
	UCS-4	UCS-4	UTF-16	UTF-16
	UCS-4	UCS-4	UTF-7	UTF-7
	UCS-4	UCS-4	UTF-8	UTF-8
	UTF-16	UTF-16	UCS-4	UCS-4

iconv_unicode(5)

UTF-16	UTF-16	UTF-8	UTF-8
UTF-7	UTF-7	UCS-2	UCS-2
UTF-7	UTF-7	UCS-4	UCS-4
UTF-7	UTF-7	UTF-8	UTF-8
UTF-8	UTF-8	ISO 8859-1 (Latin 1)	8859-1
UTF-8	UTF-8	ISO 8859-2 (Latin 2)	8859-2
UTF-8	UTF-8	ISO 8859-3 (Latin 3)	8859-3
UTF-8	UTF-8	ISO 8859-4 (Latin 4)	8859-4
UTF-8	UTF-8	ISO 8859-5 (Cyrillic)	8859-5
UTF-8	UTF-8	ISO 8859-6 (Arabic)	8859-6
UTF-8	UTF-8	ISO 8859-7 (Greek)	8859-7
UTF-8	UTF-8	ISO 8859-8 (Hebrew)	8859-8
UTF-8	UTF-8	ISO 8859-9 (Latin 5)	8859-9
UTF-8	UTF-8	ISO 8859-10 (Latin 6)	8859-10
UTF-8	UTF-8	Japanese EUC	eucJP
UTF-8	UTF-8	Chinese/PRC EUC (GB 2312-1980)	gb2312
UTF-8	UTF-8	ISO-2022	iso2022
UTF-8	UTF-8	KOI8-R (Cyrillic)	KOI8-R
UTF-8	UTF-8	PC Kanji (SJIS)	PCK
UTF-8	UTF-8	PC Kanji (SJIS)	SJIS
UTF-8	UTF-8	UCS-2	UCS-2
UTF-8	UTF-8	UCS-4	UCS-4
UTF-8	UTF-8	UTF-16	UTF-16
UTF-8	UTF-8	UTF-7	UTF-7
UTF-8	UTF-8	Chinese/PRC EUC (GB 2312-1980)	zh_CN.euc

CODE SET CONVERSIONS SUPPORTED

FROM Code Set Code	FROM Filename Element	TO Code Set Target Code	TO Filename Element
UTF-8	UTF-8	ISO 2022-CN	zh_CN.iso2022-7
UTF-8	UTF-8	Chinese/Taiwan Big5	zh_TW-big5
UTF-8	UTF-8	Chinese/Taiwan EUC (CNS 11643-1992)	zh_TW-euc
UTF-8	UTF-8	ISO 2022-TW	zh_TW-iso2022-7
Chinese/PRC EUC (GB 2312-1980)	zh_CN.euc	UTF-8	UTF-8
ISO 2022-CN	zh_CN.iso2022-7	UTF-8	UTF-8
Chinese/Taiwan Big5	zh_TW-big5	UTF-8	UTF-8
Chinese/Taiwan EUC (CNS 11643-1992)	zh_TW-euc	UTF-8	UTF-8
ISO 2022-TW	zh_TW-iso2022-7	UTF-8	UTF-8

EXAMPLES

EXAMPLE 1 The library module filename

In the conversion library, `/usr/lib/iconv` (see `iconv(3C)`), the library module filename is composed of two symbolic elements separated by the percent sign (%). The first symbol specifies the code set that is being converted; the second symbol specifies the *target code*, that is, the code set to which the first one is being converted.

EXAMPLE 1 The library module filename (Continued)

In the conversion table above, the first symbol is termed the "FROM Filename Element". The second symbol, representing the target code set, is the "TO Filename Element".

For example, the library module filename to convert from the *Korean EUC* code set to the *Korean UTF-8* code set is

```
ko_KR-euc%ko_KR-UTF-8
```

FILES /usr/lib/iconv/*.so conversion modules

SEE ALSO iconv(1), iconv(3C), iconv(5)

Chernov, A., *Registration of a Cyrillic Character Set*, RFC 1489, RELCOM Development Team, July 1993.

Chon, K., H. Je Park, and U. Choi, *Korean Character Encoding for Internet Messages*, RFC 1557, Solvit Chosun Media, December 1993.

Goldsmith, D., and M. Davis, *UTF-7 – A Mail-Safe Transformation Format of Unicode*, RFC 1642, Taligent, Inc., July 1994.

Lee, F., *HZ – A Data Format for Exchanging Files of Arbitrarily Mixed Chinese and ASCII characters*, RFC 1843, Stanford University, August 1995.

Murai, J., M. Crispin, and E. van der Poel, *Japanese Character Encoding for Internet Messages*, RFC 1468, Keio University, Panda Programming, June 1993.

Nussbacher, H., and Y. Bourvine, *Hebrew Character Encoding for Internet Messages*, RFC 1555, Israeli Inter-University, Hebrew University, December 1993.

Ohta, M., *Character Sets ISO-10646 and ISO-10646-J-1*, RFC 1815, Tokyo Institute of Technology, July 1995.

Ohta, M., and K. Handa, *ISO-2022-JP-2: Multilingual Extension of ISO-2022-JP*, RFC 1554, Tokyo Institute of Technology, December 1993.

Reynolds, J., and J. Postel, *ASSIGNED NUMBERS*, RFC 1700, University of Southern California/Information Sciences Institute, October 1994.

Simonson, K., *Character Mnemonics & Character Sets*, RFC 1345, Rational Almen Planlaegning, June 1992.

Spinellis, D., *Greek Character Encoding for Electronic Mail Messages*, RFC 1947, SENA S.A., May 1996.

iconv_unicode(5)

The Unicode Consortium, *The Unicode Standard, Version 2.0*, Addison Wesley Developers Press, July 1996.

Wei, Y., Y. Zhang, J. Li, J. Ding, and Y. Jiang, *ASCII Printable Characters-Based Chinese Character Encoding for Internet Messages*, RFC 1842, AsiaInfo Services Inc., Harvard University, Rice University, University of Maryland, August 1995.

Yergeau, F., *UTF-8, a transformation format of Unicode and ISO 10646*, RFC 2044, Alis Technologies, October 1996.

Zhu, H., D. Hu, Z. Wang, T. Kao, W. Chang, and M. Crispin, *Chinese Character Encoding for Internet Messages*, RFC 1922, Tsinghua University, China Information Technology Standardization Technical Committee (CITS), Institute for Information Industry (III), University of Washington, March 1996.

NOTES ISO 8859 character sets using Latin alphabetic characters are distinguished as follows:

ISO 8859-1 (Latin 1)

For most West European languages, including:

Albanian	Finnish	Italian
Catalan	French	Norwegian
Danish	German	Portuguese
Dutch	Galician	Spanish
English	Irish	Swedish
Faeroese	Icelandic	

ISO 8859-2 (Latin 2)

For most Latin-written Slavic and Central European languages:

Czech	Polish	Slovak
German	Rumanian	Slovene
Hungarian	Croatian	

ISO 8859-3 (Latin 3)

Popularly used for Esperanto, Galician, Maltese, and Turkish.

ISO 8859-4 (Latin 4)

Introduces letters for Estonian, Latvian, and Lithuanian. It is an incomplete predecessor of ISO 8859-10 (Latin 6).

iconv_unicode(5)

ISO 8859-9 (Latin 5)

Replaces the rarely needed Icelandic letters in ISO 8859-1 (Latin 1) with the Turkish ones.

ISO 8859-10 (Latin 6)

Adds the last Inuit (Greenlandic) and Sami (Lappish) letters that were not included in ISO 8859-4 (Latin 4) to complete coverage of the Nordic area.

isalist(5)

NAME	isalist – the native instruction sets known to Solaris software
DESCRIPTION	<p>The possible instruction set names returned by <code>isalist(1)</code> and the <code>SI_ISALIST</code> command of <code>sysinfo(2)</code> are listed here.</p> <p>The list is ordered within an instruction set family in the sense that later names are generally faster than earlier names; note that this is in the reverse order than listed by <code>isalist(1)</code> and <code>sysinfo(2)</code>. In the following list of values, numbered entries generally represent increasing performance; lettered entries are either mutually exclusive or cannot be ordered.</p>
SPARC Platforms	<p>Where appropriate, correspondence with a given value of the <code>-xarch</code> option of Sun's C 4.0 compiler is indicated. Other compilers may have similar options.</p> <p>1a. <code>sparc</code> Indicates the SPARC V8 instruction set, as defined in The SPARC Architecture Manual, Version 8, Prentice-Hall, Inc., 1992. Some instructions (such as integer multiply and divide, <code>FSMULD</code>, and all floating point operations on quad operands) may be emulated by the kernel on certain systems.</p> <p>1b. <code>sparcv7</code> Same as <code>sparc</code>. This corresponds to code produced with the <code>-xarch=v7</code> option of Sun's C 4.0 compiler.</p> <p>2. <code>sparcv8-fsmuld</code> Like <code>sparc</code>, except that integer multiply and divide must be executed in hardware. This corresponds to code produced with the <code>-xarch=v8a</code> option of Sun's C 4.0 compiler.</p> <p>3. <code>sparcv8</code> Like <code>sparcv8-fsmuld</code>, except that <code>FSMULD</code> must also be executed in hardware. This corresponds to code produced with the <code>-xarch=v8</code> option of Sun's C 4.0 compiler.</p> <p>4. <code>sparcv8plus</code> Indicates the SPARC V8 instruction set plus those instructions in the SPARC V9 instruction set, as defined in The SPARC Architecture Manual, Version 9, Prentice-Hall, 1994, that can be used according to The V8+ Technical Specification. This corresponds to code produced with the <code>-xarch=v8plus</code> option of Sun's C 4.0 compiler.</p> <p>5a. <code>sparcv8plus+vis</code> Like <code>sparcv8plus</code>, with the addition of those UltraSPARC I Visualization Instructions that can be used according to The V8+ Technical Specification. This corresponds to code produced with the <code>-xarch=v8plusa</code> option of Sun's C 4.0 compiler.</p> <p>5b. <code>sparcv8plus+fmuladd</code> Like <code>sparcv8plus</code>, with the addition of the Hal SPARC64 floating multiply-add and multiply-subtract instructions.</p>

	6. sparcv9	Indicates the SPARC V9 instruction set, as defined in The SPARC Architecture Manual, Version 9, Prentice-Hall, 1994.
	7a. sparcv9+vis	Like sparcv9, with the addition of the UltraSPARC I Visualization Instructions.
	7b. sparcv9+fmuladd	Like sparcv9, with the addition of the Hal SPARC64 floating multiply-add and multiply-subtract instructions.
x86 Platforms	1. i386	The Intel 80386 instruction set, as described in The i386 Microprocessor Programmer's Reference Manual.
	2. i486	The Intel 80486 instruction set, as described in The i486 Microprocessor Programmer's Reference Manual. (This is effectively i386, plus the CMPXCHG, BSWAP, and XADD instructions.)
	3. pentium	The Intel Pentium instruction set, as described in The Pentium Processor User's Manual. (This is effectively i486, plus the CPU_ID instruction, and any features that the CPU_ID instruction indicates are present.)
	4. pentium+mmx	Like pentium, with the MMX instructions guaranteed present.
	5. pentium_pro	The Intel PentiumPro instruction set, as described in The PentiumPro Family Developer's Manual. (This is effectively pentium, with the CMOVcc, FCMOVcc, FCOMI, and RDPMC instructions guaranteed present.)
	6. pentium_pro+mmx	Like pentium_pro, with the MMX instructions guaranteed present.
SEE ALSO	isalist(1), sysinfo(2)	

largefile(5)

NAME	largefile – large file status of utilities																																																																																
DESCRIPTION	A <i>large file</i> is a regular file whose size is greater than or equal to 2 Gbyte (2 ³¹ bytes). A <i>small file</i> is a regular file whose size is less than 2 Gbyte.																																																																																
Large file aware utilities	<p>A utility is called <i>large file aware</i> if it can process large files in the same manner as it does small files. A utility that is large file aware is able to handle large files as input and generate as output large files that are being processed. The exception is where additional files are used as system configuration files or support files that can augment the processing. For example, the <code>file</code> utility supports the <code>-m</code> option for an alternative "magic" file and the <code>-f</code> option for a support file that can contain a list of file names. It is unspecified whether a utility that is large file aware will accept configuration or support files that are large files. If a large file aware utility does not accept configuration or support files that are large files, it will cause no data loss or corruption upon encountering such files and will return an appropriate error.</p> <p>The following <code>/usr/bin</code> utilities are large file aware:</p> <table border="0" style="width: 100%;"> <tr><td>adb</td><td>awk</td><td>bdiff</td><td>cat</td><td>chgrp</td></tr> <tr><td>chmod</td><td>chown</td><td>cksum</td><td>cmp</td><td>compress</td></tr> <tr><td>cp</td><td>csch</td><td>csplit</td><td>cut</td><td>dd</td></tr> <tr><td>dircmp</td><td>du</td><td>egrep</td><td>fgrep</td><td>file</td></tr> <tr><td>find</td><td>ftp</td><td>getconf</td><td>grep</td><td>head</td></tr> <tr><td>join</td><td>jsh</td><td>ksh</td><td>ln</td><td>ls</td></tr> <tr><td>mdb</td><td>mkdir</td><td>mkfifo</td><td>more</td><td>mv</td></tr> <tr><td>nawk</td><td>page</td><td>paste</td><td>pathchk</td><td>pg</td></tr> <tr><td>rcp</td><td>remsh</td><td>rksh</td><td>rm</td><td>rmdir</td></tr> <tr><td>rsh</td><td>sed</td><td>sh</td><td>sort</td><td>split</td></tr> <tr><td>sum</td><td>tail</td><td>tar</td><td>tee</td><td>test</td></tr> <tr><td>touch</td><td>tr</td><td>uncompress</td><td>uudecode</td><td>uuencode</td></tr> <tr><td>wc</td><td>zcat</td><td></td><td></td><td></td></tr> </table> <p>The following <code>/usr/xpg4/bin</code> utilities are large file aware:</p> <table border="0" style="width: 100%;"> <tr><td>awk</td><td>cp</td><td>du</td><td>egrep</td><td>fgrep</td></tr> <tr><td>grep</td><td>ln</td><td>ls</td><td>more</td><td>mv</td></tr> <tr><td>rm</td><td>sed</td><td>sh</td><td>sort</td><td>tail</td></tr> </table>	adb	awk	bdiff	cat	chgrp	chmod	chown	cksum	cmp	compress	cp	csch	csplit	cut	dd	dircmp	du	egrep	fgrep	file	find	ftp	getconf	grep	head	join	jsh	ksh	ln	ls	mdb	mkdir	mkfifo	more	mv	nawk	page	paste	pathchk	pg	rcp	remsh	rksh	rm	rmdir	rsh	sed	sh	sort	split	sum	tail	tar	tee	test	touch	tr	uncompress	uudecode	uuencode	wc	zcat				awk	cp	du	egrep	fgrep	grep	ln	ls	more	mv	rm	sed	sh	sort	tail
adb	awk	bdiff	cat	chgrp																																																																													
chmod	chown	cksum	cmp	compress																																																																													
cp	csch	csplit	cut	dd																																																																													
dircmp	du	egrep	fgrep	file																																																																													
find	ftp	getconf	grep	head																																																																													
join	jsh	ksh	ln	ls																																																																													
mdb	mkdir	mkfifo	more	mv																																																																													
nawk	page	paste	pathchk	pg																																																																													
rcp	remsh	rksh	rm	rmdir																																																																													
rsh	sed	sh	sort	split																																																																													
sum	tail	tar	tee	test																																																																													
touch	tr	uncompress	uudecode	uuencode																																																																													
wc	zcat																																																																																
awk	cp	du	egrep	fgrep																																																																													
grep	ln	ls	more	mv																																																																													
rm	sed	sh	sort	tail																																																																													

tr

The following `/usr/sbin` utilities are large file aware:

install mkfile mknod mmdir swap

See the `USAGE` section of the `swap(1M)` manual page for limitations of `swap` on block devices greater than 2 Gbyte on a 32-bit operating system.

The following `/usr/ucb` utilities are large file aware:

chown from ln ls sed
sum touch

The `/usr/bin/cpio` and `/usr/bin/pax` utilities are large file aware, but cannot archive a file whose size exceeds 8 Gbyte - 1 byte.

The `/usr/bin/truss` utilities has been modified to read a dump file and display information relevant to large files, such as offsets.

cachefs file systems

The following `/usr/bin` utilities are large file aware for `cachefs` file systems:

cachefspack cachefsstat

The following `/usr/sbin` utilities are large file aware for `cachefs` file systems:

cachefslog cachefswssize cfsadmin fsck
mount umount

nfs file systems

The following utilities are large file aware for `nfs` file systems:

`/usr/lib/autofs/automountd` `/usr/sbin/mount`
`/usr/lib/nfs/rquotad`

ufs file systems

The following `/usr/bin` utility is large file aware for `ufs` file systems:

df

largefile(5)

The following `/usr/lib/nfs` utility is large file aware for `ufs` file systems:

`rquotad`

The following `/usr/xpg4/bin` utility is large file aware for `ufs` file systems:

`df`

The following `/usr/sbin` utilities are large file aware for `ufs` file systems:

<code>clri</code>	<code>dcopy</code>	<code>edquota</code>	<code>ff</code>	<code>fsck</code>
<code>fsdb</code>	<code>fsirand</code>	<code>fstyp</code>	<code>labelit</code>	<code>lockfs</code>
<code>mkfs</code>	<code>mount</code>	<code>ncheck</code>	<code>newfs</code>	<code>quot</code>
<code>quota</code>	<code>quotacheck</code>	<code>quotaoff</code>	<code>quotaon</code>	<code>repquota</code>
<code>tunefs</code>	<code>ufsdump</code>	<code>ufsrestore</code>	<code>umount</code>	

Large file safe utilities

A utility is called *large file safe* if it causes no data loss or corruption when it encounters a large file. A utility that is large file safe is unable to process properly a large file, but returns an appropriate error.

The following `/usr/bin` utilities are large file safe:

<code>audioconvert</code>	<code>audioplay</code>	<code>audiorecord</code>	<code>comm</code>	<code>diff</code>
<code>diff3</code>	<code>diffmk</code>	<code>ed</code>	<code>lp</code>	<code>mail</code>
<code>mailcompat</code>	<code>mailstats</code>	<code>mailx</code>	<code>pack</code>	<code>pcat</code>
<code>red</code>	<code>rmail</code>	<code>sdiff</code>	<code>unpack</code>	<code>vi</code>
<code>view</code>				

The following `/usr/xpg4/bin` utilities are large file safe:

<code>ed</code>	<code>vi</code>	<code>view</code>
-----------------	-----------------	-------------------

The following `/usr/sbin` utilities are large file safe:

<code>lpfilter</code>	<code>lpforms</code>
-----------------------	----------------------

The following `/usr/ucb` utilities are large file safe:

Mail lpr

The following `/usr/lib` utility is large file safe:

sendmail

SEE ALSO `lf64(5)`, `lfcompile(5)`, `lfcompile64(5)`

lf64(5)

NAME	lf64 – transitional interfaces for 64-bit file offsets																										
DESCRIPTION	The data types, interfaces, and macros described on this page provide explicit access to 64-bit file offsets. They are accessible through the transitional compilation environment described on the <code>lfcompile64(5)</code> manual page. The function prototype and semantics of a transitional interface are equivalent to those of the standard version of the call, except that relevant data types are 64-bit entities.																										
Data Types	<p>The following tables list the standard data or struct types in the left-hand column and their corresponding explicit 64-bit file offset types in the right-hand column, grouped by header. The absence of an entry in the left-hand column indicates that there is no existing explicit 32-bit type that corresponds to the 64-bit type listed in the right-hand column. Note that in a 64-bit application, the standard definition is equivalent to the 64-bit file offset definition.</p> <p>< aio.h ></p> <table border="0"> <tr> <td><code>struct aiocb</code></td> <td><code>struct aiocb64</code></td> </tr> <tr> <td><code>off_t aio_offset;</code></td> <td><code>off64_t aio_offset;</code></td> </tr> </table> <p>< sys/dirent.h ></p> <table border="0"> <tr> <td><code>struct dirent</code></td> <td><code>struct dirent64</code></td> </tr> <tr> <td><code>ino_t d_ino;</code></td> <td><code>ino64_t d_ino;</code></td> </tr> <tr> <td><code>off_t d_off;</code></td> <td><code>off64_t d_off;</code></td> </tr> </table> <p>< sys/fcntl.h ></p> <table border="0"> <tr> <td><code>struct flock</code></td> <td><code>struct flock64</code></td> </tr> <tr> <td><code>off_t l_start;</code></td> <td><code>off64_t l_start;</code></td> </tr> <tr> <td><code>off_t l_len;</code></td> <td><code>off64_t l_len;</code></td> </tr> <tr> <td><code>F_SETLK</code></td> <td><code>F_SETLK64</code></td> </tr> <tr> <td><code>F_SETLKW</code></td> <td><code>F_SETLKW64</code></td> </tr> <tr> <td><code>F_GETLK</code></td> <td><code>F_GETLK64</code></td> </tr> <tr> <td><code>F_FREESP</code></td> <td><code>F_FREESP64</code></td> </tr> <tr> <td></td> <td><code>O_LARGEFILE</code></td> </tr> </table> <p>< sys/stdio.h ></p>	<code>struct aiocb</code>	<code>struct aiocb64</code>	<code>off_t aio_offset;</code>	<code>off64_t aio_offset;</code>	<code>struct dirent</code>	<code>struct dirent64</code>	<code>ino_t d_ino;</code>	<code>ino64_t d_ino;</code>	<code>off_t d_off;</code>	<code>off64_t d_off;</code>	<code>struct flock</code>	<code>struct flock64</code>	<code>off_t l_start;</code>	<code>off64_t l_start;</code>	<code>off_t l_len;</code>	<code>off64_t l_len;</code>	<code>F_SETLK</code>	<code>F_SETLK64</code>	<code>F_SETLKW</code>	<code>F_SETLKW64</code>	<code>F_GETLK</code>	<code>F_GETLK64</code>	<code>F_FREESP</code>	<code>F_FREESP64</code>		<code>O_LARGEFILE</code>
<code>struct aiocb</code>	<code>struct aiocb64</code>																										
<code>off_t aio_offset;</code>	<code>off64_t aio_offset;</code>																										
<code>struct dirent</code>	<code>struct dirent64</code>																										
<code>ino_t d_ino;</code>	<code>ino64_t d_ino;</code>																										
<code>off_t d_off;</code>	<code>off64_t d_off;</code>																										
<code>struct flock</code>	<code>struct flock64</code>																										
<code>off_t l_start;</code>	<code>off64_t l_start;</code>																										
<code>off_t l_len;</code>	<code>off64_t l_len;</code>																										
<code>F_SETLK</code>	<code>F_SETLK64</code>																										
<code>F_SETLKW</code>	<code>F_SETLKW64</code>																										
<code>F_GETLK</code>	<code>F_GETLK64</code>																										
<code>F_FREESP</code>	<code>F_FREESP64</code>																										
	<code>O_LARGEFILE</code>																										

fpos_t	fpos64_t
<sys/resource.h>	
rlim_t	rlim64_t
struct rlimit	struct rlimit64
rlim_t rlim_cur;	rlim64_t rlim_cur;
rlim_t rlim_max;	rlim64_t rlim_max;
RLIM_INFINITY	RLIM64_INFINITY
RLIM_SAVED_MAX	RLIM64_SAVED_MAX
RLIM_SAVED_CUR	RLIM64_SAVED_CUR
<sys/stat.h>	
struct stat	struct stat64
ino_t st_ino;	ino64_t st_ino;
off_t st_size;	off64_t st_size;
blkcnt_t st_blocks;	blkcnt64_t st_blocks;
<sys/statvfs.h>	
struct statvfs	struct statvfs64
fsblkcnt_t f_blocks;	fsblkcnt64_t f_blocks;
fsblkcnt_t f_bfree;	fsblkcnt64_t f_bfree;
fsblkcnt_t f_bavail;	fsblkcnt64_t f_bavail;
fsfilcnt_t f_files;	fsfilcnt64_t f_files;
fsfilcnt_t f_ffree;	fsfilcnt64_t f_ffree;
fsfilcnt_t f_favail;	fsfilcnt64_t f_favail;
<sys/types.h>	
off_t;	off64_t;

lf64(5)

```
ino_t;                ino64_t;
blkcnt_t;             blkcnt64_t;
fsblkcnt_t;          fsblkcnt64_t;
fsfilcnt_t;          fsfilcnt64_t;
```

<unistd.h>

```
_LFS64_LARGEFILE
_LFS64_STDIO
```

<sys/unistd.h>

```
_CS_LFS64_CFLAGS
_CS_LFS64_LDFLAGS
_CS_LFS64_LIBS
_CS_LFS64_LINTFLAGS
```

System Interfaces

The following tables display the standard API and the corresponding transitional interfaces for 64-bit file offsets. The interfaces are grouped by header. The interface name and the affected data types are displayed in courier font..

<aio.h>

```
int aio_cancel(...,      int aio_cancel64(...,
struct aiocb *);         struct aiocb64 *);
int aio_error(           int aio_error64(
const struct aiocb *);   const struct aiocb64 *);
int aio_fsync(...,      int aio_fsync64(...,
struct aiocb *);         struct aiocb64 *);
int aio_read(struct aiocb *);  int aio_read64(struct aiocb64 *);
int aio_return(struct aiocb *); int aio_return64(struct aiocb64 *);
int aio_suspend(         int aio_suspend64(
const struct aiocb *, ...); const struct aiocb64 *, ...);
```

<code>int aio_write(struct aiocb *);</code>	<code>int aio_write64(struct aiocb64 *);</code>
<code>int lio_listio(...,</code>	<code>int lio_listio64(...,</code>
<code>const struct aiocb *,...);</code>	<code>const struct aiocb64 *,...);</code>
<dirent.h>	
<code>struct dirent *readdir();</code>	<code>struct dirent64 *readdir64();</code>
<code>struct dirent *readdir_r();</code>	<code>struct dirent64 *readdir64_r();</code>
<fcntl.h>	
<code>int attropen();</code>	<code>int attropen64();</code>
<code>int creat();</code>	<code>int creat64();</code>
<code>int open();</code>	<code>int open64();</code>
<code>int openat();</code>	<code>int openat64();</code>
<ftw.h>	
<code>int ftw(...,</code>	<code>int ftw64(...,</code>
<code>const struct stat *</code>	<code>const struct stat64 *</code>
<code>...);</code>	<code>...);</code>
<code>int nftw(...,</code>	<code>int nftw64(...,</code>
<code>const struct stat *</code>	<code>const struct stat64 *</code>
<code>...);</code>	<code>...);</code>
<libgen.h>	
<code>char *copylist(..., off_t);</code>	<code>char *copylist64(...,</code>
	<code>off64_t);</code>
<stdio.h>	
<code>int fgetpos();</code>	<code>int fgetpos64();</code>

FILE *fopen();	FILE *fopen64();
FILE *freopen();	FILE *freopen64();
int fseeko(..., off_t,...);	int fseeko64(..., off64_t,...);
int fsetpos(..., const fpos_t *);	int fsetpos64(..., const fpos64_t *);
off_t ftello();	off64_t ftello64();
FILE *tmpfile();	FILE *tmpfile64();
 <stdlib.h>	
int mkstemp();	int mkstemp64();
 <sys/async.h>	
int aioread(...,off_t, ...);	int aioread64(...,off64_t, ...);
int aiowrite(...,off_t, ...);	int aiowrite64(..., off64_t,...);
 <ucbinclude/sys/dir.h>	
int alphasort(struct direct **, struct direct **);	int alphasort64(struct direct64 **, struct direct64 **);
struct direct *readdir()	struct direct64 *readdir64();
int scandir(..., struct direct *(*[]), ...);	int scandir64(..., struct direct64 *(*[]), ...);
 <sys/dirent.h>	

<code>int getdents(..., dirent);</code>	<code>int getdents64(..., dirent64);</code>
<sys/mman.h>	
<code>void mmap(..., off_t);</code>	<code>void mmap64(..., off64_t);</code>
<sys/resource.h>	
<code>int getrlimit(..., struct rlimit *);</code>	<code>int getrlimit64(..., struct rlimit64 *);</code>
<code>int setrlimit(..., const struct rlimit *);</code>	<code>int setrlimit64(..., const struct rlimit64 *);</code>
<sys/stat.h>	
<code>int fstat(..., struct stat *);</code>	<code>int fstat64(..., struct stat64 *);</code>
<code>int fstatat(..., struct stat *, int);</code>	<code>int fstatat64(..., struct stat64 *, int);</code>
<code>int lstat(..., struct stat *);</code>	<code>int lstat64(..., struct stat64 *);</code>
<code>int stat(..., struct stat *);</code>	<code>int stat64(..., struct stat64 *);</code>
<sys/statvfs.h>	
<code>int statvfs(..., struct statvfs *);</code>	<code>int statvfs64(..., struct statvfs64 *);</code>
<code>int fstatvfs(..., struct statvfs *);</code>	<code>int fstatvfs64(..., struct statvfs64 *);</code>
<unistd.h>	

lf64(5)

<code>int lockf(..., off_t);</code>	<code>int lockf64(..., off64_t);</code>
<code>off_t lseek(..., off_t, ...);</code>	<code>off64_t lseek64(..., off64_t, ...);</code>
<code>int ftruncate(..., off_t);</code>	<code>int ftruncate64(..., off64_t);</code>
<code>ssize_t pread(..., off_t);</code>	<code>ssize_t pread64(..., off64_t);</code>
<code>ssize_t pwrite(..., off_t);</code>	<code>ssize_t pwrite64(..., off64_t);</code>
<code>int truncate(..., off_t);</code>	<code>int truncate64(..., off64_t);</code>

SEE ALSO `lfcompile(5)`, `lfcompile64(5)`

NAME	Ifcompile – large file compilation environment for 32-bit applications										
DESCRIPTION	<p>All 64-bit applications can manipulate large files by default. The methods described on this page allow 32-bit applications to manipulate large files.</p> <p>In the large file compilation environment, source interfaces are bound to appropriate 64-bit functions, structures, and types. Compiling in this environment allows 32-bit applications to access files whose size is greater than or equal to 2 Gbyte (2^{31} bytes).</p> <p>Each interface named <code>xxx()</code> that needs to access 64-bit entities to access large files maps to a <code>xxx64()</code> call in the resulting binary. All relevant data types are defined to be of correct size (for example, <code>off_t</code> has a typedef definition for a 64-bit entity).</p> <p>An application compiled in this environment is able to use the <code>xxx()</code> source interfaces to access both large and small files, rather than having to explicitly utilize the transitional <code>xxx64()</code> interface calls to access large files. See the <code>lfcompile64(5)</code> manual page for information regarding the transitional compilation environment.</p> <p>Applications can be compiled in the large file compilation environment by using the following methods:</p> <ul style="list-style-type: none"> ■ Use the <code>getconf(1)</code> utility with one or more of the arguments listed in the table below. This method is recommended for portable applications. <table border="1" data-bbox="461 1010 1430 1266"> <thead> <tr> <th>argument</th> <th>purpose</th> </tr> </thead> <tbody> <tr> <td>LFS_CFLAGS</td> <td>obtain compilation flags necessary to enable the large file compilation environment</td> </tr> <tr> <td>LFS_LDFLAGS</td> <td>obtain link editor options</td> </tr> <tr> <td>LFS_LIBS</td> <td>obtain link library names</td> </tr> <tr> <td>LFS_LINTFLAGS</td> <td>obtain lint options</td> </tr> </tbody> </table> <ul style="list-style-type: none"> ■ Set the compile-time flag <code>_FILE_OFFSET_BITS</code> to 64 before including any headers. Applications may combine objects produced in the large file compilation environment with objects produced in the transitional compilation environment, but must be careful with respect to interoperability between those objects. Applications should not declare global variables of types whose sizes change between compilation environments. <p>The <code>fseek()</code> and <code>ftell()</code> functions <i>do not</i> map to functions named <code>fseek64()</code> and <code>ftell64()</code>; rather, the large file additions <code>fseeko()</code> and <code>ftello()</code>, have functionality identical to <code>fseek()</code> and <code>ftell()</code> and <i>do</i> map to the 64-bit functions <code>fseeko64()</code> and <code>ftello64()</code>. Applications wishing to access large files should use <code>fseeko()</code> and <code>ftello()</code> in place of <code>fseek()</code> and <code>ftell()</code>. See the <code>fseek(3C)</code> and <code>ftell(3C)</code> manual pages for information about <code>fseeko()</code> and <code>ftello()</code>.</p>	argument	purpose	LFS_CFLAGS	obtain compilation flags necessary to enable the large file compilation environment	LFS_LDFLAGS	obtain link editor options	LFS_LIBS	obtain link library names	LFS_LINTFLAGS	obtain lint options
argument	purpose										
LFS_CFLAGS	obtain compilation flags necessary to enable the large file compilation environment										
LFS_LDFLAGS	obtain link editor options										
LFS_LIBS	obtain link library names										
LFS_LINTFLAGS	obtain lint options										
Access to Additional Large File Interfaces											

lfcompile(5)

Applications wishing to access `fseeko()` and `ftello()` as well as the POSIX and X/Open specification-conforming interfaces should define the macro `_LARGEFILE_SOURCE` to be 1 and set whichever feature test macros are appropriate to obtain the desired environment (see `standards(5)`).

EXAMPLES In the following examples, the large file compilation environment is accessed by invoking the `getconf` utility with one of the arguments listed in the table above. The additional large file interfaces are accessed by specifying `-D_LARGEFILE_SOURCE`.

The examples that use the form of command substitution specifying the command within parentheses preceded by a dollar sign can be executed only in a POSIX-conforming shell such as the Korn Shell (see `ksh(1)`). In a shell that is not POSIX-conforming, such as the Bourne Shell (see `sh(1)`) and the C Shell (see `csh(1)`), the `getconf` calls must be enclosed within grave accent marks, as shown in the second example.

EXAMPLE 1 An example of compiling a program with a “large” `off_t`, and that uses `fseeko()`, `ftello()`, and `yacc(1)`

```
$ c89 -D_LARGEFILE_SOURCE          \
      -D_FILE_OFFSET_BITS=64 -o foo \
      $(getconf LFS_CFLAGS) y.tab.c b.o \
      $(getconf LFS_LDFLAGS)        \
      -ly $(getconf LFS_LIBS)
```

EXAMPLE 2 An example of compiling a program with a “large” `off_t` that does not use `fseeko()` and `ftello()` and has no application specific libraries:

```
% c89 -D_FILE_OFFSET_BITS=64      \
      `getconf LFS_CFLAGS` a.c     \
      `getconf LFS_LDFLAGS`       \
      `getconf LFS_LIBS`          \
```

EXAMPLE 3 An example of compiling a program with a “default” `off_t` and that uses `fseeko()` and `ftello()`:

```
$ c89 -D_LARGEFILE_SOURCE a.c
```

SEE ALSO `csh(1)`, `getconf(1)`, `ksh(1)`, `lint(1B)`, `sh(1)`, `fseek(3C)`, `ftell(3C)`, `lf64(5)`, `lfcompile64(5)`, `standards(5)`

NOTES Certain system-specific or non-portable interfaces are not usable in the large file compilation environment. Known cases are:

- Kernel data structures read from `/dev/kmem`.
- Interfaces in the kernel virtual memory library, `-lkvm`.
- Interfaces in the ELF access library, `-lelf`.
- Interfaces to `/proc` defined in `<procfs.h>`.
- The `ustat(2)` system call.

Programs that use these interfaces should not be compiled in the large file compilation environment. As a partial safeguard against making this mistake, including either of the `<libelf.h>` or `<sys/procfs.h>` header files will induce a compilation error when the large file compilation environment is enabled.

In general, caution should be exercised when using any separately-compiled library whose interfaces include data items of type `off_t` or the other redefined types either directly or indirectly, such as with `'struct stat'`. (The redefined types are `off_t`, `rlim_t`, `ino_t`, `blkcnt_t`, `fsblkcnt_t`, and `fsfilcnt_t`.) For the large file compilation environment to work correctly with such a library, the library interfaces must include the appropriate `xxx64()` binary entry points and must have them mapped to the corresponding primary functions when `_FILE_OFFSET_BITS` is set to 64.

Care should be exercised using any of the `printf()` or `scanf()` routines on variables of the types mentioned above. In the large file compilation environment, these variables should be printed or scanned using `long long` formats.

BUGS The `lint(1B)` utility will generate spurious error messages when `_FILE_OFFSET_BITS` is set to 64. This is because the binary `libc lint` library, `/usr/lib/llib-1c.ln`, is compiled only for the standard interfaces, not with `_FILE_OFFSET_BITS` set to 64. This deficiency hampers static error-checking for programs compiled in the large file compilation environment.

Symbolic formats analogous to those found in `<sys/int_fmtio.h>` do not exist for printing or scanning variables of the types that are redefined in the large file compilation environment.

lfcompile64(5)

NAME	lfcompile64 – transitional compilation environment				
DESCRIPTION	<p>All 64-bit applications can manipulate large files by default. The transitional interfaces described on this page can be used by 32-bit and 64-bit applications to manipulate large files.</p> <p>In the transitional compilation environment, explicit 64-bit functions, structures, and types are added to the API. Compiling in this environment allows both 32-bit and 64-bit applications to access files whose size is greater than or equal to 2 Gbyte (2^{31} bytes).</p> <p>The transitional compilation environment exports all the explicit 64-bit functions (<i>xxx64()</i>) and types in addition to all the regular functions (<i>xxx()</i>) and types. Both <i>xxx()</i> and <i>xxx64()</i> functions are available to the program source. A 32-bit application must use the <i>xxx64()</i> functions in order to access large files. See the <code>lf64(5)</code> manual page for a complete listing of the 64-bit transitional interfaces.</p> <p>The transitional compilation environment differs from the large file compilation environment, wherein the underlying interfaces are bound to 64-bit functions, structures, and types. An application compiled in the large file compilation environment is able to use the <i>xxx()</i> source interfaces to access both large and small files, rather than having to explicitly utilize the transitional <i>xxx64()</i> interface calls to access large files. See the <code>lfcompile(5)</code> manual page for more information regarding the large file compilation environment.</p> <p>Applications may combine objects produced in the large file compilation environment with objects produced in the transitional compilation environment, but must be careful with respect to interoperability between those objects. Applications should not declare global variables of types whose sizes change between compilation environments.</p> <p>For applications that do not wish to conform to the POSIX or X/Open specifications, the 64-bit transitional interfaces are available by default. No compile-time flags need to be set.</p> <p>Access to Additional Large File Interfaces</p> <p>Applications that wish to access the transitional interfaces as well as the POSIX or X/Open specification-conforming interfaces should use the following compilation methods and set whichever feature test macros are appropriate to obtain the desired environment (see <code>standards(5)</code>).</p> <ul style="list-style-type: none">■ Set the compile-time flag <code>_LARGEFILE64_SOURCE</code> to 1 before including any headers.■ Use the <code>getconf(1)</code> command with one or more of the following arguments: <table border="1" data-bbox="446 1623 1414 1734"><thead><tr><th>argument</th><th>purpose</th></tr></thead><tbody><tr><td><code>LFS64_CFLAGS</code></td><td>obtain compilation flags necessary to enable the transitional compilation environment</td></tr></tbody></table>	argument	purpose	<code>LFS64_CFLAGS</code>	obtain compilation flags necessary to enable the transitional compilation environment
argument	purpose				
<code>LFS64_CFLAGS</code>	obtain compilation flags necessary to enable the transitional compilation environment				

argument	purpose
LFS64_LDFLAGS	obtain link editor options
LFS64_LIBS	obtain link library names
LFS64_LINTFLAGS	obtain lint options

EXAMPLES

In the following examples, the transitional compilation environment is accessed by invoking the `getconf` utility with one of the arguments listed in the table above. The additional large file interfaces are accessed either by specifying `-D_LARGEFILE64_SOURCE` or by invoking the `getconf` utility with the arguments listed above.

The example that uses the form of command substitution specifying the command within parentheses preceded by a dollar sign can be executed only in a POSIX-conforming shell such as the Korn Shell (see `ksh(1)`). In a shell that is not POSIX-conforming, such as the Bourne Shell (see `sh(1)`) and the C Shell (see `cs(1)`), the command must be enclosed within grave accent marks.

EXAMPLE 1 An example of compiling a program using transitional interfaces such as `lseek64()` and `fopen64()`:

```
$ c89 -D_LARGEFILE64_SOURCE      \
      $(getconf LFS64_CFLAGS) a.c \
      $(getconf LFS64_LDFLAGS)   \
      $(getconf LFS64_LIBS)
```

EXAMPLE 2 An example of running lint on a program using transitional interfaces:

```
% lint -D_LARGEFILE64_SOURCE      \
      `getconf LFS64_LINTFLAGS` ... \
      `getconf LFS64_LIBS`
```

SEE ALSO `getconf(1)`, `lseek(2)`, `fopen(3C)`, `lf64(5)`, `standards(5)`

live_upgrade(5)

NAME	live_upgrade – overview of Live Upgrade feature
DESCRIPTION	<p>The Live Upgrade feature of the Solaris operating environment enables you to maintain multiple operating system images on a single system. An image—called a boot environment, or BE—represents a set of operating system and application software packages. The BEs might contain different operating system and/or application versions.</p> <p>On a system with the Solaris Live Upgrade software, your currently booted OS environment is referred to as your active, or current BE. You have one active, or current BE; all others are inactive. You can perform any number of modifications to inactive BEs on the same system, then boot from one of those BEs. If there is a failure or some undesired behavior in the newly booted BE, Live Upgrade software makes it easy for you to fall back to the previously running BE.</p> <p>Live Upgrade software includes a full suite of commands, listed below and described in individual man pages, which implement all of the Live Upgrade features and functions. The software also includes a Forms and Menu Language Interpreter-based user interface named <code>lu(1M)</code>. (See <code>fml(1)</code> for a description of the Forms and Menu Language Interpreter.) The FMLI interface implements a subset of Live Upgrade functions. Unlike the command-line interfaces, output from the FMLI interface is not internationalizable.</p> <p>The following are some of the tasks you can perform with Live Upgrade software:</p> <ul style="list-style-type: none">■ You can make one or more copies of the currently running system.■ You can upgrade to a new OS version on a second boot environment, then boot from that environment. If you choose, you can then fall back to your original boot environment or boot from yet another environment.■ You can install application or OS packages to a boot environment, then boot from that environment.■ You can install OS patches to a boot environment, then boot from that environment.■ From a flash archive, you can install an OS to a boot environment, then boot from that environment. See <code>flar(1M)</code> for information on administering flash archives.■ You can split and rejoin file systems in a new BE. For example, you can separate <code>/usr</code>, <code>/var</code>, and <code>/opt</code> from <code>/</code>, putting them on their own partitions. Conversely, you could join these file systems on a single partition under <code>/</code>.■ You can mount any or all of the filesystems of a BE that is not active, compare the files in any pair of BEs, delete or rename a BE, and perform other administrative tasks. <p>The Live Upgrade software supports upgrade from any valid Solaris installation medium, including a CD-ROM, an NFS or UFS directory, or a flash archive. (See <code>flash_archive(4)</code> for a description of the flash archive feature.)</p>

In simplest terms, a BE, for Live Upgrade, consists of the disk slice containing a root file system and the file system/device (usually disk) slice entries specified in `vfstab(4)`. This set of slices is not limited to a single disk. This means that you can have multiple BEs on a single device, or have a BE spread across slices on multiple devices.

The minimal requirement for a Live Upgrade BE is the same as for any Solaris boot environment: you must have root (/) and `usr` filesystems (which might both reside on /). All filesystems except for /, /usr, /var, and /opt can be shared among multiple BEs, if you choose.

Each BE must have a unique copy of the file systems that contain the OS—/, /usr, /var, and /opt. For Live Upgrade purposes, these are referred to as non-shareable (sometimes referred to as *critical*) file systems. With other file systems, such as /export or /home, you have the option of copying the files to a new BE or, the default, sharing them among BEs. These are referred to as shareable file systems. A BE is made up of a unique copy of one or more non-shareable file systems and zero or more copies of shareable file systems.

Live Upgrade commands support an option (-X) that enables XML output. Characteristics of the XML are specified in a DTD shipped with the product. XML output enables programmatic parsing of portions of the command output.

Live Upgrade supports the notion of a BE description, an optional attribute of a BE. A BE description can be of any length and format. It might be a text string or a binary file. See `ludes(1M)` for details.

Below is an example set of steps that you might follow in the use of Live Upgrade software. These steps specify the use of commands rather than `lu(1M)`, the FMLI interface. Many Live Upgrade functions are accessible through `lu`. Except where `lu` does not support a function, the choice between `lu` and Live Upgrade commands is a matter of your requirements and preferences. The following example is by no means exhaustive of the possibilities of the use of the Live Upgrade software.

1. You create a new BE, using `lucreate(1M)`. The first time you create a BE on a given system, you must designate the current Solaris operating environment as a BE (give it a name). You then specify a name and a set of device (disk) slices you want to use for the new BE. The `lucreate` command copies the contents of the current Solaris operating environment (now a BE) to the new BE.
After you have created additional BEs, you can use a BE other than the current BE as the source for a new BE. Also, you can create an empty BE onto which you can later install a flash archive.
2. Using `luupgrade(1M)`, you upgrade the OS version on your new BE (or on yet another BE you created with `lucreate`). The `luupgrade` enables you to upgrade an OS (from any valid Solaris installation medium, including a flash archive), add or remove packages (OS or application), and add or remove patches.
3. You use `luactivate(1M)` to make the new BE bootable. The next time you reboot your system, you will come up in the new BE.

live_upgrade(5)

4. Using `lucompare(1M)`, you compare the system files on two different BEs. This utility gives you a comprehensive list of the files that have differences.
5. Using `lumount(1M)`, you mount the filesystems of a BE that is not active, enabling you to make changes. When you are finished with the changes, use `luumount(1M)` to unmount the BE's file systems.
6. Upon booting a new BE, you discover a failure or some other undesirable behavior. Using the procedure specified in `luactivate`, you can fall back to the previous BE.
7. Using `ludelete` then `lucreate`, you reassign file systems on the now-deleted BE to different disk slices. You separate `/opt` and `/var` from `/` on the new BE. Also, you specify that swap be spread over slices on multiple disks.

The following is a summary of Live Upgrade commands. All commands require root privileges.

`lu`

FMLI-based interface for creating and administering BEs.

`luactivate`

Designate a BE as the BE to boot from upon the next reboot of the system.

`lucancel`

Cancel a previously scheduled operation.

`lucompare`

Compare the contents of two BEs.

`lucreate`

Create a BE.

`lucurr`

Display the name of the current BE.

`ludelete`

Delete a BE.

`ludesc`

Add or change BE descriptions.

`lufslist`

List the file systems on a specified BE.

`lumake`

Re-create a BE based on the active BE.

`lumount, luumount`

Mount, unmount file systems of a specified BE.

`lurename`

Rename a BE.

lustatus

For all BEs on a system, report on whether a BE is active, active upon the next reboot, in the midst of a copy operation, and whether a copy operation is scheduled for it.

luupgrade

Upgrade an OS and install application software on a BE. Such software includes flash archives, complete OS installations, OS and application packages, and OS patches.

FILES

/etc/lutab
list of BEs on the system

SEE ALSO

lu(1M), luactivate(1M), lucancel(1M), lucompare(1M), lucreate(1M), lucurr(1M), ludelete(1M), ludescl(1M), lufslis(1M), lumake(1M), lumount(1M), lurename(1M), lustatus(1M), luupgrade(1M), lutab(4)

NOTES

Correct operation of Solaris Live Upgrade requires that a limited set of patch revisions be installed for a given OS version. Before installing or running Live Upgrade, you are required to install the limited set of patch revisions. Make sure you have the most recently updated patch list by consulting <http://sunsolve.sun.com>. Search for the infodoc 72099 on the SunSolve web site.

It is possible for an operating system upgrade to remove installed patches. Prior to such an upgrade, use `analyze_patches`, as described in `luupgrade(1M)`, to determine which, if any, patches will be removed.

Live Upgrade supports the release it is distributed on and up to three marketing releases back. For example, if you obtained Live Upgrade with Solaris 9 (including a Solaris 9 upgrade), that version of Live Upgrade supports Solaris versions 2.6, Solaris 7, and Solaris 8, in addition to Solaris 9. No version of Live Upgrade supports a Solaris version prior to Solaris 2.6.

locale(5)

NAME	locale – subset of a user’s environment that depends on language and cultural conventions												
DESCRIPTION	<p>A <code>locale</code> is the definition of the subset of a user’s environment that depends on language and cultural conventions. It is made up from one or more categories. Each category is identified by its name and controls specific aspects of the behavior of components of the system. Category names correspond to the following environment variable names:</p> <table><tr><td><code>LC_CTYPE</code></td><td>Character classification and case conversion.</td></tr><tr><td><code>LC_COLLATE</code></td><td>Collation order.</td></tr><tr><td><code>LC_TIME</code></td><td>Date and time formats.</td></tr><tr><td><code>LC_NUMERIC</code></td><td>Numeric formatting.</td></tr><tr><td><code>LC_MONETARY</code></td><td>Monetary formatting.</td></tr><tr><td><code>LC_MESSAGES</code></td><td>Formats of informative and diagnostic messages and interactive responses.</td></tr></table> <p>The standard utilities base their behavior on the current locale, as defined in the <code>ENVIRONMENT</code> section for each utility. The behavior of some of the C-language functions will also be modified based on the current locale, as defined by the last call to <code>setlocale(3C)</code>.</p> <p>Locales other than those supplied by the implementation can be created by the application via the <code>localedef(1)</code> utility. The value that is used to specify a locale when using environment variables will be the string specified as the <i>name</i> operand to <code>localedef</code> when the locale was created. The strings "C" and "POSIX" are reserved as identifiers for the POSIX locale.</p> <p>Applications can select the desired locale by invoking the <code>setlocale()</code> function with the appropriate value. If the function is invoked with an empty string, such as:</p> <pre>setlocale(LC_ALL, "");</pre> <p>the value of the corresponding environment variable is used. If the environment variable is unset or is set to the empty string, the <code>setlocale()</code> function sets the appropriate environment.</p>	<code>LC_CTYPE</code>	Character classification and case conversion.	<code>LC_COLLATE</code>	Collation order.	<code>LC_TIME</code>	Date and time formats.	<code>LC_NUMERIC</code>	Numeric formatting.	<code>LC_MONETARY</code>	Monetary formatting.	<code>LC_MESSAGES</code>	Formats of informative and diagnostic messages and interactive responses.
<code>LC_CTYPE</code>	Character classification and case conversion.												
<code>LC_COLLATE</code>	Collation order.												
<code>LC_TIME</code>	Date and time formats.												
<code>LC_NUMERIC</code>	Numeric formatting.												
<code>LC_MONETARY</code>	Monetary formatting.												
<code>LC_MESSAGES</code>	Formats of informative and diagnostic messages and interactive responses.												
Locale Definition	<p>Locales can be described with the file format accepted by the <code>localedef</code> utility.</p> <p>The locale definition file must contain one or more locale category source definitions, and must not contain more than one definition for the same locale category.</p> <p>A category source definition consists of a category header, a category body and a category trailer. A category header consists of the character string naming of the category, beginning with the characters <code>LC_</code>. The category trailer consists of the string <code>END</code>, followed by one or more blank characters and the string used in the corresponding category header.</p>												

The category body consists of one or more lines of text. Each line contains an identifier, optionally followed by one or more operands. Identifiers are either keywords, identifying a particular locale element, or collating elements. Each keyword within a locale must have a unique name (that is, two categories cannot have a commonly-named keyword); no keyword can start with the characters `LC_`. Identifiers must be separated from the operands by one or more blank characters.

Operands must be characters, collating elements or strings of characters. Strings must be enclosed in double-quotes. Literal double-quotes within strings must be preceded by the *<escape character>*, described below. When a keyword is followed by more than one operand, the operands must be separated by semicolons; blank characters are allowed both before and after a semicolon.

The first category header in the file can be preceded by a line modifying the comment character. It has the following format, starting in column 1:

```
"comment_char %c\n" , <comment character>
```

The comment character defaults to the number sign (#). Blank lines and lines containing the *<comment character>* in the first position are ignored.

The first category header in the file can be preceded by a line modifying the escape character to be used in the file. It has the following format, starting in column 1:

```
"escape_char %c\n" , <escape character>
```

The escape character defaults to backslash.

A line can be continued by placing an escape character as the last character on the line; this continuation character will be discarded from the input. Although the implementation need not accept any one portion of a continued line with a length exceeding `{LINE_MAX}` bytes, it places no limits on the accumulated length of the continued line. Comment lines cannot be continued on a subsequent line using an escaped newline character.

Individual characters, characters in strings, and collating elements must be represented using symbolic names, as defined below. In addition, characters can be represented using the characters themselves or as octal, hexadecimal or decimal constants. When non-symbolic notation is used, the resultant locale definitions will in many cases not be portable between systems. The left angle bracket (<) is a reserved symbol, denoting the start of a symbolic name; when used to represent itself it must be preceded by the escape character. The following rules apply to character representation:

1. A character can be represented via a symbolic name, enclosed within angle brackets < and >. The symbolic name, including the angle brackets, must exactly match a symbolic name defined in the charmap file specified via the `localedef -f` option, and will be replaced by a character value determined from the value associated with the symbolic name in the charmap file. The use of a symbolic name not found in the charmap file constitutes an error, unless the category is `LC_CTYPE`

or LC_COLLATE, in which case it constitutes a warning condition (see `localedef(1)` for a description of action resulting from errors and warnings). The specification of a symbolic name in a `collating-element` or `collating-symbol` section that duplicates a symbolic name in the charmap file (if present) is an error. Use of the escape character or a right angle bracket within a symbolic name is invalid unless the character is preceded by the escape character.

Example:

```
<c>;<c-cedilla> "<M><a><y>"
```

2. A character can be represented by the character itself, in which case the value of the character is implementation-dependent. Within a string, the double-quote character, the escape character and the right angle bracket character must be escaped (preceded by the escape character) to be interpreted as the character itself. Outside strings, the characters

```
, ; < > escape_char
```

must be escaped to be interpreted as the character itself.

For example:

```
c beta-char "May"
```

3. A character can be represented as an octal constant. An octal constant is specified as the escape character followed by two or more octal digits. Each constant represents a byte value. Multi-byte values can be represented by concatenated constants specified in byte order with the last constant specifying the least significant byte of the character.

For example:

```
\143;\347;\143\150 "\115\141\171"
```

4. A character can be represented as a hexadecimal constant. A hexadecimal constant is specified as the escape character followed by an `x` followed by two or more hexadecimal digits. Each constant represents a byte value. Multi-byte values can be represented by concatenated constants specified in byte order with the last constant specifying the least significant byte of the character.

For example:

```
\x63;\xe7;\x63\x68 "\x4d\x61\x79"
```

5. A character can be represented as a decimal constant. A decimal constant is specified as the escape character followed by a `d` followed by two or more decimal digits. Each constant represents a byte value. Multi-byte values can be represented by concatenated constants specified in byte order with the last constant specifying the least significant byte of the character.

For example:

```
\d99;\d231;\d99\d104 "\d77\d97\d121"
```

Only characters existing in the character set for which the locale definition is created can be specified, whether using symbolic names, the characters themselves, or octal, decimal or hexadecimal constants. If a charmap file is present, only characters defined in the charmap can be specified using octal, decimal or hexadecimal constants. Symbolic names not present in the charmap file can be specified and will be ignored, as specified under item 1 above.

LC_CTYPE

The `LC_CTYPE` category defines character classification, case conversion and other character attributes. In addition, a series of characters can be represented by three adjacent periods representing an ellipsis symbol (. . .). The ellipsis specification is interpreted as meaning that all values between the values preceding and following it represent valid characters. The ellipsis specification is valid only within a single encoded character set; that is, within a group of characters of the same size. An ellipsis is interpreted as including in the list all characters with an encoded value higher than the encoded value of the character preceding the ellipsis and lower than the encoded value of the character following the ellipsis.

For example:

```
\x30;. . .;\x39;
```

includes in the character class all characters with encoded values between the endpoints.

The following keywords are recognized. In the descriptions, the term “automatically included” means that it is not an error either to include or omit any of the referenced characters.

The character classes `digit`, `xdigit`, `lower`, `upper`, and `space` have a set of automatically included characters. These only need to be specified if the character values (that is, encoding) differ from the implementation default values.

<code>cswidth</code>	Moved to <code>extensions</code> file (see <code>extensions(5)</code>).
<code>upper</code>	<p>Define characters to be classified as upper-case letters.</p> <p>In the POSIX locale, the 26 upper-case letters are included:</p> <pre>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</pre> <p>In a locale definition file, no character specified for the keywords <code>cntrl</code>, <code>digit</code>, <code>punct</code>, or <code>space</code> can be specified. The upper-case letters A to Z are automatically included in this class.</p>
<code>lower</code>	<p>Define characters to be classified as lower-case letters.</p> <p>In the POSIX locale, the 26 lower-case letters are included:</p> <pre>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</pre>

locale(5)

	<p>In a locale definition file, no character specified for the keywords <code>cntrl</code>, <code>digit</code>, <code>punct</code>, or <code>space</code> can be specified. The lower-case letters <code>a</code> to <code>z</code> of the portable character set are automatically included in this class.</p>
<code>alpha</code>	<p>Define characters to be classified as letters.</p> <p>In the POSIX locale, all characters in the classes <code>upper</code> and <code>lower</code> are included.</p> <p>In a locale definition file, no character specified for the keywords <code>cntrl</code>, <code>digit</code>, <code>punct</code>, or <code>space</code> can be specified. Characters classified as either <code>upper</code> or <code>lower</code> are automatically included in this class.</p>
<code>digit</code>	<p>Define the characters to be classified as numeric digits.</p> <p>In the POSIX locale, only</p> <p>0 1 2 3 4 5 6 7 8 9</p> <p>are included.</p> <p>In a locale definition file, only the digits 0, 1, 2, 3, 4, 5, 6, 7, 8, and 9 can be specified, and in contiguous ascending sequence by numerical value. The digits 0 to 9 of the portable character set are automatically included in this class.</p> <p>The definition of character class <code>digit</code> requires that only ten characters; the ones defining digits can be specified; alternative digits (for example, Hindi or Kanji) cannot be specified here.</p>
<code>space</code>	<p>Define characters to be classified as white-space characters.</p> <p>In the POSIX locale, at a minimum, the characters <code>SPACE</code>, <code>FORMFEED</code>, <code>NEWLINE</code>, <code>CARRIAGE RETURN</code>, <code>TAB</code>, and <code>VERTICAL TAB</code> are included.</p> <p>In a locale definition file, no character specified for the keywords <code>upper</code>, <code>lower</code>, <code>alpha</code>, <code>digit</code>, <code>graph</code>, or <code>xdigit</code> can be specified. The characters <code>SPACE</code>, <code>FORMFEED</code>, <code>NEWLINE</code>, <code>CARRIAGE RETURN</code>, <code>TAB</code>, and <code>VERTICAL TAB</code> of the portable character set, and any characters included in the class <code>blank</code> are automatically included in this class.</p>
<code>cntrl</code>	<p>Define characters to be classified as control characters.</p>

punct	<p>In the POSIX locale, no characters in classes <code>alpha</code> or <code>print</code> are included.</p> <p>In a locale definition file, no character specified for the keywords <code>upper</code>, <code>lower</code>, <code>alpha</code>, <code>digit</code>, <code>punct</code>, <code>graph</code>, <code>print</code>, or <code>xdigit</code> can be specified.</p> <p>Define characters to be classified as punctuation characters.</p> <p>In the POSIX locale, neither the space character nor any characters in classes <code>alpha</code>, <code>digit</code>, or <code>cntrl</code> are included.</p>
graph	<p>In a locale definition file, no character specified for the keywords <code>upper</code>, <code>lower</code>, <code>alpha</code>, <code>digit</code>, <code>cntrl</code>, <code>xdigit</code> or as the space character can be specified.</p> <p>Define characters to be classified as printable characters, not including the space character.</p> <p>In the POSIX locale, all characters in classes <code>alpha</code>, <code>digit</code>, and <code>punct</code> are included; no characters in class <code>cntrl</code> are included.</p>
print	<p>In a locale definition file, characters specified for the keywords <code>upper</code>, <code>lower</code>, <code>alpha</code>, <code>digit</code>, <code>xdigit</code>, and <code>punct</code> are automatically included in this class. No character specified for the keyword <code>cntrl</code> can be specified.</p> <p>Define characters to be classified as printable characters, including the space character.</p> <p>In the POSIX locale, all characters in class <code>graph</code> are included; no characters in class <code>cntrl</code> are included.</p>
xdigit	<p>In a locale definition file, characters specified for the keywords <code>upper</code>, <code>lower</code>, <code>alpha</code>, <code>digit</code>, <code>xdigit</code>, <code>punct</code>, and the space character are automatically included in this class. No character specified for the keyword <code>cntrl</code> can be specified.</p> <p>Define the characters to be classified as hexadecimal digits.</p> <p>In the POSIX locale, only:</p> <p style="margin-left: 2em;">0 1 2 3 4 5 6 7 8 9 A B C D E F a b c d e f</p> <p>are included.</p>

blank	<p>In a locale definition file, only the characters defined for the class <code>digit</code> can be specified, in contiguous ascending sequence by numerical value, followed by one or more sets of six characters representing the hexadecimal digits 10 to 15 inclusive, with each set in ascending order (for example A, B, C, D, E, F, a, b, c, d, e, f). The digits 0 to 9, the upper-case letters A to F and the lower-case letters a to f of the portable character set are automatically included in this class.</p> <p>The definition of character class <code>xdigit</code> requires that the characters included in character class <code>digit</code> be included here also.</p> <p>Define characters to be classified as blank characters.</p> <p>In the POSIX locale, only the space and tab characters are included.</p>
charclass	<p>In a locale definition file, the characters space and tab are automatically included in this class.</p> <p>Define one or more locale-specific character class names as strings separated by semi-colons. Each named character class can then be defined subsequently in the <code>LC_CTYPE</code> definition. A character class name consists of at least one and at most <code>{CHARCLASS_NAME_MAX}</code> bytes of alphanumeric characters from the portable filename character set. The first character of a character class name cannot be a digit. The name cannot match any of the <code>LC_CTYPE</code> keywords defined in this document.</p>
charclass-name	<p>Define characters to be classified as belonging to the named locale-specific character class. In the POSIX locale, the locale-specific named character classes need not exist. If a class name is defined by a <code>charclass</code> keyword, but no characters are subsequently assigned to it, this is not an error; it represents a class without any characters belonging to it. The <code>charclass-name</code> can be used as the <i>property</i> argument to the <code>wctype(3C)</code> function, in regular expression and shell pattern-matching bracket expressions, and by the <code>tr(1)</code> command.</p>
toupper	<p>Define the mapping of lower-case letters to upper-case letters.</p> <p>In the POSIX locale, at a minimum, the 26 lower-case characters:</p>

locale(5)

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

are mapped to the corresponding 26 upper-case characters:

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

In a locale definition file, the operand consists of character pairs, separated by semicolons. The characters in each character pair are separated by a comma and the pair enclosed by parentheses. The first character in each pair is the lower-case letter, the second the corresponding upper-case letter. Only characters specified for the keywords `lower` and `upper` can be specified. The lower-case letters a to z, and their corresponding upper-case letters A to Z, of the portable character set are automatically included in this mapping, but only when the `toupper` keyword is omitted from the locale definition.

`tolower`

Define the mapping of upper-case letters to lower-case letters.

In the POSIX locale, at a minimum, the 26 upper-case characters:

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

are mapped to the corresponding 26 lower-case characters:

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

In a locale definition file, the operand consists of character pairs, separated by semicolons. The characters in each character pair are separated by a comma and the pair enclosed by parentheses. The first character in each pair is the upper-case letter, the second the corresponding lower-case letter. Only characters specified for the keywords `lower` and `upper` can be specified. If the `tolower` keyword is omitted from the locale definition, the mapping will be the reverse mapping of the one specified for `toupper`.

LC_COLLATE

The `LC_COLLATE` category provides a collation sequence definition for numerous utilities (such as `sort(1)`, `uniq(1)`, and so forth), regular expression matching (see `regex(5)`), and the `strcoll(3C)`, `strxfrm(3C)`, `wscoll(3C)`, and `wcxfrm(3C)` functions.

A collation sequence definition defines the relative order between collating elements (characters and multi-character collating elements) in the locale. This order is expressed in terms of collation values; that is, by assigning each element one or more collation values (also known as collation weights). At least the following capabilities are provided:

1. **Multi-character collating elements.** Specification of multi-character collating elements (that is, sequences of two or more characters to be collated as an entity).
2. **User-defined ordering of collating elements.** Each collating element is assigned a collation value defining its order in the character (or basic) collation sequence. This ordering is used by regular expressions and pattern matching and, unless collation weights are explicitly specified, also as the collation weight to be used in sorting.
3. **Multiple weights and equivalence classes.** Collating elements can be assigned one or more (up to the limit `{COLL_WEIGHTS_MAX}`) collating weights for use in sorting. The first weight is hereafter referred to as the primary weight.
4. **One-to-Many mapping.** A single character is mapped into a string of collating elements.
5. **Equivalence class definition.** Two or more collating elements have the same collation value (primary weight).
6. **Ordering by weights.** When two strings are compared to determine their relative order, the two strings are first broken up into a series of collating elements; the elements in each successive pair of elements are then compared according to the relative primary weights for the elements. If equal, and more than one weight has been assigned, then the pairs of collating elements are recompared according to the relative subsequent weights, until either a pair of collating elements compare unequal or the weights are exhausted. The following keywords are recognized in a collation sequence definition. They are described in detail in the following sections.
7. **Define a collating-element symbol representing a multi-character collating element.** This keyword is optional.
8. **Define a collating symbol for use in collation order statements.** This keyword is optional.
9. **Define collation rules.** This statement is followed by one or more collation order statements, assigning character collation values and collation weights to collating elements.
10. **Specify the end of the collation-order statements.**

collating-element
keyword

In addition to the collating elements in the character set, the `collating-element` keyword is used to define multi-character collating elements. The syntax is:

```
"collating-element %s from \"%s\"\\n", <collating-symbol>, <string>
```

The *<collating-symbol>* operand is a symbolic name, enclosed between angle brackets (< and >), and must not duplicate any symbolic name in the current charmap file (if any), or any other symbolic name defined in this collation definition. The string operand is a string of two or more characters that collates as an entity. A *<collating-element>* defined via this keyword is only recognized with the LC_COLLATE category.

For example:

```
collating-element <ch> from "<c><h>"
collating-element <e-acute> from "<acute><e>"
collating-element <ll> from "ll"
```

collating-symbol
keyword

This keyword will be used to define symbols for use in collation sequence statements; that is, between the `order_start` and the `order_end` keywords. The syntax is:

```
"collating-symbol %s\n", <collating-symbol>
```

The *<collating-symbol>* is a symbolic name, enclosed between angle brackets (< and >), and must not duplicate any symbolic name in the current charmap file (if any), or any other symbolic name defined in this collation definition.

A *collating-symbol* defined via this keyword is only recognized with the LC_COLLATE category.

For example:

```
collating-symbol <UPPER_CASE>
collating-symbol <HIGH>
```

The *collating-symbol* keyword defines a symbolic name that can be associated with a relative position in the character order sequence. While such a symbolic name does not represent any collating element, it can be used as a weight.

order_start keyword

The `order_start` keyword must precede collation order entries and also defines the number of weights for this collation sequence definition and other collation rules.

The syntax of the `order_start` keyword is:

```
"order_start %s;%s;. . .;%s\n", <sort-rules>, <sort-rules>
```

The operands to the `order_start` keyword are optional. If present, the operands define rules to be applied when strings are compared. The number of operands define how many weights each element is assigned; if no operands are present, one forward operand is assumed. If present, the first operand defines rules to be applied when comparing strings using the first (primary) weight; the second when comparing strings using the second weight, and so on. Operands are separated by semicolons (;).

locale(5)

Each operand consists of one or more collation directives, separated by commas (,). If the number of operands exceeds the { COLL_WEIGHTS_MAX } limit, the utility will issue a warning message. The following directives will be supported:

forward	Specifies that comparison operations for the weight level proceed from start of string towards the end of string.
backward	Specifies that comparison operations for the weight level proceed from end of string towards the beginning of string.
position	Specifies that comparison operations for the weight level will consider the relative position of elements in the strings not subject to IGNORE . The string containing an element not subject to IGNORE after the fewest collating elements subject to IGNORE from the start of the compare will collate first. If both strings contain a character not subject to IGNORE in the same relative position, the collating values assigned to the elements will determine the ordering. In case of equality, subsequent characters not subject to IGNORE are considered in the same manner.

The directives `forward` and `backward` are mutually exclusive.

For example:

```
order_start forward;backward
```

If no operands are specified, a single `forward` operand is assumed. The character (and collating element) order is defined by the order in which characters and elements are specified between the `order_start` and `order_end` keywords. This character order is used in range expressions in regular expressions (see `regex(5)`). Weights assigned to the characters and elements define the collation sequence; in the absence of weights, the character order is also the collation sequence. The `position` keyword provides the capability to consider, in a compare, the relative position of characters not subject to IGNORE. As an example, consider the two strings "o-ring" and "or-ing". Assuming the hyphen is subject to IGNORE on the first pass, the two strings will compare equal, and the position of the hyphen is immaterial. On second pass, all characters except the hyphen are subject to IGNORE, and in the normal case the two strings would again compare equal. By taking position into account, the first collates before the second.

Collation Order

The `order_start` keyword is followed by collating identifier entries. The syntax for the collating element entries is

```
"%s %s;%s; . . . ;%s\n"<collating-identifier>, <weight>, <weight>, . . .
```

Each *collating-identifier* consists of either a character described in `Locale Definition` above, a *<collating-element>*, a *<collating-symbol>*, an ellipsis, or the special symbol `UNDEFINED`. The order in which collating elements are specified determines the character order sequence, such that each collating element compares less than the elements following it. The NUL character compares lower than any other character.

A *<collating-element>* is used to specify multi-character collating elements, and indicates that the character sequence specified via the *<collating-element>* is to be collated as a unit and in the relative order specified by its place.

A *<collating-symbol>* is used to define a position in the relative order for use in weights. No weights are specified with a *<collating-symbol>*.

The ellipsis symbol specifies that a sequence of characters will collate according to their encoded character values. It is interpreted as indicating that all characters with a coded character set value higher than the value of the character in the preceding line, and lower than the coded character set value for the character in the following line, in the current coded character set, will be placed in the character collation order between the previous and the following character in ascending order according to their coded character set values. An initial ellipsis is interpreted as if the preceding line specified the NUL character, and a trailing ellipsis as if the following line specified the highest coded character set value in the current coded character set. An ellipsis is treated as invalid if the preceding or following lines do not specify characters in the current coded character set.

The symbol `UNDEFINED` is interpreted as including all coded character set values not specified explicitly or via the ellipsis symbol. Such characters are inserted in the character collation order at the point indicated by the symbol, and in ascending order according to their coded character set values. If no `UNDEFINED` symbol is specified, and the current coded character set contains characters not specified in this section, the utility will issue a warning message and place such characters at the end of the character collation order.

The optional operands for each collation-element are used to define the primary, secondary, or subsequent weights for the collating element. The first operand specifies the relative primary weight, the second the relative secondary weight, and so on. Two or more collation-elements can be assigned the same weight; they belong to the same *equivalence class* if they have the same primary weight. Collation behaves as if, for each weight level, elements subject to `IGNORE` are removed, unless the `position` collation directive is specified for the corresponding level with the `order_start` keyword. Then each successive pair of elements is compared according to the relative weights for the elements. If the two strings compare equal, the process is repeated for the next weight level, up to the limit `{COLL_WEIGHTS_MAX}`.

Weights are expressed as characters described in `Locale Definition` above, *<collating-symbol>*s, *<collating-element>*s, an ellipsis, or the special symbol `IGNORE`. A single character, a *<collating-symbol>* or a *<collating-element>* represent the relative

position in the character collating sequence of the character or symbol, rather than the character or characters themselves. Thus, rather than assigning absolute values to weights, a particular weight is expressed using the relative order value assigned to a collating element based on its order in the character collation sequence.

One-to-many mapping is indicated by specifying two or more concatenated characters or symbolic names. For example, if the character `<eszet>` is given the string "`<s><s>`" as a weight, comparisons are performed as if all occurrences of the character `<eszet>` are replaced by `<s><s>` (assuming that `<s>` has the collating weight `<s>`). If it is necessary to define `<eszet>` and `<s><s>` as an equivalence class, then a collating element must be defined for the string `ss`.

All characters specified via an ellipsis will by default be assigned unique weights, equal to the relative order of characters. Characters specified via an explicit or implicit UNDEFINED special symbol will by default be assigned the same primary weight (that is, belong to the same equivalence class). An ellipsis symbol as a weight is interpreted to mean that each character in the sequence has unique weights, equal to the relative order of their character in the character collation sequence. The use of the ellipsis as a weight is treated as an error if the collating element is neither an ellipsis nor the special symbol UNDEFINED.

The special keyword IGNORE as a weight indicates that when strings are compared using the weights at the level where IGNORE is specified, the collating element is ignored; that is, as if the string did not contain the collating element. In regular expressions and pattern matching, all characters that are subject to IGNORE in their primary weight form an equivalence class.

An empty operand is interpreted as the collating element itself.

For example, the order statement:

```
<a>    <a>;<a>
```

is equal to:

```
<a>
```

An ellipsis can be used as an operand if the collating element was an ellipsis, and is interpreted as the value of each character defined by the ellipsis.

The collation order as defined in this section defines the interpretation of bracket expressions in regular expressions.

For example:

```
order_start                forward;backward
UNDEFINED                  IGNORE;IGNORE
```



```

<LOW>
<space>                <LOW>;<space>
. . .                  <LOW>;. . .
<a>                    <a>;<a>
<a-acute>              <a>;<a-acute>
<a-grave>              <a>;<a-grave>
<A>                    <a>;<A>
<A-acute>              <a>;<A-acute>
<A-grave>              <a>;<A-grave>
<ch>                   <ch>;<ch>
<Ch>                   <ch>;<Ch>
<s>                    <s>;<s>
<eszet>                "<s><s>" ; "<eszet><eszet>"
order_end

```

This example is interpreted as follows:

1. The UNDEFINED means that all characters not specified in this definition (explicitly or via the ellipsis) are ignored for collation purposes; for regular expression purposes they are ordered first.
2. All characters between <space> and <a> have the same primary equivalence class and individual secondary weights based on their ordinal encoded values.
3. All characters based on the upper- or lower-case character a belong to the same primary equivalence class.
4. The multi-character collating element <ch> is represented by the collating symbol <ch> and belongs to the same primary equivalence class as the multi-character collating element <Ch>.

order_end *keyword*

The collating order entries must be terminated with an `order_end` keyword.

LC_MONETARY

The LC_MONETARY category defines the rules and symbols that are used to format monetary numeric information. This information is available through the `localeconv(3C)` function

The following items are defined in this category of the locale. The item names are the keywords recognized by the `localedef(1)` utility when defining a locale. They are also similar to the member names of the `lconv` structure defined in `<locale.h>`. The `localeconv` function returns `{CHAR_MAX}` for unspecified integer items and the empty string ("") for unspecified or size zero string items.

locale(5)

In a locale definition file the operands are strings. For some keywords, the strings can contain only integers. Keywords that are not provided, string values set to the empty string (""), or integer keywords set to -1, are used to indicate that the value is not available in the locale.

<code>int_curr_symbol</code>	The international currency symbol. The operand is a four-character string, with the first three characters containing the alphabetic international currency symbol in accordance with those specified in the ISO 4217:1987 standard. The fourth character is the character used to separate the international currency symbol from the monetary quantity.
<code>currency_symbol</code>	The string used as the local currency symbol.
<code>mon_decimal_point</code>	The operand is a string containing the symbol that is used as the decimal delimiter (radix character) in monetary formatted quantities. In contexts where standards (such as the ISO C standard) limit the <code>mon_decimal_point</code> to a single byte, the result of specifying a multi-byte operand is unspecified.
<code>mon_thousands_sep</code>	The operand is a string containing the symbol that is used as a separator for groups of digits to the left of the decimal delimiter in formatted monetary quantities. In contexts where standards limit the <code>mon_thousands_sep</code> to a single byte, the result of specifying a multi-byte operand is unspecified.
<code>mon_grouping</code>	Define the size of each group of digits in formatted monetary quantities. The operand is a sequence of integers separated by semicolons. Each integer specifies the number of digits in each group, with the initial integer defining the size of the group immediately preceding the decimal delimiter, and the following integers defining the preceding groups. If the last integer is not -1, then the size of the previous group (if any) will be repeatedly used for the remainder of the digits. If the last integer is -1, then no further grouping will be performed. The following is an example of the interpretation of the <code>mon_grouping</code> keyword. Assuming that the value to be formatted is 123456789 and the <code>mon_thousands_sep</code> is ', then the following table shows the result. The third column shows the equivalent string in the ISO C standard that would be used by the <code>localeconv</code> function to accommodate this grouping.

mon_grouping	Formatted Value	ISO C String
3;-1	123456'789	"\3\177"
3	123'456'789	"\3"
3;2;-1	1234'56'789	"\3\2\177"
3;2	12'34'56'789	"\3\2"
-1	123456789	"\177"

In these examples, the octal value of `{CHAR_MAX}` is 177.

<code>positive_sign</code>	A string used to indicate a non-negative-valued formatted monetary quantity.
<code>negative_sign</code>	A string used to indicate a negative-valued formatted monetary quantity.
<code>int_frac_digits</code>	An integer representing the number of fractional digits (those to the right of the decimal delimiter) to be written in a formatted monetary quantity using <code>int_curr_symbol</code> .
<code>frac_digits</code>	An integer representing the number of fractional digits (those to the right of the decimal delimiter) to be written in a formatted monetary quantity using <code>currency_symbol</code> .
<code>p_cs_precedes</code>	An integer set to 1 if the <code>currency_symbol</code> or <code>int_curr_symbol</code> precedes the value for a monetary quantity with a non-negative value, and set to 0 if the symbol succeeds the value.
<code>p_sep_by_space</code>	An integer set to 0 if no space separates the <code>currency_symbol</code> or <code>int_curr_symbol</code> from the value for a monetary quantity with a non-negative value, set to 1 if a space separates the symbol from the value, and set to 2 if a space separates the symbol and the sign string, if adjacent.
<code>n_cs_precedes</code>	An integer set to 1 if the <code>currency_symbol</code> or <code>int_curr_symbol</code> precedes the value for a monetary quantity with a negative value, and set to 0 if the symbol succeeds the value.
<code>n_sep_by_space</code>	An integer set to 0 if no space separates the <code>currency_symbol</code> or <code>int_curr_symbol</code> from the value for a monetary quantity with a negative value,

locale(5)

set to 1 if a space separates the symbol from the value, and set to 2 if a space separates the symbol and the sign string, if adjacent.

`p_sign_posn` An integer set to a value indicating the positioning of the `positive_sign` for a monetary quantity with a non-negative value. The following integer values are recognized for both `p_sign_posn` and `n_sign_posn`:

- 0 Parentheses enclose the quantity and the `currency_symbol` or `int_curr_symbol`.
- 1 The sign string precedes the quantity and the `currency_symbol` or `int_curr_symbol`.
- 2 The sign string succeeds the quantity and the `currency_symbol` or `int_curr_symbol`.
- 3 The sign string precedes the `currency_symbol` or `int_curr_symbol`.
- 4 The sign string succeeds the `currency_symbol` or `int_curr_symbol`.

`n_sign_posn` An integer set to a value indicating the positioning of the `negative_sign` for a negative formatted monetary quantity.

The following table shows the result of various combinations:

		p_sep_by_space		
		2	1	0
<code>p_cs_precedes=1</code>	<code>p_sign_posn=0</code>	(\$1.25)	(\$1.25)	(\$1.25)
	<code>p_sign_posn=1</code>	+\$1.25	+\$1.25	+\$1.25
	<code>p_sign_posn=2</code>	\$1.25+	\$1.25+	\$1.25+
	<code>p_sign_posn=3</code>	+\$1.25	+\$1.25	+\$1.25
	<code>p_sign_posn=4</code>	\$+1.25	\$+1.25	\$+1.25
<code>p_cs_precedes=0</code>	<code>p_sign_posn=0</code>	(1.25 \$)	(1.25 \$)	(1.25\$)
	<code>p_sign_posn=1</code>	+1.25 \$	+1.25 \$	+1.25\$
	<code>p_sign_posn=2</code>	1.25\$ +	1.25 \$+	1.25\$+
	<code>p_sign_posn=3</code>	1.25+ \$	1.25 +\$	1.25+\$

```
p_sign_posn=4      1.25$ +      1.25 $+      1.25$+
```

The monetary formatting definitions for the POSIX locale follow; the code listing depicting the `localedef(1)` input, the table representing the same information with the addition of `localeconv(3C)` and `nl_langinfo(3C)` formats. All values are unspecified in the POSIX locale.

```
LC_MONETARY
# This is the POSIX locale definition for
# the LC_MONETARY category.
#

int_curr_symbol      ""
currency_symbol      ""
mon_decimal_point    ""
mon_thousands_sep   ""
mon_grouping         -1
positive_sign        ""
negative_sign        ""
int_frac_digits      -1
p_cs_precedes        -1
p_sep_by_space       -1
n_cs_precedes        -1
n_sep_by_space       -1
p_sign_posn         -1
n_sign_posn         -1
# END LC_MONETARY
```

The entry n/a indicates that the value is not available in the POSIX locale.

LC_NUMERIC

The `LC_NUMERIC` category defines the rules and symbols that will be used to format non-monetary numeric information. This information is available through the `localeconv(3C)` function.

The following items are defined in this category of the locale. The item names are the keywords recognized by the `localedef` utility when defining a locale. They are also similar to the member names of the `lconv` structure defined in `<locale.h>`. The `localeconv()` function returns `{CHAR_MAX}` for unspecified integer items and the empty string (`" "`) for unspecified or size zero string items.

locale(5)

In a locale definition file the operands are strings. For some keywords, the strings only can contain integers. Keywords that are not provided, string values set to the empty string (""), or integer keywords set to -1, will be used to indicate that the value is not available in the locale. The following keywords are recognized:

<code>decimal_point</code>	The operand is a string containing the symbol that is used as the decimal delimiter (radix character) in numeric, non-monetary formatted quantities. This keyword cannot be omitted and cannot be set to the empty string. In contexts where standards limit the <code>decimal_point</code> to a single byte, the result of specifying a multi-byte operand is unspecified.
<code>thousands_sep</code>	The operand is a string containing the symbol that is used as a separator for groups of digits to the left of the decimal delimiter in numeric, non-monetary formatted monetary quantities. In contexts where standards limit the <code>thousands_sep</code> to a single byte, the result of specifying a multi-byte operand is unspecified.
<code>grouping</code>	Define the size of each group of digits in formatted non-monetary quantities. The operand is a sequence of integers separated by semicolons. Each integer specifies the number of digits in each group, with the initial integer defining the size of the group immediately preceding the decimal delimiter, and the following integers defining the preceding groups. If the last integer is not -1, then the size of the previous group (if any) will be repeatedly used for the remainder of the digits. If the last integer is -1, then no further grouping will be performed. The non-monetary numeric formatting definitions for the POSIX locale follow; the code listing depicting the <code>localedef</code> input, the table representing the same information with the addition of <code>localeconv</code> values and <code>nl_langinfo</code> constants.

```
LC_NUMERIC
# This is the POSIX locale definition for
# the LC_NUMERIC category.
#
decimal_point    "<period>"
thousands_sep   ""
grouping         -1
#
END LC_NUMERIC
```

	POSIX locale	langinfo	localeconv()	localedef
Item	Value	Constant	Value	Value

decimal_point	"."	RADIXCHAR	"."	.
thousands_sep	n/a	THOUSEP	""	""
grouping	n/a	-	""	-1

The entry n/a indicates that the value is not available in the POSIX locale.

LC_TIME

The `LC_TIME` category defines the interpretation of the field descriptors supported by `date(1)` and affects the behavior of the `strptime(3C)`, `wcsftime(3C)`, `strptime(3C)`, and `nl_langinfo(3C)` functions. Because the interfaces for C-language access and locale definition differ significantly, they are described separately. For locale definition, the following mandatory keywords are recognized:

abday	Define the abbreviated weekday names, corresponding to the %a field descriptor (conversion specification in the <code>strptime()</code> , <code>wcsftime()</code> , and <code>strptime()</code> functions). The operand consists of seven semicolon-separated strings, each surrounded by double-quotes. The first string is the abbreviated name of the day corresponding to Sunday, the second the abbreviated name of the day corresponding to Monday, and so on.
day	Define the full weekday names, corresponding to the %A field descriptor. The operand consists of seven semicolon-separated strings, each surrounded by double-quotes. The first string is the full name of the day corresponding to Sunday, the second the full name of the day corresponding to Monday, and so on.
abmon	Define the abbreviated month names, corresponding to the %b field descriptor. The operand consists of twelve semicolon-separated strings, each surrounded by double-quotes. The first string is the abbreviated name of the first month of the year (January), the second the abbreviated name of the second month, and so on.
mon	Define the full month names, corresponding to the %B field descriptor. The operand consists of twelve semicolon-separated strings, each surrounded by double-quotes. The first string is the full name of the first month of the year (January), the second the full name of the second month, and so on.
d_t_fmt	Define the appropriate date and time representation, corresponding to the %c field descriptor. The operand consists of a string, and can contain any combination of characters and field descriptors. In addition, the string can contain the escape sequences <code>\\</code> , <code>\a</code> , <code>\b</code> , <code>\f</code> , <code>\n</code> , <code>\r</code> , <code>\t</code> , <code>\v</code> .
date_fmt	Define the appropriate date and time representation, corresponding to the %C field descriptor. The operand consists of a string, and can contain any combination of characters and field

locale(5)

	<p>descriptors. In addition, the string can contain the escape sequences <code>\\</code>, <code>\a</code>, <code>\b</code>, <code>\f</code>, <code>\n</code>, <code>\r</code>, <code>\t</code>, <code>\v</code>.</p>
<code>d_fmt</code>	<p>Define the appropriate date representation, corresponding to the <code>%x</code> field descriptor. The operand consists of a string, and can contain any combination of characters and field descriptors. In addition, the string can contain the escape sequences <code>\\</code>, <code>\a</code>, <code>\b</code>, <code>\f</code>, <code>\n</code>, <code>\r</code>, <code>\t</code>, <code>\v</code>.</p>
<code>t_fmt</code>	<p>Define the appropriate time representation, corresponding to the <code>%X</code> field descriptor. The operand consists of a string, and can contain any combination of characters and field descriptors. In addition, the string can contain the escape sequences <code>\\</code>, <code>\a</code>, <code>\b</code>, <code>\f</code>, <code>\n</code>, <code>\r</code>, <code>\t</code>, <code>\v</code>.</p>
<code>am_pm</code>	<p>Define the appropriate representation of the <i>ante meridiem</i> and <i>post meridiem</i> strings, corresponding to the <code>%p</code> field descriptor. The operand consists of two strings, separated by a semicolon, each surrounded by double-quotes. The first string represents the <i>ante meridiem</i> designation, the last string the <i>post meridiem</i> designation.</p>
<code>t_fmt_ampm</code>	<p>Define the appropriate time representation in the 12-hour clock format with <code>am_pm</code>, corresponding to the <code>%r</code> field descriptor. The operand consists of a string and can contain any combination of characters and field descriptors. If the string is empty, the 12-hour format is not supported in the locale.</p>
<code>era</code>	<p>Define how years are counted and displayed for each era in a locale. The operand consists of semicolon-separated strings. Each string is an era description segment with the format:</p> <p><i>direction:offset:start_date:end_date:era_name:era_format</i></p> <p>according to the definitions below. There can be as many era description segments as are necessary to describe the different eras.</p> <p>The start of an era might not be the earliest point For example, the Christian era B.C. starts on the day before January 1, A.D. 1, and increases with earlier time.</p> <p><i>direction</i> Either a + or a – character. The + character indicates that years closer to the <i>start_date</i> have lower numbers than those closer to the <i>end_date</i>. The – character indicates that years closer to the <i>start_date</i> have higher numbers than those closer to the <i>end_date</i>.</p> <p><i>offset</i> The number of the year closest to the <i>start_date</i> in the era, corresponding to the <code>%Eg</code> and <code>%Ey</code> field descriptors.</p>

<i>start_date</i>	A date in the form <i>yyyy/mm/dd</i> , where <i>yyyy</i> , <i>mm</i> , and <i>dd</i> are the year, month and day numbers respectively of the start of the era. Years prior to A.D. 1 are represented as negative numbers.
<i>end_date</i>	The ending date of the era, in the same format as the <i>start_date</i> , or one of the two special values <i>-*</i> or <i>+</i> . The value <i>-*</i> indicates that the ending date is the beginning of time. The value <i>+</i> indicates that the ending date is the end of time.
<i>era_name</i>	A string representing the name of the era, corresponding to the %EC field descriptor.
<i>era_format</i>	A string for formatting the year in the era, corresponding to the %EG and %EY field descriptors.
<i>era_d_fmt</i>	Define the format of the date in alternative era notation, corresponding to the %Ex field descriptor.
<i>era_t_fmt</i>	Define the locale's appropriate alternative time format, corresponding to the %EX field descriptor.
<i>era_d_t_fmt</i>	Define the locale's appropriate alternative date and time format, corresponding to the %Ec field descriptor.
<i>alt_digits</i>	Define alternative symbols for digits, corresponding to the %O field descriptor modifier. The operand consists of semicolon-separated strings, each surrounded by double-quotes. The first string is the alternative symbol corresponding with zero, the second string the symbol corresponding with one, and so on. Up to 100 alternative symbol strings can be specified. The %O modifier indicates that the string corresponding to the value specified by means of the field descriptor will be used instead of the value.

LC_TIME
C-language Access

The following information can be accessed. These correspond to constants defined in `<langinfo.h>` and used as arguments to the `nl_langinfo(3C)` function.

<i>ABDAY_x</i>	The abbreviated weekday names (for example Sun), where <i>x</i> is a number from 1 to 7.
<i>DAY_x</i>	The full weekday names (for example Sunday), where <i>x</i> is a number from 1 to 7.
<i>ABMON_x</i>	The abbreviated month names (for example Jan), where <i>x</i> is a number from 1 to 12.
<i>MON_x</i>	The full month names (for example January), where <i>x</i> is a number from 1 to 12.

locale(5)

D_T_FMT	The appropriate date and time representation.
D_FMT	The appropriate date representation.
T_FMT	The appropriate time representation.
AM_STR	The appropriate ante-meridiem affix.
PM_STR	The appropriate post-meridiem affix.
T_FMT_AMPM	The appropriate time representation in the 12-hour clock format with AM_STR and PM_STR.
ERA	<p>The era description segments, which describe how years are counted and displayed for each era in a locale. Each era description segment has the format:</p> <p><i>direction : offset : start_date : end_date : era_name : era_format</i></p> <p>according to the definitions below. There will be as many era description segments as are necessary to describe the different eras. Era description segments are separated by semicolons.</p> <p>The start of an era might not be the earliest point For example, the Christian era B.C. starts on the day before January 1, A.D. 1, and increases with earlier time.</p> <p><i>direction</i> Either a + or a – character. The + character indicates that years closer to the start_date have lower numbers than those closer to the end_date. The – character indicates that years closer to the start_date have higher numbers than those closer to the end_date.</p> <p><i>offset</i> The number of the year closest to the start_date in the era.</p> <p><i>start_date</i> A date in the form <i>yyyy/mm/dd</i>, where <i>yyyy</i>, <i>mm</i>, and <i>dd</i> are the year, month and day numbers respectively of the start of the era. Years prior to AD 1 are represented as negative numbers.</p> <p><i>end_date</i> The ending date of the era, in the same format as the <i>start_date</i>, or one of the two special values –* or +*. The value –* indicates that the ending date is the beginning of</p>

locale(5)

time. The value `+` indicates that the ending date is the end of time.

era_name The era, corresponding to the `%EC` conversion specification.

era_format The format of the year in the era, corresponding to the `%EY` and `%EY` conversion specifications.

`ERA_D_FMT` The era date format.

`ERA_T_FMT` The locale's appropriate alternative time format, corresponding to the `%EX` field descriptor.

`ERA_D_T_FMT` The locale's appropriate alternative date and time format, corresponding to the `%Ec` field descriptor.

`ALT_DIGITS` The alternative symbols for digits, corresponding to the `%O` conversion specification modifier. The value consists of semicolon-separated symbols. The first is the alternative symbol corresponding to zero, the second is the symbol corresponding to one, and so on. Up to 100 alternative symbols may be specified. The following table displays the correspondence between the items described above and the conversion specifiers used by `date(1)` and the `strftime(3C)`, `wcsftime(3C)`, and `strptime(3C)` functions.

localedef Keyword	langinfo Constant	Conversion Specifier
<code>abday</code>	<code>ABDAY_x</code>	<code>%a</code>
<code>day</code>	<code>DAY_x</code>	<code>%A</code>
<code>abmon</code>	<code>ABMON_x</code>	<code>%b</code>
<code>mon</code>	<code>MON</code>	<code>%B</code>
<code>d_t_fmt</code>	<code>D_T_FMT</code>	<code>%c</code>
<code>date_fmt</code>	<code>DATE_FMT</code>	<code>%C</code>
<code>d_fmt</code>	<code>D_FMT</code>	<code>%x</code>
<code>t_fmt</code>	<code>T_FMT</code>	<code>%X</code>
<code>am_pm</code>	<code>AM_STR</code>	<code>%p</code>
<code>am_pm</code>	<code>PM_STR</code>	<code>%p</code>
<code>t_fmt_ampm</code>	<code>T_FMT_AMP</code>	<code>%r</code>

locale(5)

localedef Keyword	langinfo Constant	Conversion Specifier
era	ERA	%EC, %Eg, %EG, %EY, %EY
era_d_fmt	ERA_D_FMT	%Ex
era_t_fmt	ERA_T_FMT	%EX
era_d_t_fmt	ERA_D_T_FMT	%Ec
alt_digits	ALT_DIGITS	%O

LC_TIME *General Information*

Although certain of the field descriptors in the POSIX locale (such as the name of the month) are shown with initial capital letters, this need not be the case in other locales. Programs using these fields may need to adjust the capitalization if the output is going to be used at the beginning of a sentence.

The LC_TIME descriptions of `abday`, `day`, `mon`, and `abmon` imply a Gregorian style calendar (7-day weeks, 12-month years, leap years, and so forth). Formatting time strings for other types of calendars is outside the scope of this document set.

As specified under `date` in *Locale Definition* and `strftime(3C)`, the field descriptors corresponding to the optional keywords consist of a modifier followed by a traditional field descriptor (for instance `%Ex`). If the optional keywords are not supported by the implementation or are unspecified for the current locale, these field descriptors are treated as the traditional field descriptor. For instance, assume the following keywords:

```
alt_digits    "0th" ; "1st" ; "2nd" ; "3rd" ; "4th" ; "5th" ; \
              "6th" ; "7th" ; "8th" ; "9th" ; "10th"
d_fmt        "The %d day of %B in %Y"
```

On 7/4/1776, the `%x` field descriptor would result in "The 4th day of July in 1776" while 7/14/1789 would come out as "The 14 day of July in 1789" It can be noted that the above example is for illustrative purposes only; the `%O` modifier is primarily intended to provide for Kanji or Hindi digits in `date` formats.

LC_MESSAGES

The LC_MESSAGES category defines the format and values for affirmative and negative responses.

The following keywords are recognized as part of the locale definition file. The `nl_langinfo(3C)` function accepts upper-case versions of the first four keywords.

`yesexpr` The operand consists of an extended regular expression (see `regex(5)`) that describes the acceptable affirmative response to a question expecting an affirmative or negative response.

noexpr The operand consists of an extended regular expression that describes the acceptable negative response to a question expecting an affirmative or negative response.

yesstr The operand consists of a fixed string (not a regular expression) that can be used by an application for composition of a message that lists an acceptable affirmative response, such as in a prompt.

nostr The operand consists of a fixed string that can be used by an application for composition of a message that lists an acceptable negative response. The format and values for affirmative and negative responses of the POSIX locale follow; the code listing depicting the `localedef` input, the table representing the same information with the addition of `nl_langinfo()` constants.

```
LC_MESSAGES
# This is the POSIX locale definition for
# the LC_MESSAGES category.
#
yesexpr "<circumflex><left-square-bracket><y><Y><right-square-bracket>"
#
noexpr "<circumflex><left-square-bracket><n><N><right-square-bracket>"
#
yesstr "yes"
nostr "no"
END LC_MESSAGES
```

localedef Keyword	langinfo Constant	POSIX Locale Value
yesexpr	YESEXPR	"^[yY]"
noexpr	NOEXPR	"^[nN]"
yesstr	YESSTR	"yes"
nostr	NOSTR	"no"

SEE ALSO `date(1)`, `locale(1)`, `localedef(1)`, `sort(1)`, `tr(1)`, `uniq(1)`, `localeconv(3C)`, `nl_langinfo(3C)`, `setlocale(3C)`, `strcoll(3C)`, `strftime(3C)`, `strptime(3C)`, `strxfrm(3C)`, `wscoll(3C)`, `wcsftime(3C)`, `wcsxfrm(3C)`, `wctype(3C)`, `attributes(5)`, `charmap(5)`, `extensions(5)`, `regex(5)`

man(5)

NAME	man – macros to format Reference Manual pages																																												
SYNOPSIS	nroff -man <i>filename...</i> troff -man <i>filename...</i>																																												
DESCRIPTION	<p>These macros are used to lay out the reference pages in this manual. Note: if <i>filename</i> contains format input for a preprocessor, the commands shown above must be piped through the appropriate preprocessor. This is handled automatically by the man(1) command. See the “Conventions” section.</p> <p>Any text argument <i>t</i> may be zero to six words. Quotes may be used to include SPACE characters in a “word”. If <i>text</i> is empty, the special treatment is applied to the next input line with text to be printed. In this way . I may be used to italicize a whole line, or .SB may be used to make small bold letters.</p> <p>A prevailing indent distance is remembered between successive indented paragraphs, and is reset to default value upon reaching a non-indented paragraph. Default units for indents <i>i</i> are ens.</p> <p>Type font and size are reset to default values before each paragraph, and after processing font and size setting macros.</p> <p>These strings are predefined by -man:</p> <p>*R ‘®’, ‘(Reg)’ in nroff. *S Change to default type size.</p>																																												
Requests	<p>* n.t.l. = next text line; p.i. = prevailing indent</p> <table><thead><tr><th><i>Request</i></th><th><i>Cause</i></th><th><i>If no</i></th><th><i>Explanation</i></th></tr></thead><tbody><tr><td></td><td><i>Break</i></td><td><i>Argument</i></td><td></td></tr><tr><td>.B <i>t</i></td><td>no</td><td><i>t=n.t.l.*</i></td><td>Text is in bold font.</td></tr><tr><td>.BI <i>t</i></td><td>no</td><td><i>t=n.t.l.</i></td><td>Join words, alternating bold and italic.</td></tr><tr><td>.BR <i>t</i></td><td>no</td><td><i>t=n.t.l.</i></td><td>Join words, alternating bold and roman.</td></tr><tr><td>.DT</td><td>no</td><td>.5i 1i...</td><td>Restore default tabs.</td></tr><tr><td>.HP <i>i</i></td><td>yes</td><td><i>i=p.i.*</i></td><td>Begin paragraph with hanging indent. Set prevailing indent to <i>i</i>.</td></tr><tr><td>.I <i>t</i></td><td>no</td><td><i>t=n.t.l.</i></td><td>Text is italic.</td></tr><tr><td>.IB <i>t</i></td><td>no</td><td><i>t=n.t.l.</i></td><td>Join words, alternating italic and bold.</td></tr><tr><td>.IP <i>x i</i></td><td>yes</td><td><i>x=""</i></td><td>Same as .TP with tag <i>x</i>.</td></tr><tr><td>.IR <i>t</i></td><td>no</td><td><i>t=n.t.l.</i></td><td>Join words, alternating italic and roman.</td></tr></tbody></table>	<i>Request</i>	<i>Cause</i>	<i>If no</i>	<i>Explanation</i>		<i>Break</i>	<i>Argument</i>		.B <i>t</i>	no	<i>t=n.t.l.*</i>	Text is in bold font.	.BI <i>t</i>	no	<i>t=n.t.l.</i>	Join words, alternating bold and italic.	.BR <i>t</i>	no	<i>t=n.t.l.</i>	Join words, alternating bold and roman.	.DT	no	.5i 1i...	Restore default tabs.	.HP <i>i</i>	yes	<i>i=p.i.*</i>	Begin paragraph with hanging indent. Set prevailing indent to <i>i</i> .	.I <i>t</i>	no	<i>t=n.t.l.</i>	Text is italic.	.IB <i>t</i>	no	<i>t=n.t.l.</i>	Join words, alternating italic and bold.	.IP <i>x i</i>	yes	<i>x=""</i>	Same as .TP with tag <i>x</i> .	.IR <i>t</i>	no	<i>t=n.t.l.</i>	Join words, alternating italic and roman.
<i>Request</i>	<i>Cause</i>	<i>If no</i>	<i>Explanation</i>																																										
	<i>Break</i>	<i>Argument</i>																																											
.B <i>t</i>	no	<i>t=n.t.l.*</i>	Text is in bold font.																																										
.BI <i>t</i>	no	<i>t=n.t.l.</i>	Join words, alternating bold and italic.																																										
.BR <i>t</i>	no	<i>t=n.t.l.</i>	Join words, alternating bold and roman.																																										
.DT	no	.5i 1i...	Restore default tabs.																																										
.HP <i>i</i>	yes	<i>i=p.i.*</i>	Begin paragraph with hanging indent. Set prevailing indent to <i>i</i> .																																										
.I <i>t</i>	no	<i>t=n.t.l.</i>	Text is italic.																																										
.IB <i>t</i>	no	<i>t=n.t.l.</i>	Join words, alternating italic and bold.																																										
.IP <i>x i</i>	yes	<i>x=""</i>	Same as .TP with tag <i>x</i> .																																										
.IR <i>t</i>	no	<i>t=n.t.l.</i>	Join words, alternating italic and roman.																																										

<i>Request</i>	<i>Cause</i>	<i>If no</i>	<i>Explanation</i>
	<i>Break</i>	<i>Argument</i>	
.IX <i>t</i>	no	-	Index macro, for SunSoft internal use.
.LP	yes	-	Begin left-aligned paragraph. Set prevailing indent to .5i.
.P	yes	-	Same as .LP.
.PD <i>d</i>	no	<i>d=.4v</i>	Set vertical distance between paragraphs.
.PP	yes	-	Same as .LP.
.RE	yes	-	End of relative indent. Restores prevailing indent.
.RB <i>t</i>	no	<i>t=n.t.l.</i>	Join words, alternating roman and bold.
.RI <i>t</i>	no	<i>t=n.t.l.</i>	Join words, alternating roman and italic.
.RS <i>i</i>	yes	<i>i=p.i.</i>	Start relative indent, increase indent by <i>i</i> . Sets prevailing indent to .5i for nested indents.
.SB <i>t</i>	no	-	Reduce size of text by 1 point, make text bold.
.SH <i>t</i>	yes	-	Section Heading.
.SM <i>t</i>	no	<i>t=n.t.l.</i>	Reduce size of text by 1 point.
.SS <i>t</i>	yes	<i>t=n.t.l.</i>	Section Subheading.
.TH <i>n s d f m</i>	yes	-	Begin reference page <i>n</i> , of of section <i>s</i> ; <i>d</i> is the date of the most recent change. If present, <i>f</i> is the left page footer; <i>m</i> is the main page (center) header. Sets prevailing indent and tabs to .5i.
.TP <i>i</i>	yes	<i>i=p.i.</i>	Begin indented paragraph, with the tag given on the next text line. Set prevailing indent to <i>i</i> .
.TX <i>t p</i>	no	-	Resolve the title abbreviation <i>t</i> ; join to punctuation mark (or text) <i>p</i> .

Conventions

When formatting a manual page, man examines the first line to determine whether it requires special processing. For example a first line consisting of:

```
' \" t
```

indicates that the manual page must be run through the `tbl(1)` preprocessor.

A typical manual page for a command or function is laid out as follows:

.TH <i>title</i> [1-9]	The name of the command or function, which serves as the title of the manual page. This is followed by the number of the section in which it appears.
.SH NAME	The name, or list of names, by which the command is called, followed by a dash and then a one-line summary of the action performed. All in roman font, this section contains no <code>troff(1)</code> commands or escapes, and no macro requests. It is used to generate the <code>windex</code> database, which is used by the <code>whatism(1)</code> command.
.SH SYNOPSIS	<p>Commands:</p> <p>The syntax of the command and its arguments, as typed on the command line. When in boldface, a word must be typed exactly as printed. When in italics, a word can be replaced with an argument that you supply. References to bold or italicized items are not capitalized in other sections, even when they begin a sentence.</p> <p>Syntactic symbols appear in roman face:</p> <p>[] An argument, when surrounded by brackets is optional.</p> <p> Arguments separated by a vertical bar are exclusive. You can supply only one item from such a list.</p> <p>... Arguments followed by an ellipsis can be repeated. When an ellipsis follows a bracketed set, the expression within the brackets can be repeated.</p> <p>Functions:</p> <p>If required, the data declaration, or <code>#include</code> directive, is shown first, followed by the function declaration. Otherwise, the function declaration is shown.</p>
.SH DESCRIPTION	A narrative overview of the command or function's external behavior. This includes how it interacts with files or data, and how it handles the standard input, standard output and standard error. Internals and implementation details are normally omitted. This section attempts to provide a succinct overview in answer to the question, "what does it do?"

Literal text from the synopsis appears in constant width, as do literal filenames and references to items that appear elsewhere in the reference manuals. Arguments are italicized.

If a command interprets either subcommands or an input grammar, its command interface or input grammar is normally described in a USAGE section, which follows the OPTIONS section. The DESCRIPTION section only describes the behavior of the command itself, not that of subcommands.

.SH OPTIONS	The list of options along with a description of how each affects the command's operation.
.SH RETURN VALUES	A list of the values the library routine will return to the calling program and the conditions that cause these values to be returned.
.SH EXIT STATUS	A list of the values the utility will return to the calling program or shell, and the conditions that cause these values to be returned.
.SH FILES	A list of files associated with the command or function.
.SH SEE ALSO	A comma-separated list of related manual pages, followed by references to other published materials.
.SH DIAGNOSTICS	A list of diagnostic messages and an explanation of each.
.SH BUGS	A description of limitations, known defects, and possible problems associated with the command or function.

FILES /usr/share/lib/tmac/an
/usr/share/man/windex

SEE ALSO man(1), nroff(1), troff(1), whatis(1)

Dale Dougherty and Tim O'Reilly, *Unix Text Processing*

mansun(5)

NAME	mansun – macros to format Reference Manual pages																																												
SYNOPSIS	nroff -mansun <i>filename...</i> troff -mansun <i>filename...</i>																																												
DESCRIPTION	<p>These macros are used to lay out the reference pages in this manual. Note: if <i>filename</i> contains format input for a preprocessor, the commands shown above must be piped through the appropriate preprocessor. This is handled automatically by <code>man(1)</code>. See the “Conventions” section.</p> <p>Any text argument <i>t</i> may be zero to six words. Quotes may be used to include SPACE characters in a “word”. If <i>text</i> is empty, the special treatment is applied to the next input line with text to be printed. In this way <code>.I</code> may be used to italicize a whole line, or <code>.SB</code> may be used to make small bold letters.</p> <p>A prevailing indent distance is remembered between successive indented paragraphs, and is reset to default value upon reaching a non-indented paragraph. Default units for indents <i>i</i> are ens.</p> <p>Type font and size are reset to default values before each paragraph, and after processing font and size setting macros.</p> <p>These strings are predefined by <code>-mansun</code>:</p> <p><code>*R</code> ‘®’, ‘(Reg)’ in <code>nroff</code>.</p> <p><code>*S</code> Change to default type size.</p>																																												
Requests	<p>* n.t.l. = next text line; p.i. = prevailing indent</p> <table><thead><tr><th><i>Request</i></th><th><i>Cause</i></th><th><i>If no</i></th><th><i>Explanation</i></th></tr><tr><td></td><td><i>Break</i></td><td><i>Argument</i></td><td></td></tr></thead><tbody><tr><td><code>.B t</code></td><td>no</td><td><code>t=n.t.l.*</code></td><td>Text is in bold font.</td></tr><tr><td><code>.BI t</code></td><td>no</td><td><code>t=n.t.l.</code></td><td>Join words, alternating bold and italic.</td></tr><tr><td><code>.BR t</code></td><td>no</td><td><code>t=n.t.l.</code></td><td>Join words, alternating bold and Roman.</td></tr><tr><td><code>.DT</code></td><td>no</td><td><code>.5i 1i...</code></td><td>Restore default tabs.</td></tr><tr><td><code>.HP i</code></td><td>yes</td><td><code>i=p.i.*</code></td><td>Begin paragraph with hanging indent. Set prevailing indent to <i>i</i>.</td></tr><tr><td><code>.I t</code></td><td>no</td><td><code>t=n.t.l.</code></td><td>Text is italic.</td></tr><tr><td><code>.IB t</code></td><td>no</td><td><code>t=n.t.l.</code></td><td>Join words, alternating italic and bold.</td></tr><tr><td><code>.IP x i</code></td><td>yes</td><td><code>x=""</code></td><td>Same as <code>.TP</code> with tag <i>x</i>.</td></tr><tr><td><code>.IR t</code></td><td>no</td><td><code>t=n.t.l.</code></td><td>Join words, alternating italic and Roman.</td></tr></tbody></table>	<i>Request</i>	<i>Cause</i>	<i>If no</i>	<i>Explanation</i>		<i>Break</i>	<i>Argument</i>		<code>.B t</code>	no	<code>t=n.t.l.*</code>	Text is in bold font.	<code>.BI t</code>	no	<code>t=n.t.l.</code>	Join words, alternating bold and italic.	<code>.BR t</code>	no	<code>t=n.t.l.</code>	Join words, alternating bold and Roman.	<code>.DT</code>	no	<code>.5i 1i...</code>	Restore default tabs.	<code>.HP i</code>	yes	<code>i=p.i.*</code>	Begin paragraph with hanging indent. Set prevailing indent to <i>i</i> .	<code>.I t</code>	no	<code>t=n.t.l.</code>	Text is italic.	<code>.IB t</code>	no	<code>t=n.t.l.</code>	Join words, alternating italic and bold.	<code>.IP x i</code>	yes	<code>x=""</code>	Same as <code>.TP</code> with tag <i>x</i> .	<code>.IR t</code>	no	<code>t=n.t.l.</code>	Join words, alternating italic and Roman.
<i>Request</i>	<i>Cause</i>	<i>If no</i>	<i>Explanation</i>																																										
	<i>Break</i>	<i>Argument</i>																																											
<code>.B t</code>	no	<code>t=n.t.l.*</code>	Text is in bold font.																																										
<code>.BI t</code>	no	<code>t=n.t.l.</code>	Join words, alternating bold and italic.																																										
<code>.BR t</code>	no	<code>t=n.t.l.</code>	Join words, alternating bold and Roman.																																										
<code>.DT</code>	no	<code>.5i 1i...</code>	Restore default tabs.																																										
<code>.HP i</code>	yes	<code>i=p.i.*</code>	Begin paragraph with hanging indent. Set prevailing indent to <i>i</i> .																																										
<code>.I t</code>	no	<code>t=n.t.l.</code>	Text is italic.																																										
<code>.IB t</code>	no	<code>t=n.t.l.</code>	Join words, alternating italic and bold.																																										
<code>.IP x i</code>	yes	<code>x=""</code>	Same as <code>.TP</code> with tag <i>x</i> .																																										
<code>.IR t</code>	no	<code>t=n.t.l.</code>	Join words, alternating italic and Roman.																																										

<i>Request</i>	<i>Cause</i>	<i>If no</i>	<i>Explanation</i>
	<i>Break</i>	<i>Argument</i>	
.IX <i>t</i>	no	-	Index macro, for SunSoft internal use.
.LP	yes	-	Begin left-aligned paragraph. Set prevailing indent to .5i.
.P	yes	-	Same as .LP.
.PD <i>d</i>	no	<i>d=.4v</i>	Set vertical distance between paragraphs.
.PP	yes	-	Same as .LP.
.RE	yes	-	End of relative indent. Restores prevailing indent.
.RB <i>t</i>	no	<i>t=n.t.l.</i>	Join words, alternating Roman and bold.
.RI <i>t</i>	no	<i>t=n.t.l.</i>	Join words, alternating Roman and italic.
.RS <i>i</i>	yes	<i>i=p.i.</i>	Start relative indent, increase indent by <i>i</i> . Sets prevailing indent to .5i for nested indents.
.SB <i>t</i>	no	-	Reduce size of text by 1 point, make text bold.
.SH <i>t</i>	yes	-	Section Heading.
.SM <i>t</i>	no	<i>t=n.t.l.</i>	Reduce size of text by 1 point.
.SS <i>t</i>	yes	<i>t=n.t.l.</i>	Section Subheading.
.TH <i>n s d f m</i>	yes	-	Begin reference page <i>n</i> , of of section <i>s</i> ; <i>d</i> is the date of the most recent change. If present, <i>f</i> is the left page footer; <i>m</i> is the main page (center) header. Sets prevailing indent and tabs to .5i.
.TP <i>i</i>	yes	<i>i=p.i.</i>	Begin indented paragraph, with the tag given on the next text line. Set prevailing indent to <i>i</i> .
.TX <i>t p</i>	no	-	Resolve the title abbreviation <i>t</i> ; join to punctuation mark (or text) <i>p</i> .

Conventions When formatting a manual page, mansun examines the first line to determine whether it requires special processing. For example a first line consisting of:

```
' \ " t
```

indicates that the manual page must be run through the `tbl(1)` preprocessor.

A typical manual page for a command or function is laid out as follows:

.TH <i>title</i> [1-8]	The name of the command or function, which serves as the title of the manual page. This is followed by the number of the section in which it appears.
.SH NAME	The name, or list of names, by which the command is called, followed by a dash and then a one-line summary of the action performed. All in Roman font, this section contains no <code>troff(1)</code> commands or escapes, and no macro requests. It is used to generate the <code>windex</code> database, which is used by the <code>what is(1)</code> command.
.SH SYNOPSIS	<p>Commands:</p> <p>The syntax of the command and its arguments, as typed on the command line. When in boldface, a word must be typed exactly as printed. When in italics, a word can be replaced with an argument that you supply. References to bold or italicized items are not capitalized in other sections, even when they begin a sentence.</p> <p>Syntactic symbols appear in Roman face:</p> <p>[] An argument, when surrounded by brackets is optional.</p> <p> Arguments separated by a vertical bar are exclusive. You can supply only one item from such a list.</p> <p>... Arguments followed by an ellipsis can be repeated. When an ellipsis follows a bracketed set, the expression within the brackets can be repeated.</p> <p>Functions:</p> <p>If required, the data declaration, or <code>#include</code> directive, is shown first, followed by the function declaration. Otherwise, the function declaration is shown.</p>
.SH DESCRIPTION	A narrative overview of the command or function's external behavior. This includes how it interacts with files or data, and how it handles the standard input, standard output and standard error. Internals and implementation details are normally omitted. This section attempts to provide a succinct overview in answer to the question, "what does it do?"

Literal text from the synopsis appears in constant width, as do literal filenames and references to items that appear elsewhere in the reference manuals. Arguments are italicized.

If a command interprets either subcommands or an input grammar, its command interface or input grammar is normally described in a USAGE section, which follows the OPTIONS section. The DESCRIPTION section only describes the behavior of the command itself, not that of subcommands.

- .SH OPTIONS The list of options along with a description of how each affects the command's operation.
- .SH FILES A list of files associated with the command or function.
- .SH SEE ALSO A comma-separated list of related manual pages, followed by references to other published materials.
- .SH DIAGNOSTICS A list of diagnostic messages and an explanation of each.
- .SH BUGS A description of limitations, known defects, and possible problems associated with the command or function.

FILES /usr/share/lib/tmac/ansun
 /usr/share/man/windex

SEE ALSO man(1), nroff(1), troff(1), whatis(1)

Dale Dougherty and Tim O'Reilly, *Unix Text Processing*

me(5)

NAME	me – macros for formatting papers																																				
SYNOPSIS	nroff -me [<i>options</i>] <i>filename...</i> troff -me [<i>options</i>] <i>filename...</i>																																				
DESCRIPTION	<p>This package of nroff and troff macro definitions provides a canned formatting facility for technical papers in various formats. When producing 2-column output on a terminal, filter the output through col(1).</p> <p>The macro requests are defined below. Many nroff and troff requests are unsafe in conjunction with this package, however, these requests may be used with impunity after the first .pp:</p> <p>.bp begin new page .br break output line here .sp <i>n</i> insert <i>n</i> spacing lines .ls <i>n</i> (line spacing) <i>n</i>=1 single, <i>n</i>=2 double space .na no alignment of right margin .ce <i>n</i> center next <i>n</i> lines .ul <i>n</i> underline next <i>n</i> lines .sz +<i>n</i> add <i>n</i> to point size</p> <p>Output of the eqn(1), neqn(1), refer(1), and tbl(1) preprocessors for equations and tables is acceptable as input.</p>																																				
REQUESTS	<p>In the following list, “initialization” refers to the first .pp, .lp, .ip, .np, .sh, or .uh macro. This list is incomplete.</p> <table><thead><tr><th><i>Request</i></th><th><i>Initial Value</i></th><th><i>Cause Break</i></th><th><i>Explanation</i></th></tr></thead><tbody><tr><td>. (c</td><td>-</td><td>yes</td><td>Begin centered block.</td></tr><tr><td>. (d</td><td>-</td><td>no</td><td>Begin delayed text.</td></tr><tr><td>. (f</td><td>-</td><td>no</td><td>Begin footnote.</td></tr><tr><td>. (l</td><td>-</td><td>yes</td><td>Begin list.</td></tr><tr><td>. (q</td><td>-</td><td>yes</td><td>Begin major quote.</td></tr><tr><td>. (xx</td><td>-</td><td>no</td><td>Begin indexed item in index <i>x</i>.</td></tr><tr><td>. (z</td><td>-</td><td>no</td><td>Begin floating keep.</td></tr><tr><td>.) c</td><td>-</td><td>yes</td><td>End centered block.</td></tr></tbody></table>	<i>Request</i>	<i>Initial Value</i>	<i>Cause Break</i>	<i>Explanation</i>	. (c	-	yes	Begin centered block.	. (d	-	no	Begin delayed text.	. (f	-	no	Begin footnote.	. (l	-	yes	Begin list.	. (q	-	yes	Begin major quote.	. (xx	-	no	Begin indexed item in index <i>x</i> .	. (z	-	no	Begin floating keep.	.) c	-	yes	End centered block.
<i>Request</i>	<i>Initial Value</i>	<i>Cause Break</i>	<i>Explanation</i>																																		
. (c	-	yes	Begin centered block.																																		
. (d	-	no	Begin delayed text.																																		
. (f	-	no	Begin footnote.																																		
. (l	-	yes	Begin list.																																		
. (q	-	yes	Begin major quote.																																		
. (xx	-	no	Begin indexed item in index <i>x</i> .																																		
. (z	-	no	Begin floating keep.																																		
.) c	-	yes	End centered block.																																		

<i>Request</i>	<i>Initial Value</i>	<i>Cause Break</i>	<i>Explanation</i>
.) d	-	yes	End delayed text.
.) f	-	yes	End footnote.
.) l	-	yes	End list.
.) q	-	yes	End major quote.
.) x	-	yes	End index item.
.) z	-	yes	End floating keep.
.++ <i>m H</i>	-	no	Define paper section. <i>m</i> defines the part of the paper, and can be C (chapter), A (appendix), P (preliminary, for instance, abstract, table of contents, etc.), B (bibliography), RC (chapters renumbered from page one each chapter), or RA (appendix renumbered from page one).
.+c <i>T</i>	-	yes	Begin chapter (or appendix, etc., as set by .++). <i>T</i> is the chapter title.
.1c	1	yes	One column format on a new page.
.2c	1	yes	Two column format.
.EN	-	yes	Space after equation produced by eqn or neqn.
.EQ <i>x y</i>	-	yes	Precede equation; break out and add space. Equation number is <i>y</i> . The optional argument <i>x</i> may be <i>I</i> to indent equation (default), <i>L</i> to left-adjust the equation, or <i>C</i> to center the equation.
.GE	-	yes	End <i>gremlin</i> picture.

me(5)

<i>Request</i>	<i>Initial Value</i>	<i>Cause Break</i>	<i>Explanation</i>
.GS	-	yes	Begin <i>gremlin</i> picture.
.PE	-	yes	End <i>pic</i> picture.
.PS	-	yes	Begin <i>pic</i> picture.
.TE	-	yes	End table.
.TH	-	yes	End heading section of table.
.TS <i>x</i>	-	yes	Begin table; if <i>x</i> is <i>H</i> table has repeated heading.
.ac <i>A N</i>	-	no	Set up for ACM style output. <i>A</i> is the Author's name(s), <i>N</i> is the total number of pages. Must be given before the first initialization.
.b <i>x</i>	no	no	Print <i>x</i> in boldface; if no argument switch to boldface.
.ba <i>+n</i>	0	yes	Augments the base indent by <i>n</i> . This indent is used to set the indent on regular text (like paragraphs).
.bc	no	yes	Begin new column.
.bi <i>x</i>	no	no	Print <i>x</i> in bold italics (nofill only).
.bu	-	yes	Begin bulleted paragraph.
.bx <i>x</i>	no	no	Print <i>x</i> in a box (nofill only).
.ef ' <i>x'y'z</i>	''''	no	Set even footer to <i>x y z</i> .
.eh ' <i>x'y'z</i>	''''	no	Set even header to <i>x y z</i> .
.fo ' <i>x'y'z</i>	''''	no	Set footer to <i>x y z</i> .
.hx	-	no	Suppress headers and footers on next page.
.he ' <i>x'y'z</i>	''''	no	Set header to <i>x y z</i> .
.hl	-	yes	Draw a horizontal line.

<i>Request</i>	<i>Initial Value</i>	<i>Cause Break</i>	<i>Explanation</i>
.i <i>x</i>	no	no	Italicize <i>x</i> ; if <i>x</i> missing, italic text follows.
.ip <i>x y</i>	no	yes	Start indented paragraph, with hanging tag <i>x</i> . Indentation is <i>y</i> ens (default 5).
.lp	yes	yes	Start left-blocked paragraph.
.lo	-	no	Read in a file of local macros of the form . * <i>x</i> . Must be given before initialization.
.np	1	yes	Start numbered paragraph.
.of ' <i>x'y'z</i>	''''	no	Set odd footer to <i>x y z</i> .
.oh ' <i>x'y'z</i>	''''	no	Set odd header to <i>x y z</i> .
.pd	-	yes	Print delayed text.
.pp	no	yes	Begin paragraph. First line indented.
.r	yes	no	Roman text follows.
.re	-	no	Reset tabs to default values.
.sc	no	no	Read in a file of special characters and diacritical marks. Must be given before initialization.
.sh <i>n x</i>	-	yes	Section head follows, font automatically bold. <i>n</i> is level of section, <i>x</i> is title of section.
.sk	no	no	Leave the next page blank. Only one page is remembered ahead.
.sm <i>x</i>	-	no	Set <i>x</i> in a smaller pointsize.
.sz <i>+n</i>	10p	no	Augment the point size by <i>n</i> points.
.th	no	no	Produce the paper in thesis format. Must be given before initialization.

me(5)

<i>Request</i>	<i>Initial</i>	<i>Cause</i>	<i>Explanation</i>
	<i>Value</i>	<i>Break</i>	
.tp	no	yes	Begin title page.
.u x	-	no	Underline argument (even in troff). (Nofill only).
.uh	-	yes	Like .sh but unnumbered.
.xp x	-	no	Print index x.

FILES /usr/share/lib/tmac/e

/usr/share/lib/tmac/*.me

SEE ALSO col(1), eqn(1), nroff(1), refer(1), tbl(1), troff(1)

NAME	mm – text formatting (memorandum) macros																				
SYNOPSIS	nroff -mm [<i>options</i>] <i>filename...</i> troff -mm [<i>options</i>] <i>filename...</i>																				
DESCRIPTION	<p>This package of <code>nroff(1)</code> and <code>troff(1)</code> macro definitions provides a formatting facility for various styles of articles, theses, and books. When producing 2-column output on a terminal or lineprinter, or when reverse line motions are needed, filter the output through <code>col(1)</code>. All external -mm macros are defined below.</p> <p>Note: this -mm macro package is an extended version written at Berkeley and is a superset of the standard -mm macro packages as supplied by Bell Labs. Some of the Bell Labs macros have been removed; for instance, it is assumed that the user has little interest in producing headers stating that the memo was generated at Whippany Labs.</p> <p>Many <code>nroff</code> and <code>troff</code> requests are unsafe in conjunction with this package. However, the first four requests below may be used with impunity after initialization, and the last two may be used even before initialization:</p> <p><code>.bp</code> begin new page <code>.br</code> break output line <code>.spn</code> insert n spacing lines <code>.cen</code> center next n lines <code>.lsn</code> line spacing: $n=1$ single, $n=2$ double space <code>.na</code> no alignment of right margin</p> <p>Font and point size changes with <code>\f</code> and <code>\s</code> are also allowed; for example, <code>\fIword\fR</code> will italicize <i>word</i>. Output of the <code>tbl(1)</code>, <code>eqn(1)</code> and <code>refer(1)</code> preprocessors for equations, tables, and references is acceptable as input.</p>																				
REQUESTS	<p>Here is a table of macros.</p> <table border="1"> <thead> <tr> <th>Macro Name</th> <th>Initial Value</th> <th>Break? Reset?</th> <th>Explanation</th> </tr> </thead> <tbody> <tr> <td><code>.1C</code></td> <td>on</td> <td>y,y</td> <td>one column format on a new page</td> </tr> <tr> <td><code>.2C [l]</code></td> <td>–</td> <td>y,y</td> <td>two column format l=line length</td> </tr> <tr> <td><code>.AE</code></td> <td>–</td> <td>y</td> <td>end abstract</td> </tr> <tr> <td><code>.AL [t][i][s]</code></td> <td>$t=1; i=.Li; s=0$</td> <td>y</td> <td>Start automatic list type $t=[1,A,a,I,i]$ 1=arabic numbers; A=uppercase letters a=lowercase letters; I=uppercase Roman numerals; i=lowercase Roman numerals indentation i; separation s</td> </tr> </tbody> </table>	Macro Name	Initial Value	Break? Reset?	Explanation	<code>.1C</code>	on	y,y	one column format on a new page	<code>.2C [l]</code>	–	y,y	two column format l =line length	<code>.AE</code>	–	y	end abstract	<code>.AL [t][i][s]</code>	$t=1; i=.Li; s=0$	y	Start automatic list type $t=[1,A,a,I,i]$ 1=arabic numbers; A=uppercase letters a=lowercase letters; I=uppercase Roman numerals; i=lowercase Roman numerals indentation i ; separation s
Macro Name	Initial Value	Break? Reset?	Explanation																		
<code>.1C</code>	on	y,y	one column format on a new page																		
<code>.2C [l]</code>	–	y,y	two column format l =line length																		
<code>.AE</code>	–	y	end abstract																		
<code>.AL [t][i][s]</code>	$t=1; i=.Li; s=0$	y	Start automatic list type $t=[1,A,a,I,i]$ 1=arabic numbers; A=uppercase letters a=lowercase letters; I=uppercase Roman numerals; i=lowercase Roman numerals indentation i ; separation s																		

mm(5)

Macro Name	Initial Value	Break? Reset?	Explanation
.AS <i>m</i> [<i>n</i>]	<i>n</i> =0	y	begin abstract
.AU	–	y	author's name
.AV <i>x</i>	–	y	signature and date line of verifier <i>x</i>
.B <i>x</i>	–	n	embolden <i>x</i> ; if no <i>x</i> , switch to boldface
.BE	–	y	end block text
.BI <i>x y</i>	–	n	embolden <i>x</i> and underline <i>y</i>
.BL	–	y	bullet list
.BR <i>x y</i>	–	n	embolden <i>x</i> and use Roman font for <i>y</i>
.BS	–	n	start block text
.CN	–	y	same as .DE (nroff)
.CS	–	y	cover sheet
.CW	–	n	same as .DS I (nroff)
.DE	–	y	end display
.DF [<i>p</i>] [<i>f</i>] [<i>rp</i>]	<i>p</i> =L; <i>f</i> =N	y	start floating display; position <i>p</i> =[L,C,CB] L=left; I=indent; C=center; CB=center block fill <i>f</i> =[N,Y]; right position <i>rp</i> (fill only)
.DL [<i>i</i>] [<i>s</i>]	–	y	start dash list
.DS [<i>p</i>] [<i>f</i>] [<i>rp</i>]	<i>p</i> =L; <i>f</i> =N	y	begin static display (see .DF for argument descriptions)
.EC <i>x</i> [<i>n</i>]	<i>n</i> =1	y	equation title; equation <i>x</i> ; number <i>n</i>
.EF <i>x</i>	–	n	even footer appears at the bottom of even-numbered pages; <i>x</i> ="l' c' r" l=left; c=center; r=right
.EH <i>x</i>	–	n	even header appears at the top of even-numbered pages; <i>x</i> ="l' c' r" l=left; c=center; r=right
.EN	–	y	end displayed equation produced by eqn
.EQ	–	y	break out equation produced by eqn
.EX <i>x</i> [<i>n</i>]	<i>n</i> =1	y	exhibit title; exhibit <i>x</i>

Macro Name	Initial Value	Break? Reset?	Explanation
			number <i>n</i>
.FD [<i>f</i>] [<i>r</i>]	<i>f</i> =10; <i>r</i> =1	n	set footnote style format <i>f</i> =[0-11]; renumber <i>r</i> =[0,1]
.FE	–	y	end footnote
.FG <i>x</i> [<i>n</i>]	<i>n</i> =1	y	figure title; figure <i>x</i> ; number <i>n</i>
.FS	–	n	start footnote
.H <i>l</i> [<i>t</i>]	–	y	produce numbered heading level <i>l</i> =[1-7]; title <i>t</i>
.HU <i>t</i>	–	y	produce unnumbered heading; title <i>t</i>
.I <i>x</i>	–	n	underline <i>x</i>
.IB <i>x y</i>	–	n	underline <i>x</i> and embolden <i>y</i>
.IR <i>x y</i>	–	n	underline <i>x</i> and use Roman font on <i>y</i>
.LE [<i>s</i>]	<i>s</i> =0	y	end list; separation <i>s</i>
.LI [<i>m</i>] [<i>p</i>]	–	y	start new list item; mark <i>m</i> prefix <i>p</i> (mark only)
.ML <i>m</i> [<i>i</i>] [<i>s</i>]	<i>s</i> =0	y	start marked list; mark <i>m</i> indentation <i>i</i> ; separation <i>s</i> =[0,1]
.MT <i>x</i>		y	memo title; title <i>x</i>
.ND <i>x</i>		n	no date in page footer; <i>x</i> is date on cover
.NE	–	y	end block text
.NS	–	y	start block text
.OF <i>x</i>	–	n	odd footer appears at the bottom of odd-numbered pages; <i>x</i> ="l' c' r" <i>l</i> =left; <i>c</i> =center; <i>r</i> =right
.OF <i>x</i>	–	n	odd header appears at the top of odd-numbered pages; <i>x</i> ="l' c' r" <i>l</i> =left; <i>c</i> =center; <i>r</i> =right
.OP	–	y	skip to the top of an odd-number page
.P [<i>t</i>]	<i>t</i> =0	y,y	begin paragraph; <i>t</i> =[0,1] 0=justified; 1=indented

mm(5)

Macro Name	Initial Value	Break? Reset?	Explanation
.PF <i>x</i>	–	n	page footer appears at the bottom of every page; <i>x</i> ="l' c' r" l=left; c=center; r=right
.PH <i>x</i>	–	n	page header appears at the top of every page; <i>x</i> ="l' c' r" l=left; c=center; r=right
.R	on	n	return to Roman font
.RB <i>x y</i>	–	n	use Roman on <i>x</i> and embolden <i>y</i>
.RI <i>x y</i>	–	n	use Roman on <i>x</i> and underline <i>y</i>
.RP <i>x</i>	–	y,y	released paper format ? <i>x</i> =no stops title on first
.RS	5n	y,y	right shift: start level of relative indentation
.S <i>m n</i>	–	n	set character point size & vertical space character point size <i>m</i> ; vertical space <i>n</i>
.SA <i>x</i>	<i>x</i> =1	n	justification; <i>x</i> =[0,1]
.SK <i>x</i>	–	y	skip <i>x</i> pages
.SM	–	n	smaller; decrease point size by 2
.SP [<i>x</i>]	–	y	leave <i>x</i> blank lines
.TB <i>x</i> [<i>n</i>]	<i>n</i> =1	y	table title; table <i>x</i> ; number <i>n</i>
.TC	–	y	print table of contents (put at end of input file)
.TE	–	y	end of table processed by tbl
.TH	–	y	end multi-page header of table
.TL	–	n	title in boldface and two points larger
.TM	–	n	UC Berkeley thesis mode
.TP <i>i</i>	y	y	<i>i</i> =p.i. Begin indented paragraph, with the tag given on the next text line. Set prevailing indent to <i>i</i> .
.TS <i>x</i>	–	y,y	begin table; if <i>x</i> =H table has multi-page header
.TY	–	y	display centered title CONTENTS

Macro Name	Initial Value	Break? Reset?	Explanation
<code>.VL i [m] [s]</code>	$m=0;s=0$	y	start variable-item list; indentation i mark-indentation m ; separation s

REGISTERS

Formatting distances can be controlled in -mm by means of built-in number registers. For example, this sets the line length to 6.5 inches:

```
.nr LL 6.5i
```

Here is a table of number registers and their default values:

Name	Register Controls	Takes Effect	Default
Cl	contents level	table of contents	2
De	display eject	display	0
Df	display floating	display	5
Ds	display spacing	display	1v
Hb	heading break	heading	2
Hc	heading centering	heading	0
Hi	heading indent	heading	1
Hi	heading spacing	heading	1
Hu	heading unnumbered	heading	2
Li	list indentation	list	6 (nroff) 5 (troff)
Ls	list spacing	list	6
Pi	paragraph indent	paragraph	5
Pt	paragraph type	paragraph	1
Si	static indent	display	5 (nroff) 3 (troff)

When resetting these values, make sure to specify the appropriate units. Setting the line length to 7, for example, will result in output with one character per line. Setting `Pi` to 0 suppresses paragraph indentation

Here is a list of string registers available in -mm; they may be used anywhere in the text:

mm(5)

Name	String's Function
*Q	quote (" in nroff, `` in troff)
*U	unquote (" in nroff, ' ' in troff)
*-	dash (-- in nroff, — in troff)
*(MO	month (month of the year)
*(DY	day (current date)
**	automatically numbered footnote
*'	acute accent (before letter)
*`	grave accent (before letter)
*^	circumflex (before letter)
*,	cedilla (before letter)
*:	umlaut (before letter)
*~	tilde (before letter)
\(BU	bullet item
\(DT	date (<i>month day, yr</i>)
\(EM	em dash
\(Lf	LIST OF FIGURES title
\(Lt	LIST OF TABLES title
\(Lx	LIST OF EXHIBITS title
\(Le	LIST OF EQUATIONS title
\(Rp	REFERENCES title
\(Tm	trademark character (TM)

When using the extended accent mark definitions available with .AM, these strings should come after, rather than before, the letter to be accented.

FILES

/usr/share/lib/tmac/m

/usr/share/lib/tmac/mm.[nt] nroff and troff definitions of mm.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWdoc

SEE ALSO `col(1)`, `eqn(1)`, `nroff(1)`, `refer(1)`, `tbl(1)`, `troff(1)`, `attributes(5)`

BUGS Floating keeps and regular keeps are diverted to the same space, so they cannot be mixed together with predictable results.

ms(5)

NAME	ms – text formatting macros																																				
SYNOPSIS	nroff -ms [<i>options</i>] <i>filename...</i> troff -ms [<i>options</i>] <i>filename...</i>																																				
DESCRIPTION	<p>This package of <code>nroff(1)</code> and <code>troff(1)</code> macro definitions provides a formatting facility for various styles of articles, theses, and books. When producing 2-column output on a terminal or lineprinter, or when reverse line motions are needed, filter the output through <code>col(1)</code>. All external <code>-ms</code> macros are defined below.</p> <p>Note: this <code>-ms</code> macro package is an extended version written at Berkeley and is a superset of the standard <code>-ms</code> macro packages as supplied by Bell Labs. Some of the Bell Labs macros have been removed; for instance, it is assumed that the user has little interest in producing headers stating that the memo was generated at Whippy Labs.</p> <p>Many <code>nroff</code> and <code>troff</code> requests are unsafe in conjunction with this package. However, the first four requests below may be used with impunity after initialization, and the last two may be used even before initialization:</p> <p><code>.bp</code> begin new page <code>.br</code> break output line <code>.sp <i>n</i></code> insert <i>n</i> spacing lines <code>.ce <i>n</i></code> center next <i>n</i> lines <code>.ls <i>n</i></code> line spacing: <i>n</i>=1 single, <i>n</i>=2 double space <code>.na</code> no alignment of right margin</p> <p>Font and point size changes with <code>\f</code> and <code>\s</code> are also allowed; for example, <code>\fIword\fR</code> will italicize <i>word</i>. Output of the <code>tbl(1)</code>, <code>eqn(1)</code> and <code>refer(1)</code> preprocessors for equations, tables, and references is acceptable as input.</p>																																				
REQUESTS	<table border="1"><thead><tr><th>Macro Name</th><th>Initial Value</th><th>Break? Reset?</th><th>Explanation</th></tr></thead><tbody><tr><td><code>.AB <i>x</i></code></td><td>–</td><td>y</td><td>begin abstract; if <i>x</i>=no do not label abstract</td></tr><tr><td><code>.AE</code></td><td>–</td><td>y</td><td>end abstract</td></tr><tr><td><code>.AI</code></td><td>–</td><td>y</td><td>author's institution</td></tr><tr><td><code>.AM</code></td><td>–</td><td>n</td><td>better accent mark definitions</td></tr><tr><td><code>.AU</code></td><td>–</td><td>y</td><td>author's name</td></tr><tr><td><code>.B <i>x</i></code></td><td>–</td><td>n</td><td>embolden <i>x</i>; if no <i>x</i>, switch to boldface</td></tr><tr><td><code>.B1</code></td><td>–</td><td>y</td><td>begin text to be enclosed in a box</td></tr><tr><td><code>.B2</code></td><td>–</td><td>y</td><td>end boxed text and print it</td></tr></tbody></table>	Macro Name	Initial Value	Break? Reset?	Explanation	<code>.AB <i>x</i></code>	–	y	begin abstract; if <i>x</i> =no do not label abstract	<code>.AE</code>	–	y	end abstract	<code>.AI</code>	–	y	author's institution	<code>.AM</code>	–	n	better accent mark definitions	<code>.AU</code>	–	y	author's name	<code>.B <i>x</i></code>	–	n	embolden <i>x</i> ; if no <i>x</i> , switch to boldface	<code>.B1</code>	–	y	begin text to be enclosed in a box	<code>.B2</code>	–	y	end boxed text and print it
Macro Name	Initial Value	Break? Reset?	Explanation																																		
<code>.AB <i>x</i></code>	–	y	begin abstract; if <i>x</i> =no do not label abstract																																		
<code>.AE</code>	–	y	end abstract																																		
<code>.AI</code>	–	y	author's institution																																		
<code>.AM</code>	–	n	better accent mark definitions																																		
<code>.AU</code>	–	y	author's name																																		
<code>.B <i>x</i></code>	–	n	embolden <i>x</i> ; if no <i>x</i> , switch to boldface																																		
<code>.B1</code>	–	y	begin text to be enclosed in a box																																		
<code>.B2</code>	–	y	end boxed text and print it																																		

Macro Name	Initial Value	Break? Reset?	Explanation
.BT	date	n	bottom title, printed at foot of page
.BX <i>x</i>	–	n	print word <i>x</i> in a box
.CM	if t	n	cut mark between pages
.CT	–	<i>y,y</i>	chapter title: page number moved to CF (TM only)
.DA <i>x</i>	if n	n	force date <i>x</i> at bottom of page; today if no <i>x</i>
.DE	–	y	end display (unfilled text) of any kind
.DS <i>x y</i>	l	y	begin display with keep; <i>x</i> =I, L, C, B; <i>y</i> =indent
.ID <i>y</i>	8n,.5i	y	indented display with no keep; <i>y</i> =indent
.LD	–	y	left display with no keep
.CD	–	y	centered display with no keep
.BD	–	y	block display; center entire block
.EF <i>x</i>	–	n	even page footer <i>x</i> (3 part as for .t1)
.EH <i>x</i>	–	n	even page header <i>x</i> (3 part as for .t1)
.EN	–	y	end displayed equation produced by eqn
.EQ <i>x y</i>	–	y	break out equation; <i>x</i> =L,I,C; <i>y</i> =equation number
.FE	–	n	end footnote to be placed at bottom of page
.FP	–	n	numbered footnote paragraph; may be redefined
.FS <i>x</i>	–	n	start footnote; <i>x</i> is optional footnote label
.HD	undef	n	optional page header below header margin
.I <i>x</i>	–	n	italicize <i>x</i> ; if no <i>x</i> , switch to italics
.IP <i>x y</i>	–	<i>y,y</i>	indented paragraph, with hanging tag <i>x</i> ; <i>y</i> =indent
.IX <i>x y</i>	–	y	index words <i>x y</i> and so on (up to 5 levels)
.KE	–	n	end keep of any kind
.KF	–	n	begin floating keep; text fills remainder of page
.KS	–	y	begin keep; unit kept together on a single page
.LG	–	n	larger; increase point size by 2
.LP	–	<i>y,y</i>	left (block) paragraph.
.MC <i>x</i>	–	<i>y,y</i>	multiple columns; <i>x</i> =column width
.ND <i>x</i>	if t	n	no date in page footer; <i>x</i> is date on cover

ms(5)

Macro Name	Initial Value	Break? Reset?	Explanation
.NH <i>x y</i>	–	<i>y,y</i>	numbered header; <i>x</i> =level, <i>x</i> =0 resets, <i>x</i> =S sets to <i>y</i>
.NL	10p	n	set point size back to normal
.OF <i>x</i>	–	n	odd page footer <i>x</i> (3 part as for .t1)
.OH <i>x</i>	–	n	odd page header <i>x</i> (3 part as for .t1)
.P1	if TM	n	print header on first page
.PP	–	<i>y,y</i>	paragraph with first line indented
.PT	- % -	n	page title, printed at head of page
.PX <i>x</i>	–	<i>y</i>	print index (table of contents); <i>x</i> =no suppresses title
.QP	–	<i>y,y</i>	quote paragraph (indented and shorter)
.R	on	n	return to Roman font
.RE	5n	<i>y,y</i>	retreat: end level of relative indentation
.RP <i>x</i>	–	n	released paper format; <i>x</i> =no stops title on first page
.RS	5n	<i>y,y</i>	right shift: start level of relative indentation
.SH	–	<i>y,y</i>	section header, in boldface
.SM	–	n	smaller; decrease point size by 2
.TA	8n,5n	n	set TAB characters to 8n 16n ... (nroff) or 5n 10n ... (troff)
.TC <i>x</i>	–	<i>y</i>	print table of contents at end; <i>x</i> =no suppresses title
.TE	–	<i>y</i>	end of table processed by t1
.TH	–	<i>y</i>	end multi-page header of table
.TL	–	<i>y</i>	title in boldface and two points larger
.TM	off	n	UC Berkeley thesis mode
.TS <i>x</i>	–	<i>y,y</i>	begin table; if <i>x</i> =H table has multi-page header
.UL <i>x</i>	–	n	underline <i>x</i> , even in troff
.UX <i>x</i>	–	n	UNIX; trademark message first time; <i>x</i> appended
.XA <i>x y</i>	–	<i>y</i>	another index entry; <i>x</i> =page or no for none; <i>y</i> =indent
.XE	–	<i>y</i>	end index entry (or series of .IX entries)
.XP	–	<i>y,y</i>	paragraph with first line indented, others indented
.XS <i>x y</i>	–	<i>y</i>	begin index entry; <i>x</i> =page or no for none; <i>y</i> =indent

Macro Name	Initial Value	Break? Reset?	Explanation
.1C	on	y,y	one column format, on a new page
.2C	-	y,y	begin two column format
.] -	-	n	beginning of <i>refer</i> reference
. [0	-	n	end of unclassifiable type of reference
. [N	-	n	N= 1:journal-article, 2:book, 3:book-article, 4:report

REGISTERS

Formatting distances can be controlled in -ms by means of built-in number registers. For example, this sets the line length to 6.5 inches:

```
.nr LL 6.5i
```

Here is a table of number registers and their default values:

Name	Register Controls	Takes Effect	Default
PS	point size	paragraph	10
VS	vertical spacing	paragraph	12
LL	line length	paragraph	6i
LT	title length	next page	same as LL
FL	footnote length	next .FS	5.5i
PD	paragraph distance	paragraph	1v (if n), .3v (if t)
DD	display distance	displays	1v (if n), .5v (if t)
PI	paragraph indent	paragraph	5n
QI	quote indent	next .QP	5n
FI	footnote indent	next .FS	2n
PO	page offset	next page	0 (if n), ≈1i (if t)
HM	header margin	next page	1i
FM	footer margin	next page	1i
FF	footnote format	next .FS	0 (1, 2, 3 available)

When resetting these values, make sure to specify the appropriate units. Setting the line length to 7, for example, will result in output with one character per line. Setting FF to 1 suppresses footnote superscripting; setting it to 2 also suppresses indentation of the first line; and setting it to 3 produces an .IP-like footnote paragraph.

ms(5)

Here is a list of string registers available in `-ms`; they may be used anywhere in the text:

Name	String's Function
<code>*Q</code>	quote (" in <code>nroff</code> , " in <code>troff</code>)
<code>*U</code>	unquote (" in <code>nroff</code> , " in <code>troff</code>)
<code>*-</code>	dash (-- in <code>nroff</code> , - in <code>troff</code>)
<code>*(MO</code>	month (month of the year)
<code>*(DY</code>	day (current date)
<code>**</code>	automatically numbered footnote
<code>*' </code>	acute accent (before letter)
<code>*` </code>	grave accent (before letter)
<code>*^ </code>	circumflex (before letter)
<code>*, </code>	cedilla (before letter)
<code>*: </code>	umlaut (before letter)
<code>*~ </code>	tilde (before letter)

When using the extended accent mark definitions available with `.AM`, these strings should come after, rather than before, the letter to be accented.

FILES `/usr/share/lib/tmac/s`

`/usr/share/lib/tmac/ms.???`

SEE ALSO `col(1)`, `eqn(1)`, `nroff(1)`, `refer(1)`, `tbl(1)`, `troff(1)`

BUGS Floating keeps and regular keeps are diverted to the same space, so they cannot be mixed together with predictable results.

NAME	nfssec – overview of NFS security modes
DESCRIPTION	<p>The <code>mount_nfs(1M)</code> and <code>share_nfs(1M)</code> commands each provide a way to specify the security mode to be used on an NFS file system through the <code>sec=mode</code> option. <code>mode</code> can be either <code>sys</code>, <code>dh</code>, <code>krb5</code>, <code>krb5i</code>, <code>krb5p</code>, or <code>none</code>. These security modes may also be added to the automount maps. Note that <code>mount_nfs(1M)</code> and <code>automount(1M)</code> do not support <code>sec=none</code> at this time.</p> <p>The <code>sec=mode</code> option on the <code>share_nfs(1M)</code> command line establishes the security mode of NFS servers. If the NFS connection uses the NFS Version 3 protocol, the NFS clients must query the server for the appropriate <code>mode</code> to use. If the NFS connection uses the NFS Version 2 protocol, then the NFS client will use the default security mode, which is currently <code>sys</code>. NFS clients may force the use of a specific security mode by specifying the <code>sec=mode</code> option on the command line. However, if the file system on the server is not shared with that security mode, the client may be denied access.</p> <p>If the NFS client wants to authenticate the NFS server using a particular (stronger) security mode, the client will want to specify the security mode to be used, even if the connection uses the NFS Version 3 protocol. This guarantees that an attacker masquerading as the server does not compromise the client.</p> <p>The NFS security modes are described below. Of these, the <code>krb5</code>, <code>krb5i</code>, <code>krb5p</code> modes use the Kerberos V5 protocol for authenticating and protecting the shared filesystems. Before these can be used, the system must be configured to be part of a Kerberos realm (see <code>SEAM(5)</code>).</p> <p><code>sys</code> Use <code>AUTH_SYS</code> authentication. The user's UNIX user-id and group-ids are passed in the clear on the network, unauthenticated by the NFS server. This is the simplest security method and requires no additional administration. It is the default used by Solaris NFS Version 2 clients and Solaris NFS servers.</p> <p><code>dh</code> Use a Diffie-Hellman public key system (<code>AUTH_DES</code>, which is referred to as <code>AUTH_DH</code> in the forthcoming Internet RFC).</p> <p><code>krb5</code> Use Kerberos V5 protocol to authenticate users before granting access to the shared filesystem.</p> <p><code>krb5i</code> Use Kerberos V5 authentication with integrity checking (checksums) to verify that the data has not been tampered with.</p> <p><code>krb5p</code> User Kerberos V5 authentication, integrity checksums, and privacy protection (encryption) on the shared filesystem. This provides the most secure filesystem sharing, as all traffic is encrypted. It should be noted that performance might suffer on some systems when using <code>krb5p</code>, depending on the computational intensity of the encryption algorithm and the amount of data being transferred.</p> <p><code>none</code> Use null authentication (<code>AUTH_NONE</code>). NFS clients using <code>AUTH_NONE</code> have no identity and are mapped to the anonymous user <code>nobody</code> by NFS</p>

nfssec(5)

servers. A client using a security mode other than the one with which a Solaris NFS server shares the file system will have its security mode mapped to AUTH_NONE. In this case, if the file system is shared with *sec=none*, users from the client will be mapped to the anonymous user. The NFS security mode none is supported by *share_nfs(1M)*, but not by *mount_nfs(1M)* or *automount(1M)*.

FILES /etc/nfssec.conf NFS security service configuration file.

ATTRIBUTES See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWnfscr

SEE ALSO *automount(1M)*, *mount_nfs(1M)*, *share_nfs(1M)*, *rpc_clnt_auth(3NSL)*, *secure_rpc(3NSL)*, *attributes(5)*

NOTES /etc/nfssec.conf lists the NFS security services. Do not edit this file. It is not intended to be user-configurable.

NAME	pam_authtok_check – authentication and password management module						
SYNOPSIS	pam_authtok_check.so.1						
DESCRIPTION	<p>pam_authtok_check provides functionality to the Password Management stack. The implementation of pam_sm_chauthtok(), performs a number of checks on the construction of the newly entered password. pam_sm_chauthtok() is invoked twice by the PAM framework, once with flags set to PAM_PRELIM_CHECK, and once with flags set to PAM_UPDATE_AUTH Tok. This module only performs its checks during the first invocation. This module expects the current authentication token in the PAM_OLDAUTH Tok item, the new (to be checked) password in the PAM_AUTH Tok item, and the login name in the PAM_USER item. The checks performed by this module are:</p> <p>length The password length should not be less than the minimum specified in /etc/default/passwd.</p> <p>circular shift The password should not be a circular shift of the login name.</p> <p>complexity The password should contain at least two alpha characters and one numeric or special character.</p> <p>variation The old and new passwords must differ by at least three positions.</p> <p>The following option may be passed to the module:</p> <p>debug syslog(3C) debugging information at the LOG_DEBUG level</p>						
ERRORS	If the password in PAM_AUTH Tok passes all tests, PAM_SUCCESS is returned. If any of the tests fail, PAM_AUTH Tok_ERR is returned.						
FILES	/etc/default/passwd Contains the value for PASSELENGTH, the default minimal password length.						
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> <tr> <td>MT Level</td> <td>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving	MT Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
MT Level	MT-Safe with exceptions						
SEE ALSO	passwd(1), pam(3PAM), pam_chauthtok(3PAM), syslog(3C), libpam(3LIB), pam.conf(4), attributes(5), pam_authtok_get(5), pam_authtok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix(5), pam_unix_account(5), pam_unix_auth(5), pam_unix_session(5)						
NOTES	The interfaces in libpam(3LIB) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.						

pam_authtok_check(5)

The pam_unix(5) module might not be supported in a future release. Similar functionality is provided by pam_authtok_check(5), pam_authtok_get(5), pam_authtok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix_account(5), pam_unix_auth(5), and pam_unix_session(5).

NAME	pam_authtok_get – authentication and password management module
SYNOPSIS	pam_authtok_get.so.1
DESCRIPTION	The pam_authtok_get service module provides password prompting functionality to the PAM stack. It implements pam_sm_authenticate() and pam_sm_chauthtok(), providing functionality to both the Authentication Stack and the Password Management Stack.
Authentication Service	The implementation of pam_sm_authenticate(3PAM) prompts the user name if not set and then tries to get the authentication token from the pam handle. If the token is not set, it then prompts the user for a password and stores it in the PAM item PAM_AUTHTOK. This module is meant to be the first module on an authentication stack where users are to authenticate using a keyboard.
Password Management Service	<p>Due to the nature of the PAM Password Management stack traversal mechanism, the pam_sm_chauthtok(3PAM) function is called twice. Once with the PAM_PRELIM_CHECK flag, and one with the PAM_UPDATE_AUTHTOK flag.</p> <p>In the first (PRELIM) invocation, the implementation of pam_sm_chauthtok(3PAM) moves the contents of the PAM_AUTHTOK (current authentication token) to PAM_OLDAUTHTOK, and subsequently prompts the user for a new password. This new password is stored in PAM_AUTHTOK.</p> <p>If a previous module has set PAM_AUTHTOK prior to the invocation of pam_authtok_get, this module turns into a NO-OP and immediately returns PAM_SUCCESS.</p> <p>In the second (UPDATE) invocation, the user is prompted to Re-enter his password. The pam_sm_chauthtok implementation verifies this reentered password with the password stored in PAM_AUTHTOK. If the passwords match, the module returns PAM_SUCCESS.</p> <p>The following option can be passed to the module:</p> <p>debug syslog(3C) debugging information at the LOG_DEBUG level</p>
ERRORS	<p>The authentication service returns the following error codes:</p> <p>PAM_SUCCESS Successfully obtains authentication token</p> <p>PAM_SYSTEM_ERR Fails to retrieve username, username is NULL or empty</p> <p>The password management service returns the following error codes:</p> <p>PAM_SUCCESS Successfully obtains authentication token</p> <p>PAM_AUTHTOK_ERR Authentication token manipulation error</p>

pam_authtok_get(5)

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving
MT Level	MT-Safe with exceptions

SEE ALSO `pam(3PAM)`, `pam_authenticate(3PAM)`, `syslog(3C)`, `libpam(3LIB)`, `pam.conf(4)`, `attributes(5)`, `pam_authtok_check(5)`, `pam_authtok_get(5)`, `pam_authtok_store(5)`, `pam_dhkeys(5)`, `pam_passwd_auth(5)`, `pam_unix(5)`, `pam_unix_account(5)`, `pam_unix_auth(5)`, `pam_unix_session(5)`

NOTES The interfaces in `libpam(3LIB)` are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

The `pam_unix(5)` module might not be supported in a future release. Similar functionality is provided by `pam_authtok_check(5)`, `pam_authtok_get(5)`, `pam_authtok_store(5)`, `pam_dhkeys(5)`, `pam_passwd_auth(5)`, `pam_unix_account(5)`, `pam_unix_auth(5)`, and `pam_unix_session(5)`.

NAME	pam_authtok_store – password management module						
SYNOPSIS	pam_authtok_store.so.1						
DESCRIPTION	<p>pam_authtok_store provides functionality to the PAM password management stack. It provides one function: pam_sm_chauthtok().</p> <p>When invoked with flags set to PAM_UPDATE_AUTH Tok, this module updates the authentication token for the user specified by PAM_USER.</p> <p>The authentication token PAM_OLDAUTH Tok can be used to authenticate the user against repositories that need updating (NIS, LDAP). After successful updates, the new authentication token stored in PAM_AUTH Tok is the user's valid password.</p> <p>This module honors the PAM_REPOSITORY item, which, if set, specifies which repository is to be updated. If PAM_REPOSITORY is unset, it follows the nsswitch.conf(4).</p> <p>The following option can be passed to the module:</p> <table border="0"> <tr> <td style="padding-right: 20px;">debug</td> <td>syslog(3C) debugging information at the LOG_DEBUG level</td> </tr> <tr> <td>server_policy</td> <td>If the account authority for the user, as specified by PAM_USER, is a server, do not encrypt the authentication token before updating.</td> </tr> </table>	debug	syslog(3C) debugging information at the LOG_DEBUG level	server_policy	If the account authority for the user, as specified by PAM_USER, is a server, do not encrypt the authentication token before updating.		
debug	syslog(3C) debugging information at the LOG_DEBUG level						
server_policy	If the account authority for the user, as specified by PAM_USER, is a server, do not encrypt the authentication token before updating.						
ERRORS	<table border="0"> <tr> <td style="padding-right: 20px;">PAM_SUCCESS</td> <td>Successfully obtains authentication token</td> </tr> <tr> <td>PAM_SYSTEM_ERR</td> <td>Fails to get username, service name, old password or new password, user name null or empty, or password null.</td> </tr> </table>	PAM_SUCCESS	Successfully obtains authentication token	PAM_SYSTEM_ERR	Fails to get username, service name, old password or new password, user name null or empty, or password null.		
PAM_SUCCESS	Successfully obtains authentication token						
PAM_SYSTEM_ERR	Fails to get username, service name, old password or new password, user name null or empty, or password null.						
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:						
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> <tr> <td>MT Level</td> <td>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving	MT Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
MT Level	MT-Safe with exceptions						
SEE ALSO	pam(3PAM), pam_authenticate(3PAM), pam_chauthtok(3PAM), syslog(3C), libpam(3LIB), pam.conf(4), attributes(5), pam_authtok_check(5), pam_authtok_get(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix(5), pam_unix_account(5), pam_unix_auth(5), pam_unix_session(5)						
NOTES	<p>The interfaces in libpam(3LIB) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.</p> <p>The pam_unix(5) module might not be supported in a future release. Similar functionality is provided by pam_authtok_check(5), pam_authtok_get(5), pam_authtok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix_account(5), pam_unix_auth(5), and pam_unix_session(5).</p>						

pam_dhkeys(5)

NAME	pam_dhkeys – authentication Diffie-Hellman keys management module												
SYNOPSIS	pam_dhkeys.so.1												
DESCRIPTION	<p>The pam_dhkeys.so.1 service module provides functionality to two PAM services: Secure RPC authentication and Secure RPC authentication token management.</p> <p>Secure RPC authentication differs from regular unix authentication because NIS+ and other ONC RPCs use Secure RPC as the underlying security mechanism.</p> <p>The following options may be passed to the module:</p> <p>debug syslog(3C)debugging information at LOG_DEBUG level</p> <p>nowarn Turn off warning messages</p>												
Authentication Services	<p>If the user has Diffie-Hellman keys, pam_sm_authenticate() establishes secret keys for the user specified by the PAM_USER (equivalent to running keylogin(1)), using the authentication token found in the PAM_AUTHOK item. Not being able to establish the secret keys results in an authentication error if the NIS+ repository is used to authenticate the user and the NIS+ table permissions require secure RPC credentials to access the password field. If pam_sm_setcred() is called with PAM_ESTABLISH_CRED and the user's secure RPC credentials need to be established, these credentials are set. This is equivalent to running keylogin(1).</p> <p>If the credentials could not be set and PAM_SILENT is not specified, a diagnostic message is displayed. If pam_setcred() is called with PAM_DELETE_CRED, the user's secure RPC credentials are unset. This is equivalent to running keylogout(1).</p> <p>PAM_REINITIALIZE_CRED and PAM_REFRESH_CRED are not supported and return PAM_IGNORE.</p>												
Authentication Token Management	<p>The pam_sm_chauthtok() implementation checks whether the old login password decrypts the users secret keys. If it doesn't this module prompts the user for an old Secure RPC password and stores it in a pam data item called SUNW_OLDRPCPASS. This data item can be used by the store module to effectively update the users secret keys.</p>												
ERRORS	<p>The authentication service returns the following error codes:</p> <table><tr><td>PAM_SUCCESS</td><td>Credentials set successfully.</td></tr><tr><td>PAM_IGNORE</td><td>Credentials not needed to access the password repository.</td></tr><tr><td>PAM_USER_UNKNOWN</td><td>PAM_USER is not set, or the user is unknown.</td></tr><tr><td>PAM_AUTH_ERR</td><td>No secret keys were set. PAM_AUTHOK is not set, no credentials are present or there is a wrong password.</td></tr><tr><td>PAM_BUF_ERR</td><td>Module ran out of memory.</td></tr><tr><td>PAM_SYSTEM_ERR</td><td>NIS+ subsystem failed .</td></tr></table>	PAM_SUCCESS	Credentials set successfully.	PAM_IGNORE	Credentials not needed to access the password repository.	PAM_USER_UNKNOWN	PAM_USER is not set, or the user is unknown.	PAM_AUTH_ERR	No secret keys were set. PAM_AUTHOK is not set, no credentials are present or there is a wrong password.	PAM_BUF_ERR	Module ran out of memory.	PAM_SYSTEM_ERR	NIS+ subsystem failed .
PAM_SUCCESS	Credentials set successfully.												
PAM_IGNORE	Credentials not needed to access the password repository.												
PAM_USER_UNKNOWN	PAM_USER is not set, or the user is unknown.												
PAM_AUTH_ERR	No secret keys were set. PAM_AUTHOK is not set, no credentials are present or there is a wrong password.												
PAM_BUF_ERR	Module ran out of memory.												
PAM_SYSTEM_ERR	NIS+ subsystem failed .												

pam_dhkeys(5)

The authentication token management returns the following error codes:

PAM_SUCCESS	Old rpc password is set in SUNW_OLDRPCPASS
PAM_USER_UNKNOWN	User in PAM_USER is unknown.
PAM_AUTHOK_ERR	User did not provide a password that decrypts the secret keys.
PAM_BUF_ERR	Module ran out of memory.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving
MT Level	MT-Safe with exceptions

SEE ALSO `keylogin(1)`, `keylogout(1)`, `pam(3PAM)`, `pam_authenticate(3PAM)`, `pam_chautok(3PAM)`, `pam_setcred(3PAM)`, `pam_get_item(3PAM)`, `pam_set_data(3PAM)`, `pam_get_data(3PAM)`, `syslog(3C)`, `libpam(3LIB)`, `pam.conf(4)`, `attributes(5)`, `pam_authok_check(5)`, `pam_authok_get(5)`, `pam_authok_store(5)`, `pam_passwd_auth(5)`, `pam_unix(5)`, `pam_unix_account(5)`, `pam_unix_auth(5)`, `pam_unix_session(5)`

NOTES The interfaces in `libpam(3LIB)` are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

The `pam_unix(5)` module might not be supported in a future release. Similar functionality is provided by `pam_authok_check(5)`, `pam_authok_get(5)`, `pam_authok_store(5)`, `pam_dhkeys(5)`, `pam_passwd_auth(5)`, `pam_unix_account(5)`, `pam_unix_auth(5)`, and `pam_unix_session(5)`.

pam_dial_auth(5)

NAME pam_dial_auth – authentication management PAM module for dialups

SYNOPSIS /usr/lib/security/pam_dial_auth.so.1

DESCRIPTION The dialup PAM module, /usr/lib/security/pam_dial_auth.so.1, authenticates a user according to the /etc/dialups and /etc/d_passwd files. Only pam_sm_authenticate() is implemented within this module. pam_sm_setcred() is a null function. /usr/lib/security/pam_dial_auth.so.1 is designed to be stacked immediately below the /usr/lib/security/pam_unix.so.1 module for the login service.

pam_sm_authenticate() performs authentication only if both the /etc/dialups and /etc/d_passwd files exist. The user's terminal line is checked against entries in the /etc/dialups file. If there is a match, the user's shell is compared against entries in the /etc/d_passwd file. If there is a matching entry, the user is prompted for a password which is validated against the entry in the /etc/d_passwd file. If the passwords match, the user is authenticated. The following option may be passed in to this service module:

debug syslog(3C) debugging information at LOG_DEBUG level.

ATTRIBUTES See attributes(5) for description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	MT-Safe with exceptions

SEE ALSO pam(3PAM), pam_authenticate(3PAM), d_passwd(4), dialups(4), libpam(3LIB), pam.conf(4), attributes(5), pam_authtok_check(5), pam_authtok_get(5), pam_authtok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix(5), pam_unix_account(5), pam_unix_auth(5), pam_unix_session(5)

NOTES The interfaces in libpam(3LIB) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

The pam_unix(5) module might not be supported in a future release. Similar functionality is provided by pam_authtok_check(5), pam_authtok_get(5), pam_authtok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix_account(5), pam_unix_auth(5), and pam_unix_session(5).

NAME	pam_krb5 – authentication, account, session, and password management PAM modules for Kerberos V5										
SYNOPSIS	<code>/usr/lib/security/pam_krb5.so.1</code>										
DESCRIPTION	<p>The Kerberos V5 service module for PAM, <code>/usr/lib/security/pam_krb5.so.1</code>, provides functionality for all four PAM modules: authentication, account management, session management, and password management. The <code>pam_krb5.so.1</code> module is a shared object that can be dynamically loaded to provide the necessary functionality upon demand. Its path is specified in the PAM configuration file.</p>										
Kerberos Authentication Module	<p>The Kerberos V5 authentication component provides functions to verify the identity of a user, <code>pam_sm_authenticate()</code>, and to refresh the Kerberos credentials cache, <code>pam_sm_setcred()</code>. <code>pam_sm_authenticate()</code> authenticates a user principal through the Kerberos authentication service. If the authentication request is successful, the authentication service sends a ticket-granting ticket (tgt) back to the <code>pam_krb5.so.1</code> module, which then verifies that the TGT came from a valid KDC by attempting to get a service ticket for the local host service. For this to succeed, the local host's keytab file (<code>/etc/krb5/krb5.keytab</code>) must contain the entry for the local host service (for example, <code>host/hostname.com@REALM</code> where <code>hostname.com</code> is the fully qualified local hostname and <code>REALM</code> is the default realm of the local host as defined in <code>/etc/krb5/krb5.conf</code>). Once the TGT is verified, it is stored in the credentials cache for later use by Kerberized network applications. If the host entry is not found in the keytab file, the authentication fails.</p> <p>The following options can be passed to the Kerberos V5 authentication module:</p> <table border="0"> <tr> <td style="vertical-align: top;"><code>acceptor</code></td> <td>Prevents the PAM module from performing the authentication service exchange used to obtain the initial ticket-granting ticket. This should be used on Kerberos application servers since the initial ticket is not needed.</td> </tr> <tr> <td style="vertical-align: top;"><code>debug</code></td> <td>Provides <code>syslog(3C)</code> debugging information at <code>LOG_DEBUG</code> level.</td> </tr> <tr> <td style="vertical-align: top;"><code>nowarn</code></td> <td>Turns off warning messages.</td> </tr> <tr> <td style="vertical-align: top;"><code>use_first_pass</code></td> <td>Requests Kerberos V5 authentication with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If Kerberos V5 authentication fails, or if no password has been entered, it quits and does not prompt the user for a password. This option should only be used if the authentication service is designated as <code>optional</code> in the <code>pam.conf</code> configuration file.</td> </tr> <tr> <td style="vertical-align: top;"><code>try_first_pass</code></td> <td>Requests Kerberos V5 authentication with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If</td> </tr> </table>	<code>acceptor</code>	Prevents the PAM module from performing the authentication service exchange used to obtain the initial ticket-granting ticket. This should be used on Kerberos application servers since the initial ticket is not needed.	<code>debug</code>	Provides <code>syslog(3C)</code> debugging information at <code>LOG_DEBUG</code> level.	<code>nowarn</code>	Turns off warning messages.	<code>use_first_pass</code>	Requests Kerberos V5 authentication with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If Kerberos V5 authentication fails, or if no password has been entered, it quits and does not prompt the user for a password. This option should only be used if the authentication service is designated as <code>optional</code> in the <code>pam.conf</code> configuration file.	<code>try_first_pass</code>	Requests Kerberos V5 authentication with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If
<code>acceptor</code>	Prevents the PAM module from performing the authentication service exchange used to obtain the initial ticket-granting ticket. This should be used on Kerberos application servers since the initial ticket is not needed.										
<code>debug</code>	Provides <code>syslog(3C)</code> debugging information at <code>LOG_DEBUG</code> level.										
<code>nowarn</code>	Turns off warning messages.										
<code>use_first_pass</code>	Requests Kerberos V5 authentication with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If Kerberos V5 authentication fails, or if no password has been entered, it quits and does not prompt the user for a password. This option should only be used if the authentication service is designated as <code>optional</code> in the <code>pam.conf</code> configuration file.										
<code>try_first_pass</code>	Requests Kerberos V5 authentication with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If										

pam_krb5(5)

		Kerberos V5 authentication fails, or if no password has been entered, the user is prompted for a password with the prompt "Kerberos Password:".
	use_xfn_pass	Requests Kerberos V5 authentication with a mapped password that has been stored under XFN. If Kerberos V5 authentication fails, or if no password has been entered, it quits and does not prompt the user for a password. This option should only be used if the authentication service is designated as optional in the pam.conf configuration file.
	try_xfn_Pass	Requests Kerberos V5 authentication with a mapped password that has been stored under XFN. If Kerberos V5 authentication fails, or if no password has been stored, the user is prompted for a password with the prompt "Kerberos Password:".
Kerberos V5 Account Management Module		The Kerberos account management component provides a function to perform account management, <code>pam_sm_acct_mgmt()</code> . This function checks to see if the <code>pam_krb5</code> authentication module has noted that the user's password has not expired. The following options may be passed in to the Kerberos V5 service module: debug Provides <code>syslog(3C)</code> debugging information at LOG_DEBUG level nowarn Turn off warning messages.
Kerberos V5 Session Management Module		The Kerberos V5 session management component provides functions to initiate <code>pam_sm_open_session()</code> and terminate <code>pam_sm_close_session()</code> Kerberos V5 sessions. For Kerberos V5, <code>pam_sm_open_session</code> is a null function. <code>pam_close_session</code> destroys a principal's credential cache as well as the kernel Kerberos credentials if the session being closed is the last open session on this server for the calling principal.
Kerberos V5 Password Management Module		The Kerberos V5 password management component provides a function to change passwords <code>pam_sm_chauthtok()</code> in the Key Distribution Center (KDC) database. The following options can be passed in to the Kerberos V5 password module: debug Provides <code>syslog(3C)</code> debugging information at LOG_DEBUG level. nowarn Turns off warning messages. use_first_pass Requests Kerberos V5 authentication with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If Kerberos V5 authentication fails, or if no password has been entered, it quits and does not prompt the user for a password. If authentication succeeds, the user is prompted by "New KRB5 password:" for a new

try_first_pass	password. The user is then prompted a second time for the new password for verification and the KDC database is updated with the new password if both responses match.
use_xfn_pass	Requests Kerberos V5 authentication with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If Kerberos V5 authentication fails, or if no password has been entered, the user is prompted for a password with the prompt "Old KRB5 Password:". If authentication succeeds, the user is prompted by "New KRB5 password:" for a new password. The user is then prompted a second time for the new password for verification and the KDC database is updated with the new password if both responses match.
try_xfn_pass	Requests Kerberos V5 authentication with a mapped password that has been stored under XFN. If Kerberos V5 authentication fails, or if no password has been stored, it quits and does not prompt the user for a password. If authentication succeeds, the user is prompted by "New KRB5 password:" for a new password. The user is then prompted a second time for the new password for verification and the KDC database is updated with the new password if both responses match.

Sample pam.conf File

The following is a sample `pam.conf` configuration file with Kerberos V5 support. Please note that this is only intended to give the flavor of the `pam.conf` Kerberos V5 entries and is not complete.

```
#
# Authentication management
#
login auth required /usr/lib/security/$ISA/pam_unix.so.1
login auth optional /usr/lib/security/$ISA/pam_krb5.so.1 try_first_pass
#
# Account management
```

pam_krb5(5)

```
#
dtlogin account required /usr/lib/security/$ISA/pam_unix.so.1
dtlogin account optional /usr/lib/security/$ISA/pam_krb5.so.1
#
# Session management
#
other session required /usr/lib/security/$ISA/pam_unix.so.1
other session optional /usr/lib/security/$ISA/pam_krb5.so.1
#
# Password management
#
other password required /usr/lib/security/$ISA/pam_unix.so.1
other password optional /usr/lib/security/$ISA/pam_krb5.so.1 try_first_pass
```

The Kerberos V5 module entries typically follow the Unix module entries. Thus, the Kerberos V5 modules are "stacked" behind the Unix module. For the login service, the Kerberos V5 authentication module runs after the Unix module. Its entry is `optional`, so the user can still login if it fails, assuming that the previous Unix module succeeded. If the entry designates `required` instead of `optional`, the user cannot login if Kerberos V5 authentication fails. Because the `try_first_pass` option is designated, it tries the user's password entered for the Unix module. If Kerberos V5 authentication fails, or no password has been entered, the user is prompted for the Kerberos V5 password. For all session related services, the Kerberos V5 session module runs after the Unix module. For the `dtlogin` service, the Kerberos V5 account management module runs after the Unix module. For all password changing related services, the Kerberos V5 module runs after the Unix module. Because the `try_first_pass` option is designated, if the initial password entered for the Unix module authenticates Kerberos V5 successfully, the old Kerberos V5 password is not requested from the user; only the new Kerberos V5 password is requested.

ATTRIBUTES See `attributes(5)` for description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	MT-Safe with exceptions
Interface Stability	Evolving

SEE ALSO `keylogin(1)`, `ktutil(1)`, `pam(3PAM)`, `pam_authenticate(3PAM)`, `syslog(3C)`, `libpam(3LIB)`, `pam.conf(4)`, `attributes(5)`, `SEAM(5)`, `pam_authtok_check(5)`, `pam_authtok_get(5)`, `pam_authtok_store(5)`, `pam_dhkeys(5)`, `pam_passwd_auth(5)`, `pam_unix(5)`, `pam_unix_account(5)`, `pam_unix_auth(5)`, `pam_unix_session(5)`

NOTES The interfaces in `libpam(3LIB)` are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

pam_krb5(5)

The pam_unix(5) module might not be supported in a future release. Similar functionality is provided by pam_authok_check(5), pam_authok_get(5), pam_authok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix_account(5), pam_unix_auth(5), and pam_unix_session(5).

pam_ldap(5)

NAME	pam_ldap – authentication, account, and password management PAM module for LDAP
SYNOPSIS	/usr/lib/security/pam_ldap.so.1
DESCRIPTION	<p>The pam_ldap module implements pam_sm_authenticate, pam_sm_setcred, pam_sm_acct_mgmt, and pam_sm_chauthtok, the functions that provide functionality for the PAM authentication, account management and password management stacks. The pam_ldap module ties the authentication, account management and password change functionality to the functionality of the supporting LDAP server. For authentication, pam_ldap can authenticate the user directly to any LDAP directory server by using any supported authentication mechanism, such as DIGEST-MD5. However, the account management and password change components of pam_ldap will only work with the bundled Sun ONE Directory Server. The Sun ONE Directory Server user account management, that is, password and account lockout policy, must be properly configured on the server before it can be used by pam_ldap to provide the account management, password aging, and password syntax checking controls. Refer to the Sun ONE Directory Server Administrator's Guide that is cited in the NOTES section.</p> <p>pam_ldap must be used in conjunction with the modules that support the UNIX authentication, password, and account management., which are pam_authtok_check(5), pam_authtok_get(5), pam_authtok_store(5), pam_passwd_auth(5), pam_unix_account(5), and pam_unix_auth(5). pam_ldap is designed to be stacked directly below these modules. If other modules are designed to be stacked in this manner, the modules can be stacked below the pam_ldap module. The EXAMPLES section shows how the UNIX modules are stacked with pam_ldap. When stacked together, the UNIX modules are used to control local accounts, such as root. pam_ldap is used for control with the network accounts, that is, LDAP users. For the stacks to work, pam_unix_auth, pam_unix_account, pam_passwd_auth, and pam_authtok_store must to configured with the binding control flag and the server_policy option. This configuration allows local account override of a network account.</p>
LDAP Authentication Module	<p>The LDAP authentication module verifies the identity of a user. The pam_sm_authenticate function uses the password entered by the user to attempt to authenticate to the LDAP server. If successful, the user is authenticated.</p> <p>The authentication method used is either defined in the client profile , or the authentication method is configured by using the ldapclient(1M) command. To determine the authentication method to use, this module first attempts to use the authentication method that is defined, for service pam_ldap, for example, serviceAuthenticationMethod:pam_ldap:sasl/DIGEST-MD5. If no authentication method is defined, pam_ldap uses the default authentication method. If neither are set, the authentication fails. This module skips the configured authentication method if the authentication method is set to none.</p> <p>The pam_sm_setcred(3PAM) function does nothing. This function always returns PAM_IGNORE.</p>

The following options may be passed to the LDAP service module:

debug	syslog(3C) debugging information at LOG_DEBUG level.
nowarn	Turn off warning messages.

These options are case sensitive, and the options must be used exactly as presented here.

LDAP Account Management Module

The LDAP account management module validates the user's account. The `pam_sm_acct_mgmt(3PAM)` function authenticates to the LDAP server to verify that the user's password has not expired, or that the user's account has not been locked. The following options may be passed to the LDAP service module:

debug	syslog(3C) debugging information at LOG_DEBUG level.
nowarn	Turn off warning messages.

These options are case sensitive, and the options must be used exactly as presented here.

LDAP Password Management Module

The preferred way to configure password management for LDAP is by using the `pam_authtok_store(5)` module and by specifying the `server_policy` option. Use the `pam_authtok_store` function instead of `pam_ldap` for password change. When password management is configured this way, both the local and LDAP accounts are handled. `pam_authtok_store(5)` updates the passwords in all the repositories configured by `nsswitch.conf(4)`. `pam_ldap` updates only the password in the LDAP password database.

The LDAP password management module provides the `pam_sm_chauthtok()` function to change passwords in the LDAP database.

The following options may be passed to the LDAP service module:

debug	syslog(3C) debugging information at LOG_DEBUG level.
nowarn	Turn off warning messages.

These options are case sensitive , and the options must be used exactly as presented here.

ERRORS

The authentication service returns the following error codes:

PAM_SUCCESS	Authentication successful
PAM_MAXTRIES	Maximum number of authentication attempts exceeded
PAM_AUTH_ERR	Authentication failure
PAM_USER_UNKNOWN	No account present for user

pam_ldap(5)

PAM_BUF_ERR	Memory buffer error
PAM_SYSTEM_ERR	System error
The account management service returns the following error codes:	
PAM_SUCCESS	User allowed access to account
PAM_NEW_AUTHTOK_REQD	New authentication token required
PAM_ACCT_EXPIRED	User account has expired
PAM_PERM_DENIED	User denied access to account at this time
PAM_USER_UNKNOWN	No account present for user
PAM_BUF_ERROR	Memory buffer error
PAM_SYSTEM_ERR	System error

The password management service returns the following values:

PAM_SUCCESS	Successfully updates authentication token
PAM_PERM_DENIED	No permission to update authentication token
PAM_AUTHTOK_ERR	Authentication token manipulation error
PAM_USER_UNKNOWN	No account present for user
PAM_BUF_ERR	Memory buffer error
PAM_SYSTEM_ERR	System error

EXAMPLES

EXAMPLE 1 Using pam_ldap With Authentication

The following is a configuration for the login service when using pam_ldap. The service name login can be substituted for any other authentication service such as dtlogin or su. Lines that begin with the # symbol are comments, and these lines ignored.

```
# Authentication management for login service is stacked.
# If pam_unix_auth succeeds, pam_ldap is not invoked.
# The control flag "binding" provides a local overriding
# remote (LDAP) control. The "server_policy" option is used
# to tell pam_unix_auth.so.1 to ignore the LDAP users.

login  auth requisite pam_authtok_get.so.1
login  auth required  pam_dhkeys.so.1
login  auth binding   pam_unix_auth.so.1 server_policy
login  auth required  pam_ldap.so.1
```

EXAMPLE 2 Using pam_ldap With Account Management

The following is a configuration for account management when using pam_ldap. Lines that begin with the # symbol are ignored.

EXAMPLE 2 Using pam_ldap With Account Management (Continued)

```

# Account management for all services is stacked
# If pam_unix_account succeeds, pam_ldap is not invoked.
# The control flag "binding" provides a local overriding
# remote (LDAP) control. The "server_policy" option is used
# to tell pam_unix_account.so.1 to ignore the LDAP users.

other account requisite pam_roles.so.1
other account required pam_projects.so.1
other account binding pam_unix_account.so.1 server_policy
other account required pam_ldap.so.1

```

EXAMPLE 3 Using pam_authtok_store With Password Management For Both Local and LDAP Accounts

The following is a configuration for password management when using pam_authtok_store instead of pam_ldap. This configuration works because pam_authtok_store updates password in all the repositories configured by nsswitch.conf(4). Lines that begin with the # symbol are comments, and the lines are ignored.

```

# Password management (authentication)
passwd auth binding pam_passwd_auth.so.1 server_policy
passwd auth required pam_ldap.so.1

# Password management (updates)
# This is the preferred stack, since it updates
# passwords stored both in the local /etc files and
# in the LDAP directory. The "server_policy"
# option is used to tell pam_authtok_store to
# follow the LDAP server's policy when updating
# passwords stored in the LDAP directory

other password required pam_dhkeys.so.1
other password requisite pam_authtok_get.so.1
other password requisite pam_authtok_check.so.1
other password required pam_authtok_store.so.1 server_policy

```

EXAMPLE 4 Using pam_ldap With Password Management if There are no Local Accounts

Use the following configuration for password management when using pam_ldap. Lines that begin with the # symbol are comments, and the comments are ignored.

```

# Password management (authentication)
# The control flag "binding" provides a local overriding
# remote (LDAP) control. The server_policy option is used
# to tell pam_passwd_auth.so.1 to ignore the LDAP users.

passwd auth binding pam_passwd_auth.so.1 server_policy
passwd auth required pam_ldap.so.1

# Password management (updates)
# This stack is limited to updating password stored in the

```

pam_ldap(5)

EXAMPLE 4 Using pam_ldap With Password Management if There are no Local Accounts *(Continued)*

```
# LDAP directory. The preferred method is shown in Example 3.  
other password required pam_ldap.so.1
```

FILES /var/ldap/ldap_client_file
/var/ldap/ldap_client_cred The LDAP configuration files of the client. Do not manually modify these files, as these files may not be human readable. Use ldapclient(1M) to update these files.

/etc/pam.conf PAM configuration file.

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe with exceptions
Stability Level	Evolving

SEE ALSO ldap(1), idsconfig(1M), ldap_cachemgr(1M), ldapclient(1M), libpam(3LIB), pam(3PAM), pam_sm_authenticate(3PAM), pam_sm_chauthtok(3PAM), pam_sm_setcred(3PAM), syslog(3C), pam.conf(4), attributes(5), pam_authtok_check(5), pam_authtok_get(5), pam_authtok_store(5), pam_passwd_auth(5), pam_unix_account(5), pam_unix_auth(5)

NOTES The interfaces in libpam(3LIB) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

For information on how to configure the user account management, including password and account lockout policy for the bundled Sun ONE Directory Server, please browse the html file
</usr/iplanet/ds5/manual/en/slapd/ag/password.htm>.

NAME	pam_passwd_auth – authentication module for password								
SYNOPSIS	pam_passwd_auth.so.1								
DESCRIPTION	<p>pam_passwd_auth provides authentication functionality to the password service as implemented by passwd(1). It differs from the standard PAM authentication modules in its prompting behavior.</p> <p>The name of the user whose password attributes are to be updated must be present in the PAM_USER item. This can be accomplished due to a previous call to pam_start(3PAM), or explicitly set by pam_set_item(3PAM). Based on the current user-id and the repository that is to be updated, the module determines whether a password is necessary for a successful update of the password repository, and if so, which password is required.</p> <p>The following option can be passed to the module:</p> <table border="0"> <tr> <td style="padding-right: 20px;">debug</td> <td>syslog(3C) debugging information at the LOG_DEBUG level</td> </tr> <tr> <td>nowarn</td> <td>Turn off warning messages</td> </tr> <tr> <td>server_policy</td> <td>If the account authority for the user, as specified by PAM_USER, is a server, do not apply the Unix policy from the passwd entry in the name service switch.</td> </tr> </table>	debug	syslog(3C) debugging information at the LOG_DEBUG level	nowarn	Turn off warning messages	server_policy	If the account authority for the user, as specified by PAM_USER, is a server, do not apply the Unix policy from the passwd entry in the name service switch.		
debug	syslog(3C) debugging information at the LOG_DEBUG level								
nowarn	Turn off warning messages								
server_policy	If the account authority for the user, as specified by PAM_USER, is a server, do not apply the Unix policy from the passwd entry in the name service switch.								
ERRORS	<p>The following error codes are returned:</p> <table border="0"> <tr> <td style="padding-right: 20px;">PAM_BUF_ERR</td> <td>Memory buffer error</td> </tr> <tr> <td>PAM_IGNORE</td> <td>Ignore module, not participating in result</td> </tr> <tr> <td>PAM_SUCCESS</td> <td>Successfully obtains authentication token</td> </tr> <tr> <td>PAM_SYSTEM_ERR</td> <td>System error</td> </tr> </table>	PAM_BUF_ERR	Memory buffer error	PAM_IGNORE	Ignore module, not participating in result	PAM_SUCCESS	Successfully obtains authentication token	PAM_SYSTEM_ERR	System error
PAM_BUF_ERR	Memory buffer error								
PAM_IGNORE	Ignore module, not participating in result								
PAM_SUCCESS	Successfully obtains authentication token								
PAM_SYSTEM_ERR	System error								
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:								
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> <tr> <td>MT Level</td> <td>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving	MT Level	MT-Safe with exceptions		
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Interface Stability	Evolving								
MT Level	MT-Safe with exceptions								
SEE ALSO	passwd(1), pam(3PAM), pam_authenticate(3PAM), pam_start(3PAM), pam_set_item(3PAM), syslog(3C), libpam(3LIB), pam.conf(4), attributes(5), pam_authtok_check(5), pam_authtok_get(5), pam_authtok_store(5), pam_dhkeys(5), pam_unix(5), pam_unix_account(5), pam_unix_auth(5), pam_unix_session(5)								
NOTES	The interfaces in libpam(3LIB) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.								

pam_passwd_auth(5)

This module relies on the value of the current real UID, this module is only safe for MT-applications that don't change UIDs during the call to `pam_authenticate(3PAM)`.

The `pam_unix(5)` module might not be supported in a future release. Similar functionality is provided by `pam_authtok_check(5)`, `pam_authtok_get(5)`, `pam_authtok_store(5)`, `pam_dhkeys(5)`, `pam_passwd_auth(5)`, `pam_unix_account(5)`, `pam_unix_auth(5)`, and `pam_unix_session(5)`.

NAME	pam_projects – account management PAM module for projects				
SYNOPSIS	/usr/lib/security/pam_projects.so.1				
DESCRIPTION	<p>The projects service module for PAM, /usr/lib/security/pam_projects.so.1, provides functionality for the account management PAM module. The pam_projects.so.1 module is a shared object that can be dynamically loaded to provide the necessary functionality upon demand. Its path is specified in the PAM configuration file.</p> <p>pam_projects.so.1 is designed to be stacked on top of the pam_unix_account.so.1 module for all services. This module is normally configured as “required”, implying that any user lacking a default project will be denied login.</p> <p>The project account management component provides a function to perform account management, pam_sm_acct_mgmt(). This function uses the getdefaultproj() function (see getprojent(3PROJECT)) to retrieve the user’s default project entry from the project(4) database. It then sets the project ID attribute of the calling process, using the settaskid(2) system call.</p> <p>If the user does not belong to any project defined in the project(4) database, or if the settaskid() system call failed to set the project ID attribute of the calling process, the module will display an error message and will return error code PAM_PERM_DENIED.</p>				
Projects Account Management Module					
ATTRIBUTES	<p>See attributes(5) for description of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe with exceptions				
SEE ALSO	<p>settaskid(2), getprojent(3PROJECT), libpam(3LIB), pam(3PAM), pam_acct_mgmt(3PAM), pam.conf(4), project(4), attributes(5)</p> <p>, pam_authtok_check(5), pam_authtok_get(5), pam_authtok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix(5), pam_unix_account(5), pam_unix_auth(5), pam_unix_session(5)</p>				
NOTES	<p>The interfaces in libpam(3LIB) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.</p> <p>The pam_unix(5) module might not be supported in a future release. Similar functionality is provided by pam_authtok_check(5), pam_authtok_get(5), pam_authtok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix_account(5), pam_unix_auth(5), and pam_unix_session(5).</p>				

pam_rhosts_auth(5)

NAME pam_rhosts_auth – authentication management PAM module using ruserok()

SYNOPSIS /usr/lib/security/pam_rhosts_auth.so.1

DESCRIPTION The rhosts PAM module, /usr/lib/security/pam_rhosts_auth.so.1, authenticates a user via the rlogin authentication protocol. Only pam_sm_authenticate() is implemented within this module. pam_sm_authenticate() uses the ruserok(3SOCKET) library function to authenticate the rlogin or rsh user. pam_sm_setcred() is a null function.

/usr/lib/security/pam_rhosts_auth.so.1 is designed to be stacked on top of the /usr/lib/security/pam_unix.so.1 module for both the rlogin and rsh services. This module is normally configured as *sufficient* so that subsequent authentication is performed only on failure of pam_sm_authenticate(). The following option may be passed in to this service module:

debug syslog(3C) debugging information at LOG_DEBUG level.

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	MT-Safe with exceptions

SEE ALSO pam(3PAM), pam_authenticate(3PAM), ruserok(3SOCKET), syslog(3C), libpam(3LIB), pam.conf(4), attributes(5)

NOTES The interfaces in libpam() are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

NAME	pam_roles – Role Account Management PAM module for Solaris
SYNOPSIS	<code>/usr/lib/security/pam_roles.so.1</code>
DESCRIPTION	<p>The Role Account Management module for PAM, <code>/usr/lib/security/pam_roles.so.1</code>, provides functionality for one PAM module: Account management. The <code>pam_roles.so.1</code> is a shared object that can be dynamically loaded to provide the necessary functionality upon demand. Its path is specified in the PAM configuration file.</p>
Role Account Management Module	<p>The Role account management component provides a function to check for authorization to assume a role. It prevents direct logins to a role. It uses the <code>user_attr(4)</code> database to specify which users can assume which roles.</p> <p>The following options may be passed to the Role Authentication service module:</p> <p><code>debug</code> <code>syslog(3C)</code> debugging information at <code>LOG_DEBUG</code> level.</p> <p>If <code>PAM_USER</code> (see <code>pam_set_item(3PAM)</code>) is specified as type normal in the <code>user_attr(4)</code> database, the module returns <code>PAM_IGNORE</code>.</p> <p>If <code>PAM_RUSER</code> (see <code>pam_set_item(3PAM)</code>) is not set, the uid of the process loading the module is used to determine <code>PAM_RUSER</code>.</p> <p>The module returns success if the <code>user_attr(4)</code> entry for <code>PAM_RUSER</code> has an entry in the roles field for <code>PAM_USER</code>; otherwise it returns <code>PAM_PERM_DENIED</code>.</p> <p>This module is generally stacked above the account management module <code>pam_unix.so.1</code>. The error messages indicating that roles cannot be logged into correctly are only issued if the user has entered the correct password.</p> <p>Here are some sample entries from <code>pam.conf(4)</code> demonstrating the use of the <code>pam_roles.so.1</code> module:</p> <pre>dtlogin account requisite /usr/lib/security/\$ISA/pam_roles.so.1 dtlogin account required /usr/lib/security/\$ISA/pam_unix.so.1 # su account requisite /usr/lib/security/\$ISA/pam_roles.so.1 su account requisite /usr/lib/security/\$ISA/pam_roles.so.1 # rlogin account requisite /usr/lib/security/\$ISA/pam_roles.so.1 rlogin account required /usr/lib/security/\$ISA/pam_unix.so.1 #</pre> <p>The <code>dtlogin</code> program invokes <code>pam_roles.so.1</code>. <code>PAM_RUSER</code> is the username corresponding to the uid of the <code>dtlogin</code> process, which is 0. The <code>user_attr</code> entry for root user (uid 0) is empty, so all role logins are prevented through <code>dtlogin</code>. The same rule applies to login.</p> <p>The <code>su</code> program invokes <code>pam_roles.so.1</code>. <code>PAM_RUSER</code> is the username of the userid of the shell that invokes <code>su</code>. A user needs the appropriate entry in the roles list in <code>user_attr(4)</code> to be able to <code>su</code> to another user.</p>

pam_roles(5)

In the example above, the `rlogin` program invokes the `pam_roles.so.1` module. The module checks for `PAM_RUSER` and determines whether the role being assumed, `PAM_RUSER`, is in the roles list of the `userattr` entry for `PAM_RUSER`. If it is in the roles list, the module returns `PAM_SUCCESS`; otherwise it returns `PAM_PERM_DENIED`.

SEE ALSO `keylogin(1)`, `libpam(3LIB)`, `pam(3PAM)`, `pam_acct_mgmt(3PAM)`, `pam_setcred(3PAM)`, `pam_set_item(3PAM)`, `syslog(3C)`, `pam.conf(4)`, `user_attr(4)`, `attributes(5)`, `pam_authtok_check(5)`, `pam_authtok_get(5)`, `pam_authtok_store(5)`, `pam_dhkeys(5)`, `pam_passwd_auth(5)`, `pam_unix(5)`, `pam_unix_account(5)`, `pam_unix_auth(5)`, `pam_unix_session(5)`

NOTES The interfaces in `libpam(3LIB)` are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

The `pam_unix(5)` module might not be supported in a future release. Similar functionality is provided by `pam_authtok_check(5)`, `pam_authtok_get(5)`, `pam_authtok_store(5)`, `pam_dhkeys(5)`, `pam_passwd_auth(5)`, `pam_unix_account(5)`, `pam_unix_auth(5)`, and `pam_unix_session(5)`.

NAME	pam_sample – a sample PAM module																		
SYNOPSIS	/usr/lib/security/pam_sample.so.1																		
DESCRIPTION	The SAMPLE service module for PAM is divided into four components: authentication, account management, password management, and session management. The sample module is a shared object that is dynamically loaded to provide the necessary functionality.																		
SAMPLE Authentication Component	<p>The SAMPLE authentication module, typically /usr/lib/security/pam_sample.so.1, provides functions to test the PAM framework functionality using the pam_sm_authenticate(3PAM) call. The SAMPLE module implementation of the pam_sm_authenticate(3PAM) function compares the user entered password with the password set in the pam.conf(4) file, or the string "test" if a default test password has not been set. The following options may be passed in to the SAMPLE Authentication module:</p> <table border="0"> <tr> <td style="padding-right: 20px;">debug</td> <td>Syslog debugging information at the LOG_DEBUG level.</td> </tr> <tr> <td>passwd=newone</td> <td>Sets the password to be "newone."</td> </tr> <tr> <td>first_pass_good</td> <td>The first password is always good when used with the use_first_pass or try_first_pass option.</td> </tr> <tr> <td>first_pass_bad</td> <td>The first password is always bad when used with the use_first_pass or try_first_pass option.</td> </tr> <tr> <td>always_fail</td> <td>Always returns PAM_AUTH_ERR.</td> </tr> <tr> <td>always_succeed</td> <td>Always returns PAM_SUCCESS.</td> </tr> <tr> <td>always_ignore</td> <td>Always returns PAM_IGNORE.</td> </tr> <tr> <td>use_first_pass</td> <td>Use the user's initial password (entered when the user is authenticated to the first authentication module in the stack) to authenticate with the SAMPLE module. If the passwords do not match, or if this is the first authentication module in the stack, quit and do not prompt the user for a password. It is recommended that this option only be used if the SAMPLE authentication module is designated as <i>optional</i> in the pam.conf configuration file.</td> </tr> <tr> <td>try_first_pass</td> <td>Use the user's initial password (entered when the user is authenticated to the first authentication module in the stack) to authenticate with the SAMPLE module. If the passwords do not match, or if this is the first authentication module in the stack, prompt the user for a password. The SAMPLE module pam_sm_setcred(3PAM) function always returns PAM_SUCCESS.</td> </tr> </table>	debug	Syslog debugging information at the LOG_DEBUG level.	passwd=newone	Sets the password to be "newone."	first_pass_good	The first password is always good when used with the use_first_pass or try_first_pass option.	first_pass_bad	The first password is always bad when used with the use_first_pass or try_first_pass option.	always_fail	Always returns PAM_AUTH_ERR.	always_succeed	Always returns PAM_SUCCESS.	always_ignore	Always returns PAM_IGNORE.	use_first_pass	Use the user's initial password (entered when the user is authenticated to the first authentication module in the stack) to authenticate with the SAMPLE module. If the passwords do not match, or if this is the first authentication module in the stack, quit and do not prompt the user for a password. It is recommended that this option only be used if the SAMPLE authentication module is designated as <i>optional</i> in the pam.conf configuration file.	try_first_pass	Use the user's initial password (entered when the user is authenticated to the first authentication module in the stack) to authenticate with the SAMPLE module. If the passwords do not match, or if this is the first authentication module in the stack, prompt the user for a password. The SAMPLE module pam_sm_setcred(3PAM) function always returns PAM_SUCCESS.
debug	Syslog debugging information at the LOG_DEBUG level.																		
passwd=newone	Sets the password to be "newone."																		
first_pass_good	The first password is always good when used with the use_first_pass or try_first_pass option.																		
first_pass_bad	The first password is always bad when used with the use_first_pass or try_first_pass option.																		
always_fail	Always returns PAM_AUTH_ERR.																		
always_succeed	Always returns PAM_SUCCESS.																		
always_ignore	Always returns PAM_IGNORE.																		
use_first_pass	Use the user's initial password (entered when the user is authenticated to the first authentication module in the stack) to authenticate with the SAMPLE module. If the passwords do not match, or if this is the first authentication module in the stack, quit and do not prompt the user for a password. It is recommended that this option only be used if the SAMPLE authentication module is designated as <i>optional</i> in the pam.conf configuration file.																		
try_first_pass	Use the user's initial password (entered when the user is authenticated to the first authentication module in the stack) to authenticate with the SAMPLE module. If the passwords do not match, or if this is the first authentication module in the stack, prompt the user for a password. The SAMPLE module pam_sm_setcred(3PAM) function always returns PAM_SUCCESS.																		

pam_sample(5)

SAMPLE Account Management Component

The SAMPLE Account Management Component, typically `pam_sample.so.1`, implements a simple access control scheme that limits machine access to a list of authorized users. The list of authorized users is supplied as option arguments to the entry for the SAMPLE account management PAM module in the `pam.conf` file. Note that the module always permits access to the root super user.

The option field syntax to limit access is shown below: `allow= name[,name] allow=name [allow=name]`

The example `pam.conf` show below permits only larry to login directly. `rlogin` is allowed only for don and larry. Once a user is logged in, the user can use `su` if the user are sam or eric.

login	account	require	pam_sample.so.1	allow=larry
dtlogin	account	require	pam_sample.so.1	allow=larry
rlogin	account	require	pam_sample.so.1	allow=don allow=larry
su	account	require	pam_sample.so.1	allow=sam,eric

The debug and nowarn options are also supported.

SAMPLE Password Management Component
SAMPLE Session Management Component

The SAMPLE Password Management Component function (`pam_sm_chauthtok(3PAM)`), always returns `PAM_SUCCESS`.

The SAMPLE Session Management Component functions (`pam_sm_open_session(3PAM)`, `pam_sm_close_session(3PAM)`) always return `PAM_SUCCESS`.

ATTRIBUTES

See `attributes(5)` for description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	MT-Safe with exceptions

SEE ALSO

`pam(3PAM)`, `pam_sm_authenticate(3PAM)`, `pam_sm_chauthtok(3PAM)`, `pam_sm_close_session(3PAM)`, `pam_sm_open_session(3PAM)`, `pam_sm_setcred(3PAM)`, `libpam(3LIB)`, `pam.conf(4)`, `attributes(5)`

NOTES

The interfaces in `libpam()` are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

NAME	pam_smartcard – PAM authentication module for Smart Card
SYNOPSIS	<code>/usr/lib/security/pam_smartcard.so</code>
DESCRIPTION	<p>The Smart Card service module for PAM, <code>/usr/lib/security/pam_smartcard.so</code>, provides functionality to obtain a user's information (such as user name and password) for a smart card. The <code>pam_smartcard.so</code> module is a shared object that can be dynamically loaded to provide the necessary functionality upon demand. Its path is specified in the PAM configuration file <code>pam.conf</code>. See <code>pam.conf(4)</code>.</p> <p>Smart Card Authentication Module</p> <p>The Smart Card authentication component provides the <code>pam_sm_authenticate(3PAM)</code> function to verify the identity of a smart card user.</p> <p>The <code>pam_sm_authenticate()</code> function collects as user input the PIN number. It passes this data back to its underlying layer, OCF, to perform PIN verification. If verification is successful, the module returns <code>PAM_SUCCESS</code>, and passes the username and password from the smart card to PAM modules stacked below <code>pam_smartcard</code>.</p> <p>The following options can be passed to the Smart Card service module:</p> <p><code>debug</code> <code>syslog(3C)</code> debugging information at <code>LOG_DEBUG</code> level.</p> <p><code>nowarn</code> Turn off warning messages.</p> <p><code>verbose</code> Turn on verbose authentication failure reporting to the user.</p> <p>Smart Card Module Configuration</p> <p>The PAM smart card module (<code>pam_smartcard</code>) can be configured in the PAM configuration file (<code>/etc/pam.conf</code>). For example, the following configuration on on the desktop (Common Desktop Environment) forces a user to use a smart card for logging in.</p> <p>The following are typical values set by '<code>smartcard -c enable</code>', if the command is applied to the default configuration.</p> <pre>dtlogin auth requisite pam_smartcard.so.1 dtlogin auth required pam_authtok_get.so.1 dtlogin auth required pam_dhkeys.so.1 dtlogin auth required pam_unix_auth.so.1 dtsession auth requisite pam_smartcard.so.1 dtsession auth required pam_authtok_get.so.1 dtsession auth required pam_dhkeys.so.1 dtsession auth required pam_unix_auth.so.1</pre>
SEE ALSO	<p><code>smartcard(1M)</code>, <code>libpam(3LIB)</code>, <code>pam(3PAM)</code>, <code>pam_authenticate(3PAM)</code>, <code>pam_start(3PAM)</code>, <code>pam.conf(4)</code>, <code>pam_authtok_check(5)</code>, <code>pam_authtok_get(5)</code>, <code>pam_authtok_store(5)</code>, <code>pam_dhkeys(5)</code>, <code>pam_passwd_auth(5)</code>, <code>pam_unix(5)</code>, <code>pam_unix_account(5)</code>, <code>pam_unix_auth(5)</code>, <code>pam_unix_session(5)</code></p>

pam_smartcard(5)

NOTES | The pam_unix(5) module might not be supported in a future release. Similar functionality is provided by pam_authok_check(5), pam_authok_get(5), pam_authok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix_account(5), pam_unix_auth(5), and pam_unix_session(5).

NAME	pam_unix – authentication, account, session, and password management PAM modules for UNIX								
SYNOPSIS	/usr/lib/security/pam_unix.so.1								
DESCRIPTION	The UNIX service module for PAM, /usr/lib/security/pam_unix.so.1, provides functionality for all four PAM modules: authentication, account management, session management and password management. The pam_unix.so.1 module is a shared object that can be dynamically loaded to provide the necessary functionality upon demand. Its path is specified in the PAM configuration file.								
UNIX Authentication Module	<p>The UNIX authentication component provides functions to verify the identity of a user, (pam_sm_authenticate()) and to set user specific credentials (pam_sm_setcred()). pam_sm_authenticate() compares the user entered password with the password from the UNIX password database. If the passwords match, the user is authenticated. If the user also has secure RPC credentials and the secure RPC password is the same as the UNIX password, then the secure RPC credentials are also obtained.</p> <p>The following options may be passed to the UNIX service module:</p> <table border="0"> <tr> <td style="vertical-align: top;">debug</td> <td>syslog(3C) debugging information at LOG_DEBUG level.</td> </tr> <tr> <td style="vertical-align: top;">nowarn</td> <td>Turn off warning messages.</td> </tr> <tr> <td style="vertical-align: top;">use_first_pass</td> <td>It compares the password in the password database with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If the passwords do not match, or if no password has been entered, it quits and does not prompt the user for a password. This option should only be used if the authentication service is designated as <i>optional</i> in the pam.conf configuration file.</td> </tr> <tr> <td style="vertical-align: top;">try_first_pass</td> <td> <p>It compares the password in the password database with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If the passwords do not match, or if no password has been entered, prompt the user for a password. When prompting for the current password, the UNIX authentication module will use the prompt, "password:" unless one of the following scenarios occur:</p> <ol style="list-style-type: none"> 1. The option try_first_pass is specified and the password entered for the first module in the stack fails for the UNIX module. </td> </tr> </table>	debug	syslog(3C) debugging information at LOG_DEBUG level.	nowarn	Turn off warning messages.	use_first_pass	It compares the password in the password database with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If the passwords do not match, or if no password has been entered, it quits and does not prompt the user for a password. This option should only be used if the authentication service is designated as <i>optional</i> in the pam.conf configuration file.	try_first_pass	<p>It compares the password in the password database with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If the passwords do not match, or if no password has been entered, prompt the user for a password. When prompting for the current password, the UNIX authentication module will use the prompt, "password:" unless one of the following scenarios occur:</p> <ol style="list-style-type: none"> 1. The option try_first_pass is specified and the password entered for the first module in the stack fails for the UNIX module.
debug	syslog(3C) debugging information at LOG_DEBUG level.								
nowarn	Turn off warning messages.								
use_first_pass	It compares the password in the password database with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If the passwords do not match, or if no password has been entered, it quits and does not prompt the user for a password. This option should only be used if the authentication service is designated as <i>optional</i> in the pam.conf configuration file.								
try_first_pass	<p>It compares the password in the password database with the user's initial password (entered when the user authenticated to the first authentication module in the stack). If the passwords do not match, or if no password has been entered, prompt the user for a password. When prompting for the current password, the UNIX authentication module will use the prompt, "password:" unless one of the following scenarios occur:</p> <ol style="list-style-type: none"> 1. The option try_first_pass is specified and the password entered for the first module in the stack fails for the UNIX module. 								

pam_unix(5)

2. The option `try_first_pass` is not specified, and the earlier authentication modules listed in the `pam.conf` file have prompted the user for the password.

In these two cases, the UNIX authentication module will use the prompt "SYSTEM password:". The `pam_sm_setcred()` function sets user specific credentials. If the user had secure RPC credentials, but the secure RPC password was not the same as the UNIX password, then a warning message is printed. If the user wants to get secure RPC credentials, then `keylogin(1)` needs to be run.

UNIX Account Management Module

The UNIX account management component provides a function to perform account management, `pam_sm_acct_mgmt()`. The function retrieves the user's password entry from the UNIX password database and verifies that the user's account and password have not expired. The following options may be passed in to the UNIX service module:

`debug` `syslog(3C)` debugging information at `LOG_DEBUG` level.
`nowarn` Turn off warning messages.

UNIX Session Management Module

The UNIX session management component provides functions to initiate `pam_sm_open_session()` and terminate `pam_sm_close_session()` UNIX sessions. For UNIX, `pam_open_session` updates the `/var/adm/lastlog` file. The account management module reads this file to determine the previous time the user logged in. The following options may be passed in to the UNIX service module:

`debug` `syslog(3C)` debugging information at `LOG_DEBUG` level.
`nowarn` Turn off warning messages. `pam_close_session` is a null function.

UNIX Password Management Module

The UNIX password management component provides a function to change passwords `pam_sm_chauthtok()` in the UNIX password database. This module must be *required* in `pam.conf`. It cannot be *optional* or *sufficient*. The following options may be passed in to the UNIX service module:

`debug` `syslog(3C)` Debugging information at `LOG_DEBUG` level.
`nowarn` Turn off warning messages.
`use_first_pass` It compares the password in the password database with the user's old password (entered to the first password module in the stack). If the passwords do not match, or if no password has been entered, it quits and does not prompt the user for the old password. It also attempts to use the new password (entered to the first

pam_unix(5)

password module in the stack) as the new password for this module. If the new password fails, it quits and does not prompt the user for a new password.

try_first_pass

It compares the password in the password database with the user's old password (entered to the first password module in the stack). If the passwords do not match, or if no password has been entered, it prompts the user for the old password. It also attempts to use the new password (entered to the first password module in the stack) as the new password for this module. If the new password fails, it prompts the user for a new password. If the user's password has expired, the UNIX account module saves this information in the authentication handle using `pam_set_data()`, with a unique name, `SUNW_UNIX_AUTHOK_DATA`. The UNIX password module retrieves this information from the authentication handle using `pam_get_data()` to determine whether or not to force the user to update the user's password.

ATTRIBUTES See `attributes(5)` for description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	MT-Safe with exceptions

SEE ALSO `keylogin(1)`, `pam(3PAM)`, `pam_authenticate(3PAM)`, `pam_setcred(3PAM)`, `syslog(3C)`, `libpam(3LIB)`, `pam.conf(4)`, `attributes(5)`, `pam_authok_check(5)`, `pam_authok_get(5)`, `pam_authok_store(5)`, `pam_dhkeys(5)`, `pam_passwd_auth(5)`, `pam_unix_account(5)`, `pam_unix_auth(5)`, `pam_unix_session(5)`

NOTES The interfaces in `libpam(3LIB)` are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

The `pam_unix(5)` module might not be supported in a future release. Similar functionality is provided by `pam_authok_check(5)`, `pam_authok_get(5)`, `pam_authok_store(5)`, `pam_dhkeys(5)`, `pam_passwd_auth(5)`, `pam_unix_account(5)`, `pam_unix_auth(5)`, and `pam_unix_session(5)`.

pam_unix_account(5)

NAME pam_unix_account – PAM account management module for UNIX

SYNOPSIS pam_unix_account.so.1

DESCRIPTION pam_unix_account provides functionality to the PAM account management stack. The function pam(3PAM) function retrieves password aging information from the repositories specified in nsswitch.conf(4), and verifies that the user's account and password have not expired.

The following options can be passed to the module:

debug syslog(3C) debugging information at the LOG_DEBUG level

nowarn Turn off warning messages

server_policy If the account authority for the user, as specified by PAM_USER, is a server, do not apply the Unix policy from the passwd entry in the name service switch.

ERRORS The following values are returned:

PAM_AUTHOK_EXPIRED Password expired and no longer usable

PAM_BUF_ERR Memory buffer error

PAM_IGNORE Ignore module, not participating in result

PAM_NEW_AUTHOK_REQD Obtain new authentication token from the user

PAM_SERVICE_ERR Error in underlying service module

PAM_SUCCESS Successfully obtains authentication token

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving
MT Level	MT-Safe with exceptions

SEE ALSO pam(3PAM), pam_authenticate(3PAM), syslog(3C), libpam(3LIB), pam.conf(4), nsswitch.conf(4), attributes(5), pam_authok_check(5), pam_authok_get(5), pam_authok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix(5), pam_unix_auth(5), pam_unix_session(5)

NOTES The interfaces in libpam(3LIB) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

The pam_unix(5) module might not be supported in a future release. Similar functionality is provided by pam_authok_check(5), pam_authok_get(5), pam_authok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix_account(5), pam_unix_auth(5), and pam_unix_session(5).

NAME	pam_unix_auth – PAM authentication module for UNIX														
SYNOPSIS	pam_unix_auth.so.1														
DESCRIPTION	<p>The <code>pam_unix_auth</code> module implements <code>pam_sm_authenticate()</code>, which provides functionality to the PAM authentication stack. It provides functions to verify that the password contained in the PAM item <code>PAM_AUTHTOK</code> is the correct password for the user specified in the item <code>PAM_USER</code>. If <code>PAM_REPOSITORY</code> is specified, then user's passwd is fetched from that repository. Otherwise the default <code>nsswitch.conf(4)</code> repository is searched for that user.</p> <p>The following options can be passed to the module:</p> <p><code>server_policy</code> If the account authority for the user, as specified by <code>PAM_USER</code>, is a server, do not apply the Unix policy from the <code>passwd</code> entry in the name service switch.</p>														
ERRORS	<p>The following values are returned:</p> <table border="0"> <tr> <td><code>PAM_AUTH_ERR</code></td> <td>Authentication failure</td> </tr> <tr> <td><code>PAM_BUF_ERR</code></td> <td>Memory buffer error</td> </tr> <tr> <td><code>PAM_IGNORE</code></td> <td>Ignore module, not participating in result</td> </tr> <tr> <td><code>PAM_PERM_DENIED</code></td> <td>Permission denied</td> </tr> <tr> <td><code>PAM_SUCCESS</code></td> <td>Successfully obtains authentication token</td> </tr> <tr> <td><code>PAM_SYSTEM_ERR</code></td> <td>System error</td> </tr> <tr> <td><code>PAM_USER_UNKNOWN</code></td> <td>No account present for user</td> </tr> </table>	<code>PAM_AUTH_ERR</code>	Authentication failure	<code>PAM_BUF_ERR</code>	Memory buffer error	<code>PAM_IGNORE</code>	Ignore module, not participating in result	<code>PAM_PERM_DENIED</code>	Permission denied	<code>PAM_SUCCESS</code>	Successfully obtains authentication token	<code>PAM_SYSTEM_ERR</code>	System error	<code>PAM_USER_UNKNOWN</code>	No account present for user
<code>PAM_AUTH_ERR</code>	Authentication failure														
<code>PAM_BUF_ERR</code>	Memory buffer error														
<code>PAM_IGNORE</code>	Ignore module, not participating in result														
<code>PAM_PERM_DENIED</code>	Permission denied														
<code>PAM_SUCCESS</code>	Successfully obtains authentication token														
<code>PAM_SYSTEM_ERR</code>	System error														
<code>PAM_USER_UNKNOWN</code>	No account present for user														
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:														
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> <tr> <td>MT Level</td> <td>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving	MT Level	MT-Safe with exceptions								
ATTRIBUTE TYPE	ATTRIBUTE VALUE														
Interface Stability	Evolving														
MT Level	MT-Safe with exceptions														
SEE ALSO	<code>pam(3PAM)</code> , <code>pam_authenticate(3PAM)</code> , <code>syslog(3C)</code> , <code>libpam(3LIB)</code> , <code>pam.conf(4)</code> , <code>nsswitch.conf(4)</code> , <code>attributes(5)</code> , <code>pam_authtok_check(5)</code> , <code>pam_authtok_get(5)</code> , <code>pam_authtok_store(5)</code> , <code>pam_dhkeys(5)</code> , <code>pam_passwd_auth(5)</code> , <code>pam_unix(5)</code> , <code>pam_unix_account(5)</code> , <code>pam_unix_session(5)</code>														
NOTES	The interfaces in <code>libpam(3LIB)</code> are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.														

pam_unix_auth(5)

The pam_unix(5) module might not be supported in a future release. Similar functionality is provided by pam_authok_check(5), pam_authok_get(5), pam_authok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix_account(5), pam_unix_auth(5), and pam_unix_session(5).

NAME	pam_unix_session – session management PAM module for UNIX						
SYNOPSIS	pam_unix_session.so.1						
DESCRIPTION	<p>pam_unix_session provides functions to initiate pam_sm_open_session(3PAM) and to terminate pam_sm_close_session(3PAM).</p> <p>pam_open_session updates the /var/adm/lastlog file. The account management module reads this file to determine the previous time the user logged in. pam_sm_close_session is a null function.</p> <p>The following options can be passed to the module:</p> <p>debug syslog(3C) debugging information at the LOG_DEBUG level</p> <p>nowarn Turn off warning messages</p>						
ERRORS	<p>The following values are returned:</p> <p>PAM_SUCCESS Successful completion</p> <p>PAM_SESSION_ERR Can not make or remove the entry for the specified session</p> <p>PAM_USER_UNKNOWN No account is present for <i>user</i></p>						
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> <tr> <td>MT Level</td> <td>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving	MT Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
MT Level	MT-Safe with exceptions						
SEE ALSO	pam(3PAM), pam_authenticate(3PAM), syslog(3C), libpam(3LIB), pam.conf(4), nsswitch.conf(4), attributes(5), pam_authok_check(5), pam_authok_get(5), pam_authok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix(5), pam_unix_account(5), pam_unix_auth(5),						
NOTES	<p>The interfaces in libpam(3LIB) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.</p> <p>The pam_unix(5) module might not be supported in a future release. Similar functionality is provided by pam_authok_check(5), pam_authok_get(5), pam_authok_store(5), pam_dhkeys(5), pam_passwd_auth(5), pam_unix_account(5), pam_unix_auth(5), and pam_unix_session(5).</p>						

prof(5)

NAME	prof – profile within a function
SYNOPSIS	<pre>#define MARK #include <prof.h> void MARK (<i>name</i>) ;</pre>
DESCRIPTION	<p>MARK introduces a mark called <i>name</i> that is treated the same as a function entry point. Execution of the mark adds to a counter for that mark, and program-counter time spent is accounted to the immediately preceding mark or to the function if there are no preceding marks within the active function.</p> <p><i>name</i> may be any combination of letters, numbers, or underscores. Each <i>name</i> in a single compilation must be unique, but may be the same as any ordinary program symbol.</p> <p>For marks to be effective, the symbol MARK must be defined before the header <code>prof.h</code> is included, either by a preprocessor directive as in the synopsis, or by a command line argument:</p> <pre>cc -p -DMARK work.c</pre> <p>If MARK is not defined, the <code>MARK(<i>name</i>)</code> statements may be left in the source files containing them and are ignored. <code>prof -g</code> must be used to get information on all labels.</p>
EXAMPLES	<p>In this example, marks can be used to determine how much time is spent in each loop. Unless this example is compiled with MARK defined on the command line, the marks are ignored.</p> <pre>#include <prof.h> work() { int i, j; . . . MARK(loop1); for (i = 0; i < 2000; i++) { . . . } MARK(loop2); for (j = 0; j < 2000; j++) { . . . } }</pre>
SEE ALSO	profil(2), monitor(3C)

NAME	rbac – role-based access control
DESCRIPTION	<p>The addition of role-based access control (RBAC) to the Solaris operating environment gives developers the opportunity to deliver fine-grained security in new and modified applications. RBAC is an alternative to the all-or-nothing security model of traditional superuser-based systems. With RBAC, an administrator can assign privileged functions to specific user accounts (or special accounts called roles).</p> <p>There are two ways to give applications privileges:</p> <ol style="list-style-type: none"> 1. Administrators can assign special attributes such as setUID to application binaries (executable files). 2. Administrators can assign special attributes such as setUID to applications using execution profiles. <p>Special attribute assignment along with the theory behind RBAC is discussed in detail in “Role Based Access Control” chapter of the <i>System Administration Guide: Security Services</i>. This chapter describes what authorizations are and how to code for them.</p>
Authorizations	<p>An authorization is a unique string that represents a user’s right to perform some operation or class of operations. Authorization definitions are stored in a database called <code>auth_attr(4)</code>. For programming authorization checks, only the authorization name is significant.</p> <p>Some typical values in an <code>auth_attr</code> database are shown below.</p> <pre>solaris.jobs.:::Cron and At Jobs::help=JobHeader.html solaris.jobs.grant:::Delegate Cron & At Administration::help=JobsGrant.html solaris.jobs.admin:::Manage All Jobs::help=AuthJobsAdmin.html solaris.jobs.user:::Cron & At User::help=JobsUser.html</pre> <p>Authorization name strings ending with the <code>grant</code> suffix are special authorizations that give a user the ability to delegate authorizations with the same prefix and functional area to other users.</p>
Creating Authorization Checks	<p>To check authorizations, use the <code>chkauthattr(3SECDB)</code> library function, which verifies whether or not a user has a given authorization. The synopsis is:</p> <pre>int chkauthattr(const char *authname, const char *username);</pre> <p>The <code>chkauthattr()</code> function checks the <code>policy.conf(4)</code>, <code>user_attr(4)</code>, and <code>prof_attr(4)</code> databases in order for a match to the given authorization.</p> <p>If you are modifying existing code that tests for root UID, you should find the test in the code and replace it with the <code>chkauthattr()</code> function. A typical root UID check is shown in the first code segment below. An authorization check replacing it is shown in the second code segment; it uses the <code>solaris.jobs.admin</code> authorization and a variable called <code>real_login</code> representing the user.</p> <p>EXAMPLE 1 Standard root check</p> <pre>ruid = getuid();</pre>

EXAMPLE 1 Standard root check *(Continued)*

```

if ((eflag || lflag || rflag) && argc == 1) {
    if ((pwp = getpwnam(*argv)) == NULL)
        crabort(INVALIDUSER);

    if (ruid != 0) {
        if (pwp->pw_uid != ruid)
            crabort(NOTROOT);
        else
            pp = getuser(ruid);
    } else
        pp = *argv++;
} else {

```

EXAMPLE 2 Authorization check

```

ruid = getuid();
if ((pwp = getpwuid(ruid)) == NULL)
    crabort(INVALIDUSER);

strcpy(real_login, pwp->pw_name);

if ((eflag || lflag || rflag) && argc == 1) {
    if ((pwp = getpwnam(*argv)) == NULL)
        crabort(INVALIDUSER);

    if (!chkauthattr("solaris.jobs.admin", real_login)) {
        if (pwp->pw_uid != ruid)
            crabort(NOTROOT);
        else
            pp = getuser(ruid);
    } else
        pp = *argv++;
} else {

```

For new applications, find an appropriate location for the test and use `chkauthattr()` as shown above. Typically the authorization check makes an access decision based on the identity of the calling user to determine if a privileged action (for example, a system call) should be taken on behalf of that user.

Applications that perform a test to restrict who can perform their security-relevant functionality are generally `setuid` to root. Programs that were written prior to RBAC and that are only available to the root user may not have such checks. In most cases, the kernel requires an effective user ID of root to override policy enforcement. Therefore, authorization checking is most useful in programs that are `setuid` to root.

For instance, if you want to write a program that allows authorized users to set the system date, the command must be run with an effective user ID of root. Typically, this means that the file modes for the file would be `-rwsr-xr-x` with root ownership.

Use caution, though, when making programs `setuid` to root. For example, the effective UID should be set to the real UID as early as possible in the program's initialization function. The effective UID can then be set back to root after the authorization check is performed and before the system call is made. On return from the system call, the effective UID should be set back to the real UID again to adhere to the principle of least privilege.

Another consideration is that `LD_LIBRARY` path is ignored for `setuid` programs (see SECURITY section in `ld.so.1(1)`) and that shell scripts must be modified to work properly when the effective and real UIDs are different. For example, the `-p` flag in Bourne shell is required to avoid resetting the effective UID back to the real UID.

Using an effective UID of root instead of the real UID requires extra care when writing shell scripts. For example, many shell scripts check to see if the user is root before executing their functionality. With RBAC, these shell scripts may be running with the effective UID of root and with a real UID of a user or role. Thus, the shell script should check `euid` instead of `uid`. For example,

```
WHO=`id | cut -f1 -d" "`
if [ ! "$WHO" = "uid=0(root)" ]
then
    echo "$PROG: ERROR: you must be super-user to run this script."
    exit 1
fi
```

should be changed to

```
WHO=`/usr/xpg4/bin/id -n -u`
if [ ! "$WHO" = "root" ]
then
    echo "$PROG: ERROR: you are not authorized to run this script."
    exit 1
fi
```

Authorizations can be explicitly checked in shell scripts by checking the output of the `auths(1)` utility. For example,

```
for auth in `auths | tr , " "` NOTFOUND
do
    [ "$auth" = "solaris.date" ] && break          # authorization found
done

if [ "$auth" != "solaris.date" ]
then
    echo >&2 "$PROG: ERROR: you are not authorized to set the date"
    exit 1
fi
```

SEE ALSO `ld.so.1(1)`, `chkauthattr(3SECDB)`, `auth_attr(4)`, `policy.conf(4)`, `prof_attr(4)`, `user_attr(4)`

System Administration Guide: Security Services

regex(5)

NAME	regex – internationalized basic and extended regular expression matching
DESCRIPTION	<p>Regular Expressions (REs) provide a mechanism to select specific strings from a set of character strings. The Internationalized Regular Expressions described below differ from the Simple Regular Expressions described on the <code>regex(5)</code> manual page in the following ways:</p> <ul style="list-style-type: none">■ both Basic and Extended Regular Expressions are supported■ the Internationalization features—character class, equivalence class, and multi-character collation—are supported. <p>The Basic Regular Expression (BRE) notation and construction rules described in the <code>BASIC REGULAR EXPRESSIONS</code> section apply to most utilities supporting regular expressions. Some utilities, instead, support the Extended Regular Expressions (ERE) described in the <code>EXTENDED REGULAR EXPRESSIONS</code> section; any exceptions for both cases are noted in the descriptions of the specific utilities using regular expressions. Both BREs and ERAs are supported by the Regular Expression Matching interfaces <code>regcomp(3C)</code> and <code>regex(3C)</code>.</p>
BASIC REGULAR EXPRESSIONS	
BREs Matching a Single Character	<p>A BRE ordinary character, a special character preceded by a backslash, or a period matches a single character. A bracket expression matches a single character or a single collating element. See <code>RE Bracket Expression</code>, below.</p>
BRE Ordinary Characters	<p>An ordinary character is a BRE that matches itself: any character in the supported character set, except for the BRE special characters listed in <code>BRE Special Characters</code>, below.</p> <p>The interpretation of an ordinary character preceded by a backslash (<code>\</code>) is undefined, except for:</p> <ol style="list-style-type: none">1. the characters <code>), (, {,</code> and <code>}</code>2. the digits 1 to 9 inclusive (see <code>BREs Matching Multiple Characters</code>, below)3. a character inside a bracket expression.
BRE Special Characters	<p>A BRE <i>special character</i> has special properties in certain contexts. Outside those contexts, or when preceded by a backslash, such a character will be a BRE that matches the special character itself. The BRE special characters and the contexts in which they have their special meaning are:</p> <p><code>.</code> <code>[</code> <code>\</code> The period, left-bracket, and backslash are special except when used in a bracket expression (see <code>RE Bracket Expression</code>, below). An expression containing a <code>[</code> that is not preceded by a backslash and is not part of a bracket expression produces undefined results.</p> <p><code>*</code> The asterisk is special except when used:</p> <ul style="list-style-type: none">■ in a bracket expression■ as the first character of an entire BRE (after an initial <code>^</code>, if any)

	<ul style="list-style-type: none"> ■ as the first character of a subexpression (after an initial ^, if any); see BREs Matching Multiple Characters, below. <p>^ The circumflex is special when used:</p> <ul style="list-style-type: none"> ■ as an anchor (see BRE Expression Anchoring, below). ■ as the first character of a bracket expression (see RE Bracket Expression, below). <p>\$ The dollar sign is special when used as an anchor.</p>
Periods in BREs	<p>A period (.), when used outside a bracket expression, is a BRE that matches any character in the supported character set except NUL.</p>
RE Bracket Expression	<p>A bracket expression (an expression enclosed in square brackets, []) is an RE that matches a single collating element contained in the non-empty set of collating elements represented by the bracket expression.</p> <p>The following rules and definitions apply to bracket expressions:</p> <ol style="list-style-type: none"> 1. A <i>bracket expression</i> is either a matching list expression or a non-matching list expression. It consists of one or more expressions: collating elements, collating symbols, equivalence classes, character classes, or range expressions (see rule 7 below). Portable applications must not use range expressions, even though all implementations support them. The right-bracket (]) loses its special meaning and represents itself in a bracket expression if it occurs first in the list (after an initial circumflex (^), if any). Otherwise, it terminates the bracket expression, unless it appears in a collating symbol (such as [.]]) or is the ending right-bracket for a collating symbol, equivalence class, or character class. The special characters: <pre> . * [\ </pre> (period, asterisk, left-bracket and backslash, respectively) lose their special meaning within a bracket expression. The character sequences: <pre> [. [= [: </pre> (left-bracket followed by a period, equals-sign, or colon) are special inside a bracket expression and are used to delimit collating symbols, equivalence class expressions, and character class expressions. These symbols must be followed by a valid expression and the matching terminating sequence .], [=] or :], as described in the following items. 2. A <i>matching list</i> expression specifies a list that matches any one of the expressions represented in the list. The first character in the list must not be the circumflex. For example, [abc] is an RE that matches any of the characters a, b or c. 3. A <i>non-matching list</i> expression begins with a circumflex (^), and specifies a list that matches any character or collating element except for the expressions represented in the list after the leading circumflex. For example, [^abc] is an RE that matches any character or collating element except the characters a, b, or c. The circumflex will have this special meaning only when it occurs first in the list, immediately following the left-bracket.

4. A *collating symbol* is a collating element enclosed within bracket-period ([.]) delimiters. Multi-character collating elements must be represented as collating symbols when it is necessary to distinguish them from a list of the individual characters that make up the multi-character collating element. For example, if the string `ch` is a collating element in the current collation sequence with the associated collating symbol `<ch>`, the expression `[[.ch.]]` will be treated as an RE matching the character sequence `ch`, while `[ch]` will be treated as an RE matching `c` or `h`. Collating symbols will be recognized only inside bracket expressions. This implies that the RE `[[. ch.]] * c` matches the first to fifth character in the string `chchch`. If the string is not a collating element in the current collating sequence definition, or if the collating element has no characters associated with it, the symbol will be treated as an invalid expression.
5. An *equivalence class expression* represents the set of collating elements belonging to an equivalence class. Only primary equivalence classes will be recognised. The class is expressed by enclosing any one of the collating elements in the equivalence class within bracket-equal ([=]) delimiters. For example, if `a`, and `b` belong to the same equivalence class, then `[[=a]=b]`, `[[=]b]` and `[[=]b]` will each be equivalent to `[ab]`. If the collating element does not belong to an equivalence class, the equivalence class expression will be treated as a *collating symbol*.
6. A *character class expression* represents the set of characters belonging to a character class, as defined in the `LC_CTYPE` category in the current locale. All character classes specified in the current locale will be recognized. A character class expression is expressed as a character class name enclosed within bracket-colon ([:]) delimiters.

The following character class expressions are supported in all locales:

<code>[:alnum:]</code>	<code>[:cntrl:]</code>	<code>[:lower:]</code>	<code>[:space:]</code>
<code>[:alpha:]</code>	<code>[:digit:]</code>	<code>[:print:]</code>	<code>[:upper:]</code>
<code>[:blank:]</code>	<code>[:graph:]</code>	<code>[:punct:]</code>	<code>[:xdigit:]</code>

In addition, character class expressions of the form:

`[:name:]`

are recognized in those locales where the *name* keyword has been given a `charclass` definition in the `LC_CTYPE` category.

7. A *range expression* represents the set of collating elements that fall between two elements in the current collation sequence, inclusively. It is expressed as the starting point and the ending point separated by a hyphen (-).
 Range expressions must not be used in portable applications because their behavior is dependent on the collating sequence. Ranges will be treated according to the current collating sequence, and include such characters that fall within the range based on that collating sequence, regardless of character values. This, however, means that the interpretation will differ depending on collating sequence. If, for instance, one collating sequence defines `a` as a variant of `A`, while another

defines it as a letter following `z`, then the expression `[-z]` is valid in the first language and invalid in the second.

In the following, all examples assume the collation sequence specified for the POSIX locale, unless another collation sequence is specifically defined.

The starting range point and the ending range point must be a collating element or collating symbol. An equivalence class expression used as a starting or ending point of a range expression produces unspecified results. An equivalence class can be used portably within a bracket expression, but only outside the range. For example, the unspecified expression `[[=e]-f]` should be given as `[[=e]e-f]`. The ending range point must collate equal to or higher than the starting range point; otherwise, the expression will be treated as invalid. The order used is the order in which the collating elements are specified in the current collation definition.

One-to-many mappings (see `locale(5)`) will not be performed. For example, assuming that the character `eszet` is placed in the collation sequence after `r` and `s`, but before `t`, and that it maps to the sequence `ss` for collation purposes, then the expression `[r-s]` matches only `r` and `s`, but the expression `[s-t]` matches `s`, `beta`, or `t`.

The interpretation of range expressions where the ending range point is also the starting range point of a subsequent range expression (for instance `[a-m-o]`) is undefined.

The hyphen character will be treated as itself if it occurs first (after an initial `^`, if any) or last in the list, or as an ending range point in a range expression. As examples, the expressions `[-ac]` and `[ac-]` are equivalent and match any of the characters `a`, `c`, or `-`; `^[^ac]` and `^[ac-]` are equivalent and match any characters except `a`, `c`, or `-`; the expression `[%- -]` matches any of the characters between `%` and `-` inclusive; the expression `[- -@]` matches any of the characters between `-` and `@` inclusive; and the expression `[a- -@]` is invalid, because the letter `a` follows the symbol `-` in the POSIX locale. To use a hyphen as the starting range point, it must either come first in the bracket expression or be specified as a collating symbol, for example: `[[[-.-]-0]`, which matches either a right bracket or any character or collating element that collates between hyphen and `0`, inclusive.

If a bracket expression must specify both `-` and `]`, the `]` must be placed first (after the `^`, if any) and the `-` last within the bracket expression.

Note: Latin-1 characters such as ``` or `^` are not printable in some locales, for example, the `ja` locale.

BREs Matching Multiple Characters

The following rules can be used to construct BREs matching multiple characters from BREs matching a single character:

1. The concatenation of BREs matches the concatenation of the strings matched by each component of the BRE.
2. A *subexpression* can be defined within a BRE by enclosing it between the character pairs `\(` and `\)`. Such a subexpression matches whatever it would have matched without the `\(` and `\)`, except that anchoring within subexpressions is optional behavior; see *BRE Expression Anchoring*, below. Subexpressions can be arbitrarily nested.

3. The *back-reference* expression `\n` matches the same (possibly empty) string of characters as was matched by a subexpression enclosed between `\(` and `\)` preceding the `\n`. The character `n` must be a digit from 1 to 9 inclusive, *n*th subexpression (the one that begins with the *n*th `\(` and ends with the corresponding paired `\)`). The expression is invalid if less than *n* subexpressions precede the `\n`. For example, the expression `^(.*)\1$` matches a line consisting of two adjacent appearances of the same string, and the expression `\(a\) * \1` fails to match `a`. The limit of nine back-references to subexpressions in the RE is based on the use of a single digit identifier. This does not imply that only nine subexpressions are allowed in REs. The following is a valid BRE with ten subexpressions:

```
\(\( (ab\) * c\) * d\) \(ef\) * \(gh\) \{2\} \(ij\) * \(kl\) * \(mn\) * \(op\) * \(qr\) *
```

4. When a BRE matching a single character, a subexpression or a back-reference is followed by the special character asterisk (`*`), together with that asterisk it matches what zero or more consecutive occurrences of the BRE would match. For example, `[ab] *` and `[ab] [ab]` are equivalent when matching the string `ab`.
5. When a BRE matching a single character, a subexpression, or a back-reference is followed by an *interval expression* of the format `\{m\}`, `\{m,\}` or `\{m,n\}`, together with that interval expression it matches what repeated consecutive occurrences of the BRE would match. The values of *m* and *n* will be decimal integers in the range $0 \leq m \leq n \leq \{RE_DUP_MAX\}$, where *m* specifies the exact or minimum number of occurrences and *n* specifies the maximum number of occurrences. The expression `\{m\}` matches exactly *m* occurrences of the preceding BRE, `\{m,\}` matches at least *m* occurrences and `\{m,n\}` matches any number of occurrences between *m* and *n*, inclusive.

For example, in the string `abababcccccccd`, the BRE `c\{3\}` is matched by characters seven to nine, the BRE `\(ab\) \{4,\}` is not matched at all and the BRE `c\{1,3\}d` is matched by characters ten to thirteen.

The behavior of multiple adjacent duplication symbols (`*` and intervals) produces undefined results.

BRE Precedence

The order of precedence is as shown in the following table:

BRE Precedence (from high to low)	
collation-related bracket symbols	<code>[= =] [: :] [. .]</code>
escaped characters	<code>\<special character></code>
bracket expression	<code>[]</code>
subexpressions/back-references	<code>\(\) \n</code>
single-character-BRE duplication	<code>* \{m,n\}</code>
concatenation	

anchoring	^ \$
-----------	------

BRE Expression Anchoring

A BRE can be limited to matching strings that begin or end a line; this is called *anchoring*. The circumflex and dollar sign special characters will be considered BRE anchors in the following contexts:

1. A circumflex (^) is an anchor when used as the first character of an entire BRE. The implementation may treat circumflex as an anchor when used as the first character of a subexpression. The circumflex will anchor the expression to the beginning of a string; only sequences starting at the first character of a string will be matched by the BRE. For example, the BRE ^ab matches ab in the string abcdef , but fails to match in the string cdefab. A portable BRE must escape a leading circumflex in a subexpression to match a literal circumflex.
2. A dollar sign (\$) is an anchor when used as the last character of an entire BRE. The implementation may treat a dollar sign as an anchor when used as the last character of a subexpression. The dollar sign will anchor the expression to the end of the string being matched; the dollar sign can be said to match the end-of-string following the last character.
3. A BRE anchored by both ^ and \$ matches only an entire string. For example, the BRE ^abcdef\$ matches strings consisting only of abcdef.
4. ^ and \$ are not special in subexpressions.

Note: The Solaris implementation does not support anchoring in BRE subexpressions.

EXTENDED REGULAR EXPRESSIONS

The rules specified for BREs apply to Extended Regular Expressions (EREs) with the following exceptions:

- The characters | , + , and ? have special meaning, as defined below.
- The { and } characters, when used as the duplication operator, are not preceded by backslashes. The constructs \{ and \} simply match the characters { and } , respectively.
- The back reference operator is not supported.
- Anchoring (^\$) is supported in subexpressions.

EREs Matching a Single Character

An ERE ordinary character, a special character preceded by a backslash, or a period matches a single character. A bracket expression matches a single character or a single collating element. An *ERE matching a single character* enclosed in parentheses matches the same as the ERE without parentheses would have matched.

ERE Ordinary Characters

An *ordinary character* is an ERE that matches itself. An ordinary character is any character in the supported character set, except for the ERE special characters listed in **ERE Special Characters** below. The interpretation of an ordinary character preceded by a backslash (\) is undefined.

regex(5)

ERE Special Characters	<p>An <i>ERE special character</i> has special properties in certain contexts. Outside those contexts, or when preceded by a backslash, such a character is an ERE that matches the special character itself. The extended regular expression special characters and the contexts in which they have their special meaning are:</p> <p>. [\ (</p> <p>)</p> <p>* + ? {</p> <p> </p> <p>^</p> <p>\$</p>
	<p>The period, left-bracket, backslash, and left-parenthesis are special except when used in a bracket expression (see <i>RE Bracket Expression</i>, above). Outside a bracket expression, a left-parenthesis immediately followed by a right-parenthesis produces undefined results.</p> <p>The right-parenthesis is special when matched with a preceding left-parenthesis, both outside a bracket expression.</p> <p>The asterisk, plus-sign, question-mark, and left-brace are special except when used in a bracket expression (see <i>RE Bracket Expression</i>, above). Any of the following uses produce undefined results:</p> <ul style="list-style-type: none">■ if these characters appear first in an ERE, or immediately following a vertical-line, circumflex or left-parenthesis■ if a left-brace is not part of a valid interval expression. <p>The vertical-line is special except when used in a bracket expression (see <i>RE Bracket Expression</i>, above). A vertical-line appearing first or last in an ERE, or immediately following a vertical-line or a left-parenthesis, or immediately preceding a right-parenthesis, produces undefined results.</p> <p>The circumflex is special when used:</p> <ul style="list-style-type: none">■ as an anchor (see <i>ERE Expression Anchoring</i>, below).■ as the first character of a bracket expression (see <i>RE Bracket Expression</i>, above). <p>The dollar sign is special when used as an anchor.</p>
Periods in EREs	<p>A period (.), when used outside a bracket expression, is an ERE that matches any character in the supported character set except NUL.</p>
ERE Bracket Expression	<p>The rules for ERE Bracket Expressions are the same as for Basic Regular Expressions; see <i>RE Bracket Expression</i>, above).</p>
EREs Matching Multiple Characters	<p>The following rules will be used to construct EREs matching multiple characters from EREs matching a single character:</p> <ol style="list-style-type: none">1. A <i>concatenation of EREs</i> matches the concatenation of the character sequences matched by each component of the ERE. A concatenation of EREs enclosed in parentheses matches whatever the concatenation without the parentheses matches. For example, both the ERE <i>cd</i> and the ERE <i>(cd)</i> are matched by the third and fourth character of the string <i>abcdefabcdef</i>.

2. When an ERE matching a single character or an ERE enclosed in parentheses is followed by the special character plus-sign (+), together with that plus-sign it matches what one or more consecutive occurrences of the ERE would match. For example, the ERE `b+(bc)` matches the fourth to seventh characters in the string `acabbbbcde`; `[ab]+` and `[ab][ab]*` are equivalent.
3. When an ERE matching a single character or an ERE enclosed in parentheses is followed by the special character asterisk (*), together with that asterisk it matches what zero or more consecutive occurrences of the ERE would match. For example, the ERE `b*c` matches the first character in the string `cabbbbcde`, and the ERE `b*cd` matches the third to seventh characters in the string `cabbbbcdebbbbbbbcdbc`. And, `[ab]*` and `[ab][ab]` are equivalent when matching the string `ab`.
4. When an ERE matching a single character or an ERE enclosed in parentheses is followed by the special character question-mark (?), together with that question-mark it matches what zero or one consecutive occurrences of the ERE would match. For example, the ERE `b?c` matches the second character in the string `acabbbbcde`.
5. When an ERE matching a single character or an ERE enclosed in parentheses is followed by an *interval expression* of the format `{m}`, `{m,}` or `{m,n}`, together with that interval expression it matches what repeated consecutive occurrences of the ERE would match. The values of *m* and *n* will be decimal integers in the range $0 \leq m \leq n \leq \{RE_DUP_MAX\}$, where *m* specifies the exact or minimum number of occurrences and *n* specifies the maximum number of occurrences. The expression `{m}` matches exactly *m* occurrences of the preceding ERE, `{m,}` matches at least *m* occurrences and `{m,n}` matches any number of occurrences between *m* and *n*, inclusive.

For example, in the string `abababcccccccd` the ERE `c{3}` is matched by characters seven to nine and the ERE `(ab){2,}` is matched by characters one to six.

The behavior of multiple adjacent duplication symbols (+, *, ? and intervals) produces undefined results.

ERE Alternation Two EREs separated by the special character vertical-line (|) match a string that is matched by either. For example, the ERE `a((bc)|d)` matches the string `abc` and the string `ad`. Single characters, or expressions matching single characters, separated by the vertical bar and enclosed in parentheses, will be treated as an ERE matching a single character.

ERE Precedence The order of precedence will be as shown in the following table:

ERE Precedence (from high to low)	
collation-related bracket symbols	<code>[= =] [: :] [. .]</code>
escaped characters	<code>\<special character></code>
bracket expression	<code>[]</code>

regex(5)

grouping	()
single-character-ERE duplication	* + ? {m,n}
concatenation	
anchoring	^ \$
alternation	

For example, the ERE `abba | cde` matches either the string `abba` or the string `cde` (rather than the string `abbade` or `abbcde`, because concatenation has a higher order of precedence than alternation).

ERE Expression Anchoring

An ERE can be limited to matching strings that begin or end a line; this is called *anchoring*. The circumflex and dollar sign special characters are considered ERE anchors when used anywhere outside a bracket expression. This has the following effects:

1. A circumflex (^) outside a bracket expression anchors the expression or subexpression it begins to the beginning of a string; such an expression or subexpression can match only a sequence starting at the first character of a string. For example, the EREs `^ab` and `(^ab)` match `ab` in the string `abcdef`, but fail to match in the string `cdefab`, and the ERE `a^b` is valid, but can never match because the `a` prevents the expression `^b` from matching starting at the first character.
2. A dollar sign (\$) outside a bracket expression anchors the expression or subexpression it ends to the end of a string; such an expression or subexpression can match only a sequence ending at the last character of a string. For example, the EREs `ef$` and `(ef$)` match `ef` in the string `abcdef`, but fail to match in the string `cdefab`, and the ERE `e$f` is valid, but can never match because the `f` prevents the expression `e$` from matching ending at the last character.

SEE ALSO

`localedef(1)`, `regcomp(3C)`, `attributes(5)`, `environ(5)`, `locale(5)`, `regexp(5)`

NAME	regex, compile, step, advance – simple regular expression compile and match routines
SYNOPSIS	<pre> #define INIT <i>declarations</i> #define GETC(void) <i>getc code</i> #define PEEKC(void) <i>peekc code</i> #define UNGETC(void) <i>ungetc code</i> #define RETURN(ptr) <i>return code</i> #define ERROR(val) <i>error code</i> extern char *loc1, *loc2, *locs; #include <regex.h> char *compile(char *instring, char *expbuf, const char *endfug, int eof); int step(const char *string, const char *expbuf); int advance(const char *string, const char *expbuf); </pre>
DESCRIPTION	<p>Regular Expressions (REs) provide a mechanism to select specific strings from a set of character strings. The Simple Regular Expressions described below differ from the Internationalized Regular Expressions described on the regex(5) manual page in the following ways:</p> <ul style="list-style-type: none"> ■ only Basic Regular Expressions are supported ■ the Internationalization features—character class, equivalence class, and multi-character collation—are not supported. <p>The functions <code>step()</code>, <code>advance()</code>, and <code>compile()</code> are general purpose regular expression matching routines to be used in programs that perform regular expression matching. These functions are defined by the <code><regex.h></code> header.</p> <p>The functions <code>step()</code> and <code>advance()</code> do pattern matching given a character string and a compiled regular expression as input.</p> <p>The function <code>compile()</code> takes as input a regular expression as defined below and produces a compiled expression that can be used with <code>step()</code> or <code>advance()</code>.</p>
Basic Regular Expressions	<p>A regular expression specifies a set of character strings. A member of this set of strings is said to be matched by the regular expression. Some characters have special meaning when used in a regular expression; other characters stand for themselves.</p> <p>The following <i>one-character REs</i> match a <i>single</i> character:</p> <ol style="list-style-type: none"> 1.1 An ordinary character (<i>not</i> one of those discussed in 1.2 below) is a one-character RE that matches itself. 1.2 A backslash (<code>\</code>) followed by any special character is a one-character RE that matches the special character itself. The special characters are:

- a. . , * , [, and \ (period, asterisk, left square bracket, and backslash, respectively), which are always special, *except* when they appear within square brackets ([]); see 1.4 below).
- b. ^ (caret or circumflex), which is special at the *beginning* of an *entire* RE (see 4.1 and 4.3 below), or when it immediately follows the left of a pair of square brackets ([]) (see 1.4 below).
- c. \$ (dollar sign), which is special at the end of an *entire* RE (see 4.2 below).
- d. The character used to bound (that is, delimit) an entire RE, which is special for that RE (for example, see how slash (/) is used in the g command, below.)

1.3 A period (.) is a one-character RE that matches any character except new-line.

1.4 A non-empty string of characters enclosed in square brackets ([]) is a one-character RE that matches *any one* character in that string. If, however, the first character of the string is a circumflex (^), the one-character RE matches any character *except* new-line and the remaining characters in the string. The ^ has this special meaning *only* if it occurs first in the string. The minus (-) may be used to indicate a range of consecutive characters; for example, [0-9] is equivalent to [0123456789]. The - loses this special meaning if it occurs first (after an initial ^, if any) or last in the string. The right square bracket (]) does not terminate such a string when it is the first character within it (after an initial ^, if any); for example, []a-f] matches either a right square bracket (]) or one of the ASCII letters a through f inclusive. The four characters listed in 1.2.a above stand for themselves within such a string of characters.

The following rules may be used to construct REs from one-character REs:

- 2.1 A one-character RE is a RE that matches whatever the one-character RE matches.
- 2.2 A one-character RE followed by an asterisk (*) is a RE that matches 0 or more occurrences of the one-character RE. If there is any choice, the longest leftmost string that permits a match is chosen.
- 2.3 A one-character RE followed by \{m\}, \{m,\}, or \{m,n\} is a RE that matches a *range* of occurrences of the one-character RE. The values of *m* and *n* must be non-negative integers less than 256; \{m\} matches *exactly* *m* occurrences; \{m,\} matches *at least* *m* occurrences; \{m,n\} matches *any number* of occurrences *between* *m* and *n* inclusive. Whenever a choice exists, the RE matches as many occurrences as possible.
- 2.4 The concatenation of REs is a RE that matches the concatenation of the strings matched by each component of the RE.

- 2.5 A RE enclosed between the character sequences \ (and \) is a RE that matches whatever the unadorned RE matches.
- 2.6 The expression \ n matches the same string of characters as was matched by an expression enclosed between \ (and \) *earlier* in the same RE. Here *n* is a digit; the sub-expression specified is that beginning with the *n*-th occurrence of \ (counting from the left. For example, the expression $\ (\ . \ * \) \ 1 \ \$$ matches a line consisting of two repeated appearances of the same string.

An RE may be constrained to match words.

- 3.1 \ < constrains a RE to match the beginning of a string or to follow a character that is not a digit, underscore, or letter. The first character matching the RE must be a digit, underscore, or letter.
- 3.2 \ > constrains a RE to match the end of a string or to precede a character that is not a digit, underscore, or letter.

An *entire RE* may be constrained to match only an initial segment or final segment of a line (or both).

- 4.1 A circumflex (^) at the beginning of an entire RE constrains that RE to match an *initial* segment of a line.
- 4.2 A dollar sign (\$) at the end of an entire RE constrains that RE to match a *final* segment of a line.
- 4.3 The construction $\ ^entire\ RE\ \$$ constrains the entire RE to match the entire line.

The null RE (for example, //) is equivalent to the last RE encountered.

Addressing with REs

Addresses are constructed as follows:

1. The character "." addresses the current line.
2. The character "\$" addresses the last line of the buffer.
3. A decimal number *n* addresses the *n*-th line of the buffer.
4. 'x addresses the line marked with the mark name character *x*, which must be an ASCII lower-case letter (a-z). Lines are marked with the *k* command described below.
5. A RE enclosed by slashes (/) addresses the first line found by searching *forward* from the line *following* the current line toward the end of the buffer and stopping at the first line containing a string matching the RE. If necessary, the search wraps around to the beginning of the buffer and continues up to and including the current line, so that the entire buffer is searched.
6. A RE enclosed in question marks (?) addresses the first line found by searching *backward* from the line *preceding* the current line toward the beginning of the buffer and stopping at the first line containing a string matching the RE. If necessary, the

regex(5)

search wraps around to the end of the buffer and continues up to and including the current line.

7. An address followed by a plus sign (+) or a minus sign (-) followed by a decimal number specifies that address plus (respectively minus) the indicated number of lines. A shorthand for .+5 is .5.
8. If an address begins with + or -, the addition or subtraction is taken with respect to the current line; for example, -5 is understood to mean .-5.
9. If an address ends with + or -, then 1 is added to or subtracted from the address, respectively. As a consequence of this rule and of Rule 8, immediately above, the address - refers to the line preceding the current line. (To maintain compatibility with earlier versions of the editor, the character ^ in addresses is entirely equivalent to -.) Moreover, trailing + and - characters have a cumulative effect, so -- refers to the current line less 2.
10. For convenience, a comma (,) stands for the address pair 1, \$, while a semicolon (;) stands for the pair ., \$.

Characters With Special Meaning

Characters that have special meaning except when they appear within square brackets ([]) or are preceded by \ are: ., *, [, \ . Other special characters, such as \$ have special meaning in more restricted contexts.

The character ^ at the beginning of an expression permits a successful match only immediately after a newline, and the character \$ at the end of an expression requires a trailing newline.

Two characters have special meaning only when used within square brackets. The character - denotes a range, [c-c], unless it is just after the open bracket or before the closing bracket, [-c] or [c-] in which case it has no special meaning. When used within brackets, the character ^ has the meaning *complement of* if it immediately follows the open bracket (example: [^c]); elsewhere between brackets (example: [c^]) it stands for the ordinary character ^.

The special meaning of the \ operator can be escaped only by preceding it with another \ , for example \\ .

Macros

Programs must have the following five macros declared before the #include <regex.h> statement. These macros are used by the compile() routine. The macros GETC, PEEKC, and UNGETC operate on the regular expression given as input to compile().

GETC This macro returns the value of the next character (byte) in the regular expression pattern. Successive calls to GETC should return successive characters of the regular expression.

PEEKC This macro returns the next character (byte) in the regular expression. Immediately successive calls to PEEKC should return the same character, which should also be the next character returned by GETC.

UNGETC This macro causes the argument *c* to be returned by the next call to GETC and PEEKC. No more than one character of pushback is ever needed and this character is guaranteed to be the last character read by GETC. The return value of the macro UNGETC(*c*) is always ignored.

RETURN(*ptr*) This macro is used on normal exit of the compile() routine. The value of the argument *ptr* is a pointer to the character after the last character of the compiled regular expression. This is useful to programs which have memory allocation to manage.

ERROR(*val*) This macro is the abnormal return from the compile() routine. The argument *val* is an error number (see ERRORS below for meanings). This call should never return.

compile() The syntax of the compile() routine is as follows:

```
compile(instring, expbuf, endbuf, eof)
```

The first parameter, *instring*, is never used explicitly by the compile() routine but is useful for programs that pass down different pointers to input characters. It is sometimes used in the INIT declaration (see below). Programs which call functions to input characters or have characters in an external array can pass down a value of (char *)0 for this parameter.

The next parameter, *expbuf*, is a character pointer. It points to the place where the compiled regular expression will be placed.

The parameter *endbuf* is one more than the highest address where the compiled regular expression may be placed. If the compiled expression cannot fit in (endbuf-expbuf) bytes, a call to ERROR(50) is made.

The parameter *eof* is the character which marks the end of the regular expression. This character is usually a /.

Each program that includes the <regex.h> header file must have a #define statement for INIT. It is used for dependent declarations and initializations. Most often it is used to set a register variable to point to the beginning of the regular expression so that this register variable can be used in the declarations for GETC, PEEKC, and UNGETC. Otherwise it can be used to declare external variables that might be used by GETC, PEEKC and UNGETC. (See EXAMPLES below.)

step(), advance() The first parameter to the step() and advance() functions is a pointer to a string of characters to be checked for a match. This string should be null terminated.

The second parameter, *expbuf*, is the compiled regular expression which was obtained by a call to the function compile().

regexp(5)

The function `step()` returns non-zero if some substring of *string* matches the regular expression in *expbuf* and 0 if there is no match. If there is a match, two external character pointers are set as a side effect to the call to `step()`. The variable `loc1` points to the first character that matched the regular expression; the variable `loc2` points to the character after the last character that matches the regular expression. Thus if the regular expression matches the entire input string, `loc1` will point to the first character of *string* and `loc2` will point to the null at the end of *string*.

The function `advance()` returns non-zero if the initial substring of *string* matches the regular expression in *expbuf*. If there is a match, an external character pointer, `loc2`, is set as a side effect. The variable `loc2` points to the next character in *string* after the last character that matched.

When `advance()` encounters a `*` or `\{ \}` sequence in the regular expression, it will advance its pointer to the string to be matched as far as possible and will recursively call itself trying to match the rest of the string to the rest of the regular expression. As long as there is no match, `advance()` will back up along the string until it finds a match or reaches the point in the string that initially matched the `*` or `\{ \}`. It is sometimes desirable to stop this backing up before the initial point in the string is reached. If the external character pointer `locs` is equal to the point in the string at sometime during the backing up process, `advance()` will break out of the loop that backs up and will return zero.

The external variables `circf`, `sed`, and `nbra` are reserved.

EXAMPLES

EXAMPLE 1 The following is an example of how the regular expression macros and calls might be defined by an application program:

```
#define INIT          register char *sp = instring;
#define GETC          (*sp++)
#define PEEKC         (*sp)
#define UNGETC(c)     (--sp)
#define RETURN(*c)    return;
#define ERROR(c)      regerr
#include <regexp.h>
. . .
(void) compile(*argv, expbuf, &expbuf[ESIZE], '\0');
. . .
if (step(linebuf, expbuf))
    succeed;
```

DIAGNOSTICS

The function `compile()` uses the macro `RETURN` on success and the macro `ERROR` on failure (see above). The functions `step()` and `advance()` return non-zero on a successful match and zero if there is no match. Errors are:

- 11 range endpoint too large.
- 16 bad number.
- 25 `\digit` out of range.
- 36 illegal or missing delimiter.

- 41 no remembered search string.
- 42 \ (\) imbalance.
- 43 too many \ (.
- 44 more than 2 numbers given in \{ \}.
- 45 } expected after \.
- 46 first number exceeds second in \{ \}.
- 49 [] imbalance.
- 50 regular expression overflow.

SEE ALSO regex(5)

SEAM(5)

NAME	SEAM – overview of Sun Enterprise Authentication Mechanism				
DESCRIPTION	<p>SEAM (Sun Enterprise Authentication Mechanism) authenticates clients in a network environment, allowing for secure transactions. (A client may be a user or a network service) SEAM validates the identity of a client and the authenticity of transferred data. SEAM is a <i>single-sign-on</i> system, meaning that a user needs to provide a password only at the beginning of a session. SEAM is based on the Kerberos™ system developed at MIT, and is compatible with Kerberos V5 systems over heterogeneous networks.</p> <p>SEAM works by granting clients <i>tickets</i>, which uniquely identify a client, and which have a finite lifetime. A client possessing a ticket is automatically validated for network services for which it is entitled; for example, a user with a valid SEAM ticket may rlogin into another machine running SEAM without having to identify itself. Because each client has a unique ticket, its identity is guaranteed.</p> <p>To obtain tickets, a client must first initialize the SEAM session, either by using the <code>kinit(1)</code> command or a PAM module. (See <code>pam_krb5(5)</code>). <code>kinit</code> prompts for a password, and then communicates with a <i>Key Distribution Center</i> (KDC). The KDC returns a <i>Ticket-Granting Ticket</i> (TGT) and prompts for a confirmation password. If the client confirms the password, it can use the Ticket-Granting Ticket to obtain tickets for specific network services. Because tickets are granted transparently, the user need not worry about their management. Current tickets may be viewed by using the <code>klist(1)</code> command.</p> <p>Tickets are valid according to the system <i>policy</i> set up at installation time. For example, tickets have a default lifetime for which they are valid. A policy may further dictate that privileged tickets, such as those belonging to root, have very short lifetimes. Policies may allow some defaults to be overruled; for example, a client may request a ticket with a lifetime greater or less than the default.</p> <p>Tickets can be renewed using <code>kinit</code>. Tickets are also <i>forwardable</i>, allowing you to use a ticket granted on one machine on a different host. Tickets can be destroyed by using <code>kdestroy(1)</code>. It is a good idea to include a call to <code>kdestroy</code> in your <code>.logout</code> file.</p> <p>Under SEAM, a client is referred to as a <i>principal</i>. A principal takes the following form:</p> <pre>primary/instance@REALM</pre> <table><tr><td><code>primary</code></td><td>A user, a host, or a service.</td></tr><tr><td><code>instance</code></td><td>A qualification of the primary. If the primary is a host — indicated by the keyword <code>host</code> — then the instance is the fully-qualified domain name of that host. If the primary is a user or service, then the instance is optional. Some instances, such as <code>admin</code> or <code>root</code>, are privileged.</td></tr></table>	<code>primary</code>	A user, a host, or a service.	<code>instance</code>	A qualification of the primary. If the primary is a host — indicated by the keyword <code>host</code> — then the instance is the fully-qualified domain name of that host. If the primary is a user or service, then the instance is optional. Some instances, such as <code>admin</code> or <code>root</code> , are privileged.
<code>primary</code>	A user, a host, or a service.				
<code>instance</code>	A qualification of the primary. If the primary is a host — indicated by the keyword <code>host</code> — then the instance is the fully-qualified domain name of that host. If the primary is a user or service, then the instance is optional. Some instances, such as <code>admin</code> or <code>root</code> , are privileged.				

realm The Kerberos equivalent of a domain; in fact, in most cases the realm is directly mapped to a DNS domain name. SEAM realms are given in upper-case only. For examples of principal names, see the EXAMPLES.

By taking advantage of the General Security Services API (GSS-API), SEAM offers, besides user authentication, two other types of security service: *integrity*, which authenticates the validity of transmitted data, and *privacy*, which encrypts transmitted data. Developers can take advantage of the GSS-API through the use of the RPCSEC_GSS API interface (see `rpcsec_gss(3NSL)`).

EXAMPLES **EXAMPLE 1** Examples of valid principal names

The following are examples of valid principal names:

```
joe
joe/admin
joe@ENG.ACME.COM
joe/admin@ENG.ACME.COM
rlogin/bigmachine.eng.acme.com@ENG.ACME.COM
host/bigmachine.eng.acme.com@ENG.ACME.COM
```

The first four cases are *user principals*. In the first two cases, it is assumed that the user `joe` is in the same realm as the client, so no realm is specified. Note that `joe` and `joe/admin` are different principals, even if the same user uses them; `joe/admin` has different privileges from `joe`. The fifth case is a *service principal*, while the final case is a *host principal*. The word `host` is required for host principals. With host principals, the instance is the fully qualified hostname. Note that the words `admin` and `host` are reserved keywords.

SEE ALSO `kdestroy(1)`, `kinit(1)`, `klist(1)`, `kpasswd(1)`, `krb5.conf(5)`

Sun Enterprise Authentication Mechanism Guide

NOTES If you enter your username and `kinit` responds with this message:

Principal unknown (kerberos) you haven't been registered as a SEAM user. See your system administrator or the *Sun Enterprise Authentication Mechanism Guide*.

sgml(5)

NAME	sgml, solbook – Standard Generalized Markup Language
DESCRIPTION	<p>Standard Generalized Markup Language (SGML) is the ISO standard 8879:1986 that describes a syntax for marking up documents with tags that describe the purpose of the text rather than the appearance on the page. This form of markup facilitates document interchange between different platforms and applications. SGML allows the management of information as data objects rather than text on a page.</p> <p>In an SGML document the main structural components are called <code>elements</code>. The organization and structure of a document and the meaning of elements are described in the Document Type Definition (DTD). Elements are the <code>tags</code> that identify the content. Element names may be descriptive of the content for ease of use. For example <code><para></code> for paragraphs. Elements can have <code>attributes</code> which are used to modify or refine the properties or characteristics of the element. Within the DTD a valid context for each element is defined and a framework is provided for the types of elements that constitute a compliant document.</p> <p>Another component of the DTD is <code>entities</code>. Entities are a collection of characters that can be referenced as a unit. Entities are similar to constants in a programming language such as C. They can be defined and referenced. An entity can represent one character or symbol which does not appear on a standard keyboard, a word or group of words, or an entire separate sgml marked-up file. Entities allow reuse of standard text.</p> <p>There is no single standard DTD , but the de facto standard for the computer industry is the DocBook DTD , developed and maintained by the Davenport Group. Within Sun, the SolBook DTD , which is a proper subset of DocBook DTD , is used when writing reference manual pages. The SolBook DTD contains a number of tags that are designed for the unique needs of the reference pages.</p>
SolBook Elements	Elements are defined with a hierarchical structure that gives a structure to the document. The following is a description of some of the elements from the SolBook DTD which are used for reference pages.
DOCTYPE	The first line in an SGML file that identifies the location of the DTD that is used to define the document. The <code><!DOCTYPE</code> string is what the SGML -aware <code>man(1)</code> command uses to identify that a file is formatted in SGML rather than <code>nroff(1)</code> .
RefEntry	The top layer element that contains a reference page is <code><refentry></code> . All of the text and other tags must be contained within this tag.
RefMeta	The next tag in a reference page is <code><refmeta></code> , which is a container for several other tags. They are:
	<code><refentrytitle></code> This is the title of the reference page. It is equivalent to the name of the reference page's file name, without the section number extension.
	<code><manvolnum></code> This is the section number that the reference page resides in. The contents may be a text entity reference.

<code><refmiscinfo></code>	<p>There are one or more <code><refmiscinfo></code> tags which contain <i>meta</i> information. Meta information is information about the reference page. The <code><refmiscinfo></code> tag has the <code>class</code> attribute. There are four classes that are routinely used.</p> <p><code>date</code> This is the date that the file was last modified. By consensus this date is changed only when the technical information on the page changes and not simply for an editorial change.</p> <p><code>sectdesc</code> This is the section title of the reference page; for example <code>User Commands</code>. The value of this attribute may be a text entity reference.</p> <p><code>software</code> This is the name of the software product that the topic discussed on the reference page belongs to. For example UNIX commands are part of the SunOS x.x release. The value of this attribute may be a text entity reference.</p> <p><code>arch</code> This is the architectural platform limitation of the subject discussed on the reference page. If there are no limitations the value used is <code>generic</code>. Other values are <code>sparc</code> and <code>x86</code>.</p> <p><code>copyright</code> This attribute contains the Sun Microsystems copyright. Any other copyrights that may pertain to the individual reference page file should be entered as separate <code><refmiscinfo></code> entries. The value of this attribute may be a text entity reference.</p>
RefNameDiv	<p>This tag contains the equivalent information to the <code>.TH</code> macro line in an <code>nroff(1)</code> reference page. <code><refnamediv></code> contains three tags. These tags contain the text that is before and after the <code>'-'</code> (dash) on the <code>NAME</code> line.</p> <p><code><refname></code> These are the names of the topics that are discussed in the file. There may be more than one <code><refname></code> for a page. The first <code><refname></code> must match the name of the file and the <code><refentrytitle></code>. If there are more than one <code><refname></code> tags, each is separated by a <code>' '</code> (comma). The comma is generated by the publisher of sgml files, so it should not be typed. This is referred to as <i>auto-generated</i> text.</p> <p><code><refpurpose></code> The text after the dash on the <code>NAME</code> line is contained in this tag. This is a short summary of what the object or objects described on the reference page do or are used for. The dash is also auto-generated and should not be typed in.</p> <p><code><refdescriptor></code> In some cases the <code><refentrytitle></code> is a general topic descriptor of a group of related objects that are discussed on the same page.</p>

sgml(5)

	<p>In this case the first tag after the <code><refnamediv></code> is a <code><refdescriptor></code>. The <code><refname></code> tags follow. Only one <code><refdescriptor></code> is allowed, and it should match the <code><refentrytitle></code>.</p>
RefSynopsisDiv	<p>The SYNOPSIS line of the reference page is contained by this tag. There is a <code><title></code> that usually contains an entity reference. The text is the word SYNOPSIS. There are several tags within <code><refsynopsisdiv></code> that are designed specifically for the type of synopsis that is used in the different reference page sections. The three types are:</p> <p><code><cmdsynopsis></code> Used for commands and utilities pages.</p> <p><code><funcsynopsis></code> Used for programming interface pages.</p> <p><code><synopsis></code> Used for pages that do not fall into the other two categories.</p>
RefSect1	<p>This tag is equivalent to the <code>.SH</code> nroff macro. It contains a <code><title></code> element that is the title of the reference page section. Section names are the standard names such as DESCRIPTION, OPTIONS, PARAMETERS, SEE ALSO, and others. The contents of the <code><title></code> may be a text entity reference.</p>
RefSect2	<p>This tag is equivalent to the <code>.SS</code> nroff macro. It contains a <code><title></code> element that contains the text of the sub-section heading. <code><refsect2></code> tags may also be used within a <code><refsynopsisdiv></code> as a sub-section heading for the SYNOPSIS section.</p>
Block Elements	<p>There are a number of block elements that are used for grouping text. This is a list of some of these elements.</p> <p><code><para></code> This tag is used to contain a paragraph of text.</p> <p><code><variablelist></code> This tag is used to create two column lists. For example descriptions for command options, where the first column lists the option and the second column describes the option.</p> <p><code><orderedlist></code> An list of items in a specific order.</p> <p><code><itemizedlist></code> A list of items that are marked with a character such as a bullet or a dash.</p> <p><code><literallayout></code> Formatted program output as produced by a program or command. This tag is a container for lines set off from the main text in which line breaks, tabs, and leading white space are significant.</p> <p><code><programlisting></code> A segment of program code. Line breaks and leading white space are significant.</p> <p><code><table></code> This tag contains the layout and content for tabular formatting of information. <code><table></code> has a required <code><title></code>.</p>

	<code><informatable></code>	This tag is the same as the <code><table></code> tag except the <code><title></code> is not required.
	<code><example></code>	This tag contains examples of source code or usage of commands. It contains a required <code><title></code> .
	<code><informalexample></code>	This tag is the same as the <code><example></code> tag except the <code><title></code> is not required.
Inline Elements	The inline elements are used for tagging text.	
	<code><command></code>	An executable program or the entry a user makes to execute a command.
	<code><function></code>	A subroutine in a program or external library.
	<code><literal></code>	Contains any literal string.
	<code><parameter></code>	An argument passed to a computer program by a function or routine.
	<code><inlineequation></code>	An untitled mathematical equation occurring in-line.
	<code><link></code>	A hypertext link to text within a book, in the case of the reference manual it is used to cross reference to another reference page.
	<code><olink></code>	A hypertext link used to create cross references to books other than the reference manual.
	<code><xref></code>	A cross reference to another part of the same reference page.
SEE ALSO	man(1), nroff(1), man(5)	

smartcard(5)

NAME	smartcard – overview of smart card features on Solaris
DESCRIPTION	<p>The smart card framework provides a mechanism to abstract the details of interacting with smart cards and smart cardreaders (called card terminals). The framework is based on the OpenCard Framework V1.1 (OCF) with Sun extensions to allow OCF to operate in a multi-user environment. The core OCF software protocol stack is implemented as a system service daemon. This implementation allows smart cards and card terminals to be shared cooperatively among many different clients on the system while providing access control to the smart card and card terminal resources on a per-UID basis.</p> <p>An event dispatcher is provided to inform clients of events occurring on the card and at the card terminal, such as card insertion and card removal.</p> <p>A high-level authentication mechanism is provided to allow clients to perform smart card-based authentications without requiring knowledge of specific card or reader authentication features.</p> <p>A set of applet administration tools is provided for JavaCards that support downloading Java applets (although applet build tools are not provided).</p> <p>Administration of the smart card framework is provided with the <code>smartcard(1M)</code> command line administration utility and the <code>smartcardguiadmin(1)</code> GUI administration tool.</p> <p>Support for several card terminals is provided:</p> <ul style="list-style-type: none">■ Sun External Smart Card Reader I (see <code>ocf_escr1(7D)</code>)■ Sun Internal Smart Card Reader I (see <code>ocf_iscr1(7D)</code>)■ Dallas iButton Serial Reader (see <code>ocf_ibutton(7D)</code>) <p>Additional card terminals can be supported by implementing smart card terminal interfaces in a shared library.</p> <p>Support for several smart cards is provided:</p> <ul style="list-style-type: none">■ Schlumberger Cyberflex Access JavaCard■ Schlumberger MicroPayflex■ Dallas Semiconductor Java iButton JavaCard <p>Each of the supported cards has a complete set of OCF card services that implement the necessary functionality for authentication and secure storage of data. For the two supported JavaCards, an authentication and secure data storage applet is provided that can be loaded into these cards with the supplied applet administration tools. See <code>smartcard(1M)</code>.</p> <p>A PAM smart card module is provided to allow PAM clients to use smart card-based authentication. See <code>pam_smartcard(5)</code></p>

smartcard(5)

CDE is able to use the PAM smart card module for `dtlogin` and `dtsession` authentication. CDE also uses the smart card framework event dispatcher to listen for events on the card terminal and provide corresponding visual feedback to the user.

SEE ALSO `ocfserv(1M)`, `smartcard(1M)`, `libsmartcard(3LIB)`, `pam_start(3PAM)`, `pam_smartcard(5)`, `ocf_escr1(7D)`, `ocf_ibutton(7D)`, `ocf_iscr1(7D)`, `scmi2c(7D)`

standards(5)

NAME standards, ANSI, C, C++, ISO, POSIX, POSIX.1, POSIX.2, SUS, SUSv2, SVID, SVID3, XNS, XNS4, XNS5, XPG, XPG3, XPG4, XPG4v2 – standards and specifications supported by Solaris

DESCRIPTION Solaris 9 supports IEEE Std 1003.1 and IEEE Std 1003.2, commonly known as POSIX.1 and POSIX.2, respectively. The following table lists each version of these standards with a brief description and the SunOS or Solaris release that first conformed to it.

POSIX Standard	Description	Release
POSIX.1-1988	system interfaces and headers	SunOS 4.1
POSIX.1-1990	POSIX.1-1988 update	Solaris 2.0
POSIX.1b-1993	realtime extensions	Solaris 2.4
POSIX.1c-1996	threads extensions	Solaris 2.6
POSIX.2-1992	shell and utilities	Solaris 2.5
POSIX.2a-1992	interactive shell and utilities	Solaris 2.5

Solaris 9 also supports the X/Open Common Applications Environment (CAE) Portability Guide Issue 3 (XPG3) and Issue 4 (XPG4), Single UNIX Specification (SUS, also known as XPG4v2), and Single UNIX Specification, Version 2 (SUSv2). Both XPG4 and SUS include Networking Services Issue 4 (XNS4). SUSv2 includes Networking Services Issue 5 (XNS5).

The following table lists each X/Open specification with a brief description and the SunOS or Solaris release that first conformed to it.

X/Open CAE Specification	Description	Release
XPG3	superset of POSIX.1-1988 containing utilities from SVID3	SunOS 4.1
XPG4	superset of POSIX.1-1990, POSIX.2-1992, and POSIX.2a-1992 containing extensions to POSIX standards from XPG3	Solaris 2.4
SUS (XPG4v2)	superset of XPG4 containing historical BSD interfaces widely used by common application packages	Solaris 2.6
XNS4	sockets and XTI interfaces	Solaris 2.6
SUSv2	superset of SUS extended to support POSIX.1b-1993, POSIX.1c-1996, and ISO/IEC 9899 (C Standard) Amendment 1	Solaris 7

X/Open CAE Specification	Description	Release
XNS5	superset and LP64-clean derivative of XNS4.	Solaris 7

The XNS4 specification is safe for use only in ILP32 (32-bit) environments and should not be used for LP64 (64-bit) application environments. Use XNS5, which has LP64-clean interfaces that are portable across ILP32 and LP64 environments. Solaris releases 7 through 9 support both the ILP32 and ILP64 environments.

Solaris releases 7 through 9 have been branded to conform to The Open Group's UNIX 98 Product Standard.

Solaris releases 2.0 through 9 support the interfaces specified by the System V Interface Definition, Third Edition, Volumes 1 through 4 (SVID3). Note, however, that since the developers of this specification (UNIX Systems Laboratories) are no longer in business and since this specification defers to POSIX and X/Open CAE specifications, there is some disagreement about what is currently required for conformance to this specification.

When Sun WorkShop Compiler™ C 4.2 is installed, Solaris releases 2.0 through 9 support the ANSI X3.159-1989 Programming Language - C and ISO/IEC 9899:1990 Programming Language - C (C) interfaces.

When Sun WorkShop Compiler™ C 5.0 is installed, Solaris releases 7 through 9 also support ISO/IEC 9899 Amendment 1: C Integrity.

When Sun WorkShop Compiler C++ 5.0 is installed, Solaris releases 2.5.1 through 9 support ISO/IEC 14882:1998 Programming Languages - C++. Unsupported features of that standard are described in the compiler README file. The features of the C++ standard adopted from ISO/IEC 9899 Amendment 1 are not supported on Solaris 2.5.1, and are only partially supported on Solaris 2.6.

Utilities If the behavior required by POSIX.2, POSIX.2a, XPG4, SUS, or SUSv2 conflicts with historical Solaris utility behavior, the original Solaris version of the utility is unchanged; a new version that is standard-conforming has been provided in `/usr/xpg4/bin`. For applications wishing to take advantage of POSIX.2, POSIX.2a, XPG4, SUS, or SUSv2 features, the `PATH` (`sh` or `ksh`) or `path` (`csh`) environment variables should be set with `/usr/xpg4/bin` preceding any other directories in which utilities specified by those specifications are found, such as `/bin`, `/usr/bin`, `/usr/ucb`, and `/usr/ccs/bin`.

Feature Test Macros Feature test macros are used by applications to indicate additional sets of features that are desired beyond those specified by the C standard. If an application uses only those interfaces and headers defined by a particular standard (such as POSIX or X/Open CAE), then it need only define the appropriate feature test macro specified by that standard. If the application is using interfaces and headers not defined by that standard, then in addition to defining the appropriate standard feature test macro, it

standards(5)

must also define `__EXTENSIONS__`. Defining `__EXTENSIONS__` provides the application with access to all interfaces and headers not in conflict with the specified standard. The application must define `__EXTENSIONS__` either at compile time or within the application source files.

ANSI/ISO C

No feature test macros need to be defined to indicate that an application is a conforming C application.

ANSI/ISO C++

ANSI/ISO C++ does not define any feature test macros. If the standard C++ announcement macro `__cplusplus` is predefined to value 199711 or greater, the compiler operates in a standard-conforming mode, indicating C++ standards conformance. The value 199711 indicates conformance to ISO/IEC 14882:1998, as required by that standard. (As noted above, conformance to the standard is incomplete.) A standard-conforming mode is not available with compilers prior to Sun Workshop C++ 5.0.

C++ bindings are not defined for POSIX or X/Open CAE, so specifying feature test macros such as `_POSIX_SOURCE` and `_XOPEN_SOURCE` can result in compilation errors due to conflicting requirements of standard C++ and those specifications.

POSIX

Applications that are intended to be conforming POSIX.1 applications must define the feature test macros specified by the standard before including any headers. For the standards listed below, applications must define the feature test macros listed. Application writers must check the corresponding standards for other macros that can be queried to determine if desired options are supported by the implementation.

POSIX Standard	Feature Test Macros
POSIX.1-1990	<code>_POSIX_SOURCE</code>
POSIX.1-1990 and POSIX.2-1992 C-Language Bindings Option	<code>_POSIX_SOURCE</code> and <code>_POSIX_C_SOURCE=2</code>
POSIX.1b-1993	<code>_POSIX_C_SOURCE=199309L</code>
POSIX.1c-1996	<code>_POSIX_C_SOURCE=199506L</code>

SVID3

The SVID3 specification does not specify any feature test macros to indicate that an application is written to meet SVID3 requirements. The SVID3 specification was written before the C standard was completed.

X/Open CAE

To build or compile an application that conforms to one of the X/Open CAE specifications, use the following guidelines. Applications need not set the POSIX feature test macros if they require both CAE and POSIX functionality.

XPG3	The application must define <code>_XOPEN_SOURCE</code> with a value other than 500 (preferably 1).
XPG4	The application must define <code>_XOPEN_SOURCE</code> with a value other than 500 (preferably 1) and set <code>_XOPEN_VERSION=4</code> .
SUS (XPG4v2)	The application must define <code>_XOPEN_SOURCE</code> with a value other than 500 (preferably 1) and set <code>_XOPEN_SOURCE_EXTENDED=1</code> .
SUSv2	The application must define <code>_XOPEN_SOURCE=500</code> .

Compilation

A POSIX.2-, XPG4-, SUS-, or SUSv2-conforming implementation must include an ANSI X3.159-1989 (ANSI C Language) standard-conforming compilation system and the `cc` and `c89` utilities. Solaris 7 through 9 were tested with the `cc` and `c89` utilities and the compilation system provided by Sun WorkShop Compiler™ C 5.0 in the SPARC and x86 environments. When `cc` is used to link applications, `/usr/ccs/lib/values-xpg4.o` must be specified on any link/load command line, but the preferred way to build applications is described below.

An XNS4- or XNS5-conforming application must include `-lXNS` on any link/load command line in addition to defining the feature test macros specified for SUS or SUSv2, respectively.

If the compiler supports the `redefine_extname` pragma feature (the Sun WorkShop Compiler™ C 4.2 and Sun WorkShop Compiler™ C 5.0 compilers define the macro `__PRAGMA_REDEFINE_EXTNAME` to indicate that it supports this feature), then the standard headers use `#pragma redefine_extname` directives to properly map function names onto library entry point names. This mapping provides full support for ISO C, POSIX, and X/Open namespace reservations. The Sun WorkShop Compiler™ C 5.0 compiler was used for all branding and certification tests for Solaris releases 7 through 9.

If this pragma feature is not supported by the compiler, the headers use the `#define` directive to map internal function names onto appropriate library entry point names. In this instance, applications should avoid using the explicit 64-bit file offset symbols listed on the `lf64(5)` manual page, since these names are used by the implementation to name the alternative entry points.

When using Sun WorkShop Compiler™ C 5.0, applications conforming to the specifications listed above should be compiled using the utilities and flags indicated in the following table:

Specification	Compiler/Flags	Feature Test Macros
ANSI/ISO C	<code>c89</code> <code>none</code>	

standards(5)

Specification	Compiler/Flags	Feature Test Macros
SVID3	cc -Xt	none
POSIX.1-1990	c89	_POSIX_SOURCE
POSIX.1-1990 and POSIX.2-1992 C-Language Bindings Option	c89	_POSIX_SOURCE and POSIX_C_SOURCE=2
POSIX.1b-1993	c89	_POSIX_C_SOURCE=199309L
POSIX.1c-1996	c89	_POSIX_C_SOURCE=199506L
CAE XPG3	cc -Xa	_XOPEN_SOURCE
CAE XPG4	c89	_XOPEN_SOURCE and _XOPEN_VERSION=4
SUS (CAE XPG4v2) (includes XNS4)	c89	_XOPEN_SOURCE and _XOPEN_SOURCE_EXTENDED=1
SUSv2 (includes XNS5)	c89	_XOPEN_SOURCE=500

For platforms supporting the LP64 (64-bit) programming environment where the SC5.0 Compilers have been installed, SUSv2-conforming LP64 applications using XNS5 library calls should be built with command lines of the form:

```
c89 $(getconf XBS5_LP64_OFF64_CFLAGS) -D_XOPEN_SOURCE=500 \
    $(getconf XBS5_LP64_OFF64_LDFLAGS) foo.c -o foo \
    $(getconf XBS5_LP64_OFF64_LIBS) -lxnet
```

SEE ALSO sysconf(3C), environ(5), lf64(5)

NAME	sticky – mark files for special treatment
DESCRIPTION	<p>The <i>sticky bit</i> (file mode bit 01000, see <code>chmod(2)</code>) is used to indicate special treatment of certain files and directories. A directory for which the sticky bit is set restricts deletion of files it contains. A file in a sticky directory may only be removed or renamed by a user who has write permission on the directory, and either owns the file, owns the directory, or is the super-user. This is useful for directories such as <code>/tmp</code>, which must be publicly writable, but should deny users permission to arbitrarily delete or rename the files of others.</p> <p>If the sticky bit is set on a regular file and no execute bits are set, the system's page cache will not be used to hold the file's data. This bit is normally set on swap files of diskless clients so that accesses to these files do not flush more valuable data from the system's cache. Moreover, by default such files are treated as swap files, whose inode modification times may not necessarily be correctly recorded on permanent storage.</p> <p>Any user may create a sticky directory. See <code>chmod</code> for details about modifying file modes.</p>
FILES	<code>/tmp</code>
SEE ALSO	<code>chmod(1)</code> , <code>chmod(2)</code> , <code>chown(2)</code> , <code>mkdir(2)</code>
BUGS	<code>mkdir(2)</code> will not create a directory with the sticky bit set.

term(5)

NAME	term – conventional names for terminals																					
DESCRIPTION	<p>Terminal names are maintained as part of the shell environment in the environment variable <code>TERM</code>. See <code>sh(1)</code>, <code>profile(4)</code>, and <code>environ(5)</code>. These names are used by certain commands (for example, <code>tabs</code>, <code>tput</code>, and <code>vi</code>) and certain functions (for example, see <code>curses(3CURSES)</code>).</p> <p>Files under <code>/usr/share/lib/terminfo</code> are used to name terminals and describe their capabilities. These files are in the format described in <code>terminfo(4)</code>. Entries in <code>terminfo</code> source files consist of a number of comma-separated fields. To print a description of a terminal <i>term</i>, use the command <code>infocmp -I term</code>. See <code>infocmp(1M)</code>. White space after each comma is ignored. The first line of each terminal description in the <code>terminfo</code> database gives the names by which <code>terminfo</code> knows the terminal, separated by bar () characters. The first name given is the most common abbreviation for the terminal (this is the one to use to set the environment variable <code>TERMINFO</code> in <code>\$HOME/.profile</code>; see <code>profile(4)</code>), the last name given should be a long name fully identifying the terminal, and all others are understood as synonyms for the terminal name. All names but the last should contain no blanks and must be unique in the first 14 characters; the last name may contain blanks for readability.</p> <p>Terminal names (except for the last, verbose entry) should be chosen using the following conventions. The particular piece of hardware making up the terminal should have a root name chosen, for example, for the AT&T 4425 terminal, <code>att4425</code>. This name should not contain hyphens, except that synonyms may be chosen that do not conflict with other names. Up to 8 characters, chosen from the set <code>a</code> through <code>z</code> and <code>0</code> through <code>9</code>, make up a basic terminal name. Names should generally be based on original vendors rather than local distributors. A terminal acquired from one vendor should not have more than one distinct basic name. Terminal sub-models, operational modes that the hardware can be in, or user preferences should be indicated by appending a hyphen and an indicator of the mode. Thus, an AT&T 4425 terminal in 132 column mode is <code>att4425-w</code>. The following suffixes should be used where possible:</p> <table><thead><tr><th>Suffix</th><th>Meaning</th><th>Example</th></tr></thead><tbody><tr><td><code>-w</code></td><td>Wide mode (more than 80 columns)</td><td><code>att4425-w</code></td></tr><tr><td><code>-am</code></td><td>With auto. margins (usually default)</td><td><code>vt100-am</code></td></tr><tr><td><code>-nam</code></td><td>Without automatic margins</td><td><code>vt100-nam</code></td></tr><tr><td><code>-n</code></td><td>Number of lines on the screen</td><td><code>aaa-60</code></td></tr><tr><td><code>-na</code></td><td>No arrow keys (leave them in local)</td><td><code>c100-na</code></td></tr><tr><td><code>-np</code></td><td>Number of pages of memory</td><td><code>c100-4p</code></td></tr></tbody></table>	Suffix	Meaning	Example	<code>-w</code>	Wide mode (more than 80 columns)	<code>att4425-w</code>	<code>-am</code>	With auto. margins (usually default)	<code>vt100-am</code>	<code>-nam</code>	Without automatic margins	<code>vt100-nam</code>	<code>-n</code>	Number of lines on the screen	<code>aaa-60</code>	<code>-na</code>	No arrow keys (leave them in local)	<code>c100-na</code>	<code>-np</code>	Number of pages of memory	<code>c100-4p</code>
Suffix	Meaning	Example																				
<code>-w</code>	Wide mode (more than 80 columns)	<code>att4425-w</code>																				
<code>-am</code>	With auto. margins (usually default)	<code>vt100-am</code>																				
<code>-nam</code>	Without automatic margins	<code>vt100-nam</code>																				
<code>-n</code>	Number of lines on the screen	<code>aaa-60</code>																				
<code>-na</code>	No arrow keys (leave them in local)	<code>c100-na</code>																				
<code>-np</code>	Number of pages of memory	<code>c100-4p</code>																				

-rv	Reverse video	att4415-rv
-----	---------------	------------

To avoid conflicts with the naming conventions used in describing the different modes of a terminal (for example, -w), it is recommended that a terminal's root name not contain hyphens. Further, it is good practice to make all terminal names used in the terminfo(4) database unique. Terminal entries that are present only for inclusion in other entries via the use= facilities should have a '+' in their name, as in 4415+n1.

Here are some of the known terminal names: (For a complete list, enter the command `ls -C /usr/share/lib/terminfo/?`).

2621,hp2621	Hewlett-Packard 2621 series
2631	Hewlett-Packard 2631 line printer
2631-c	Hewlett-Packard 2631 line printer, compressed mode
2631-e	Hewlett-Packard 2631 line printer, expanded mode
2640,hp2640	Hewlett-Packard 2640 series
2645,hp2645	Hewlett-Packard 2645 series
3270	IBM Model 3270
33,tty33	AT&T Teletype Model 33 KSR
35,tty35	AT&T Teletype Model 35 KSR
37,tty37	AT&T Teletype Model 37 KSR
4000a	Trendata 4000a
4014,tek4014	TEKTRONIX 4014
40,tty40	AT&T Teletype Dataspeed 40/2
43,tty43	AT&T Teletype Model 43 KSR
4410,5410	AT&T 4410/5410 in 80-column mode, version 2
4410-nfk,5410-nfk	AT&T 4410/5410 without function keys, version 1
4410-nsl,5410-nsl	AT&T 4410/5410 without pln defined
4410-w,5410-w	AT&T 4410/5410 in 132-column mode
4410v1,5410v1	AT&T 4410/5410 in 80-column mode, version 1
4410v1-w,5410v1-w	AT&T 4410/5410 in 132-column mode, version 1
4415,5420	AT&T 4415/5420 in 80-column mode
4415-nl,5420-nl	AT&T 4415/5420 without changing labels

term(5)

4415-rv,5420-rv	AT&T 4415/5420 80 columns in reverse video
4415-rv-nl,5420-rv-nl	AT&T 4415/5420 reverse video without changing labels
4415-w,5420-w	AT&T 4415/5420 in 132-column mode
4415-w-nl,5420-w-nl	AT&T 4415/5420 in 132-column mode without changing labels
4415-w-rv,5420-w-rv	AT&T 4415/5420 132 columns in reverse video
4418,5418	AT&T 5418 in 80-column mode
4418-w,5418-w	AT&T 5418 in 132-column mode
4420	AT&T Teletype Model 4420
4424	AT&T Teletype Model 4424
4424-2	AT&T Teletype Model 4424 in display function group ii
4425,5425	AT&T 4425/5425
4425-fk,5425-fk	AT&T 4425/5425 without function keys
4425-nl,5425-nl	AT&T 4425/5425 without changing labels in 80-column mode
4425-w,5425-w	AT&T 4425/5425 in 132-column mode
4425-w-fk,5425-w-fk	AT&T 4425/5425 without function keys in 132-column mode
4425-nl-w,5425-nl-w	AT&T 4425/5425 without changing labels in 132-column mode
4426	AT&T Teletype Model 4426S
450	DASI 450 (same as Diablo 1620)
450-12	DASI 450 in 12-pitch mode
500,att500	AT&T-IS 500 terminal
510,510a	AT&T 510/510a in 80-column mode
513bct,att513	AT&T 513 bct terminal
5320	AT&T 5320 hardcopy terminal
5420_2	AT&T 5420 model 2 in 80-column mode
5420_2-w	AT&T 5420 model 2 in 132-column mode
5620,dmd	AT&T 5620 terminal 88 columns
5620-24,dmd-24	AT&T Teletype Model DMD 5620 in a 24x80 layer
5620-34,dmd-34	AT&T Teletype Model DMD 5620 in a 34x80 layer

610,610bct	AT&T 610 bct terminal in 80-column mode
610-w,610bct-w	AT&T 610 bct terminal in 132-column mode
630,630MTG	AT&T 630 Multi-Tasking Graphics terminal
7300,pc7300,unix_pc	AT&T UNIX PC Model 7300
735,ti	Texas Instruments TI735 and TI725
745	Texas Instruments TI745
dumb	generic name for terminals that lack reverse line-feed and other special escape sequences
hp	Hewlett-Packard (same as 2645)
lp	generic name for a line printer
pt505	AT&T Personal Terminal 505 (22 lines)
pt505-24	AT&T Personal Terminal 505 (24-line mode)
sync	generic name for synchronous Teletype Model 4540-compatible terminals

Commands whose behavior depends on the type of terminal should accept arguments of the form `-Tterm` where *term* is one of the names given above; if no such argument is present, such commands should obtain the terminal type from the environment variable `TERM`, which, in turn, should contain *term*.

FILES /usr/share/lib/terminfo/?/* compiled terminal description database

SEE ALSO sh(1), stty(1), tabs(1), tput(1), vi(1), infocmp(1M), curses(3CURSES), profile(4), terminfo(4), environ(5)

vgrindefs(5)

NAME vgrindefs – vgrind’s language definition data base

SYNOPSIS /usr/lib/vgrindefs

DESCRIPTION vgrindefs contains all language definitions for vgrind(1). Capabilities in vgrindefs are of two types: Boolean capabilities which indicate that the language has some particular feature and string capabilities which give a regular expression or keyword list. Entries may continue onto multiple lines by giving a \ as the last character of a line. Lines starting with # are comments.

Capabilities The following table names and describes each capability.

Name	Type	Description
ab	str	Regular expression for the start of an alternate form comment
ae	str	Regular expression for the end of an alternate form comment
bb	str	Regular expression for the start of a block
be	str	Regular expression for the end of a lexical block
cb	str	Regular expression for the start of a comment
ce	str	Regular expression for the end of a comment
id	str	String giving characters other than letters and digits that may legally occur in identifiers (default ‘_’)
kw	str	A list of keywords separated by spaces
lb	str	Regular expression for the start of a character constant
le	str	Regular expression for the end of a character constant
oc	bool	Present means upper and lower case are equivalent
pb	str	Regular expression for start of a procedure
pl	bool	Procedure definitions are constrained to the lexical level matched by the ‘px’ capability
px	str	A match for this regular expression indicates that procedure definitions may occur at the next lexical level. Useful for lisp-like languages in which procedure definitions occur as subexpressions of defuns.
sb	str	Regular expression for the start of a string
se	str	Regular expression for the end of a string
tc	str	Use the named entry as a continuation of this one
tl	bool	Present means procedures are only defined at the top lexical level

Regular Expressions	<p>vgrindefs uses regular expressions similar to those of <code>ex(1)</code> and <code>lex(1)</code>. The characters '^', '\$', ':', and '\' are reserved characters and must be 'quoted' with a preceding \ if they are to be included as normal characters. The metasymbols and their meanings are:</p> <p>\$ The end of a line</p> <p>^ The beginning of a line</p> <p>\d A delimiter (space, tab, newline, start of line)</p> <p>\a Matches any string of symbols (like '.'* in lex)</p> <p>\p Matches any identifier. In a procedure definition (the 'pb' capability) the string that matches this symbol is used as the procedure name.</p> <p>() Grouping</p> <p> Alternation</p> <p>? Last item is optional</p> <p>\e Preceding any string means that the string will not match an input string if the input string is preceded by an escape character (\). This is typically used for languages (like C) that can include the string delimiter in a string by escaping it.</p> <p>Unlike other regular expressions in the system, these match words and not characters. Hence something like '(tramp steamer)flies?' would match 'tramp', 'steamer', 'trampflies', or 'steamerflies'. Contrary to some forms of regular expressions, vgrindef alternation binds very tightly. Grouping parentheses are likely to be necessary in expressions involving alternation.</p>
Keyword List	<p>The keyword list is just a list of keywords in the language separated by spaces. If the 'oc' boolean is specified, indicating that upper and lower case are equivalent, then all the keywords should be specified in lower case.</p>
EXAMPLES	<p>EXAMPLE 1 A sample program.</p> <p>The following entry, which describes the C language, is typical of a language entry.</p> <pre>C c the C programming language:\ :pb=^\\d?*?\\d?\\p\\d? (\\a?\\) (\\d {):bb={:be=}:cb=/*:ce=*/:sb=":se=\\e":\ :le=\\e':tl:\ :kw=asm auto break case char continue default do double else enum\ extern float for fortran goto if int long register return short\ sizeof static struct switch typedef union unsigned void while #define\ #else #endif #if #ifdef #ifndef #include #undef # define endif\ ifdef ifndef include undef defined:</pre> <p>Note that the first field is just the language name (and any variants of it). Thus the C language could be specified to <code>vgrind(1)</code> as 'c' or 'C'.</p>

vgrindefs(5)

FILES	/usr/lib/vgrindefs	file containing vgrind descriptions
SEE ALSO	ex(1), lex(1), troff(1), vgrind(1)	

NAME	wbem – Web-Based Enterprise Management
DESCRIPTION	<p>Web-Based Enterprise Management (WBEM) is a set of management and Internet-related technologies intended to unify the management of enterprise computing environments. Developed by the Distributed Management Task Force (DMTF), WBEM enables organizations to deliver an integrated set of standards-based management tools that support and promote World Wide Web technology. The DMTF has developed a set of standards that make up WBEM. This set of standards includes:</p> <ul style="list-style-type: none"> ■ Common Information Model (CIM) - an object-oriented data model that describes the overall management of information in an enterprise network environment. CIM consists of a CIM specification and a CIM schema: <ul style="list-style-type: none"> CIM Specification Consists of the language and methodology that describes management data. CIM Schema Provides actual model descriptions of systems, applications, large area networks, and devices. The CIM Schema enables applications from different developers on different platforms to describe management data in a standard format. As a result, a variety of management applications can share this information. ■ CIM Operations Over HyperText Transport Protocol (HTTP) 1.1 is a transport mechanism that maps CIM operations to HTTP to allow implementations of CIM to interoperate in an open, standardized manner. <p>CIM Operations Over HTTP 1.1 uses eXtensible Markup Language (XML), which is a markup language that represents management information in textual form.</p> <p>In addition to the XML representation, CIM information is also represented textually by the managed object format (MOF). These MOF representations are typically stored as text files that developers compile into a CIM Object Manager.</p>
EXTENDED DESCRIPTION WBEM Tools and Services	<p>Tools and services that enable developers to create and Services management applications and instrumentation that manage heterogeneous computer environments include:</p> <ul style="list-style-type: none"> ■ Solaris WBEM Services 2.5 ■ Solaris WBEM Software Development Kit 2.5
Solaris WBEM Services 2.5	<p>These services consist of a set of value-added Services 2.5 components. These services make it easier for developers to create management applications that run in the Solaris operating environment. They also make the Solaris operating environment easier to manage. Solaris WBEM Services 2.5 consists of:</p> <ul style="list-style-type: none"> ■ CIM Object Manager, CIM Repository, and MOF Compiler

wbem(5)

- CIM and Solaris Schema, which is an extension schema of CIM. CIM and Solaris Schema is a collection of CIM classes that describe managed elements in the Solaris operating environment. These classes are available from the CIM Object Manager at start up.

- Solaris Providers, which are programs that communicate information between the Solaris operating environment and the CIM Object Manager (providers get and set "dynamic" information about managed elements, acting as an intermediary between the CIM Object Manager and the managed elements).

Solaris software providers have been developed for a variety of areas: users, roles, file systems, and network configuration, for example. A remote provider is also available to distribute agents away from the CIM Object Manager when required. Because of the incremental development capabilities of the WBEM instrumentation framework, developers can progressively and consistently add more providers for additional Solaris software services.

- SNMP Adapter for WBEM, which enables Simple Network Management Protocol (SNMP) management applications to access system management information that is provided by Solaris WBEM Services. Used with the Solstice Enterprise Agent (SEA) Master Agent `snmpdx(1M)`, the SNMP Adapter for WBEM maps SNMP requests into equivalent WBEM Common Information Model (CIM) properties or instances.

The SNMP Adapter for WBEM also remaps the response from the CIM Object Manager into an SNMP response, which is returned to the management application.

A mapping file contains the corresponding Object Identifier (OID), class name, property name, and Abstract Syntax Notation One (ASN.1) type for each object. Developers can create their own mapping files.

- SNMP Provider, which enables WBEM services to deliver SNMP information.

**Solaris WBEM
Software
Development Kit
2.5**

This kit consists of a set of key application Software development tools that make it easier for developers to write management applications that can communicate with any WBEM-enabled management device. The Solaris WBEM Software Development Kit includes examples, documentation, and CIM Workshop, a graphical user interface through which developers can view and create classes and instances, through the remote method invocation (RMI) or the XML/HTTP protocol.

Developers can also use this kit to write providers, which are programs that communicate with managed elements to access data.

All management applications that developers create with the Solaris WBEM Software Development Kit run on the Java platform. The Solaris 9 WBEM Software Development Kit installs and runs in version 1.4 of the Java environment. Developers can use the kit to write standalone applications or applications that run in conjunction with Solaris WBEM Services.

The Solaris WBEM Software Development Kit is described in the *Solaris WBEM SDK Development Guide*. Javadoc for the WBEM application programming interface is located at `/usr/sadm/lib/wbem/doc/index.html`.

Compatibility of Solaris WBEM Services with Existing Protocols

Adapters and converters enable Solaris WBEM Services of Solaris to work compatibly with existing protocols by mapping WBEM information to these protocols. One such protocol is Simple Network Management Protocol (SNMP).

Legacy management applications can administer WBEM-enabled software in the Solaris operating environment. Developers can write agents or providers that convert information from these protocols to WBEM, and they can write adapters that convert WBEM information into these protocols.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SPARC and x86
Architecture	SUNWwbapi, SUNWwbcor, SUNWwbcou, SUNWwbdev, SUNWwbdoc, SUNWwbpro
CSI	Enabled

SEE ALSO

`appletviewer(1)`, `cimworkshop(1M)`, `init.wbem(1M)`, `mofcomp(1M)`, `mofreg(1M)`, `snmpdx(1M)`, `wbemadmin(1M)`, `wbemconfig(1M)`, `wbemlogviewer(1M)`, `attributes(5)`

wbem(5)

Index

A

- account management PAM module for projects
 - pam_projects, 285
- ANSI — standards and specifications supported by Solaris, 328
- architecture — attributes of interfaces, 18
- ascii — ASCII character set, 16
- attributes — attributes of interfaces, 18
- attributes of interfaces — architecture, 18
- attributes of interfaces — attributes, 18
- attributes of interfaces — availability, 18
- attributes of interfaces — CSI, 18
- attributes of interfaces — MT-Level, 18
- attributes of interfaces — stability, 18
- attributes — characteristics of commands, utilities, and device drivers
 - Architecture, 18
 - Availability, 18
 - Interface Stability, 19
 - MT-Level, 22
- authentication, account, session, and password management PAM modules for Kerberos V5
 - pam_krb5, 273
- authentication and password management module — pam_authtok_check, 265
- authentication and password management module — pam_authtok_get, 267
- authentication module for password — pam_passwd_auth, 283
- availability — attributes of interfaces, 18

C

- C — standards and specifications supported by Solaris, 328
- C++ — standards and specifications supported by Solaris, 328
- character set description file — charmap, 26
- charmap — character set description file, 26
 - Decimal Constants, 28
 - Declarations, 26
 - Format, 27
 - Ranges of Symbolic Names, 28
 - Symbolic Names, 26
- code set conversion tables — iconv, 122
- code set conversion tables — iconv_1250, 107
- code set conversion tables — iconv_1251, 113
- code set conversion tables — iconv_646, 126
- code set conversion tables — iconv_852, 129
- code set conversion tables — iconv_8859-1, 135
- code set conversion tables — iconv_8859-2, 141
- code set conversion tables — iconv_8859-5, 147
- code set conversion tables — iconv_dhn, 155
- code set conversion tables — iconv_koi8-r, 159
- code set conversion tables —
 - iconv_mac_cyr, 167
- code set conversion tables — iconv_maz, 175
- code set conversion tables — iconv_pc_cyr, 179
- code set conversion tables —
 - iconv_unicode, 185
- code set conversion tables
 - iconv_1250, 107
 - iconv_1251, 113
 - iconv_646, 126
 - iconv_852, 129

code set conversion tables (Continued)

- iconv_8859-1, 135
- iconv_8859-2, 141
- iconv_8859-5, 147
- iconv_dhn, 155
- iconv_koi8-r, 159
- iconv_mac_cyr, 167
- iconv_maz, 175
- iconv_pc_cyr, 179

compilation environment, transitional —

- lfcompile64, 206

crypt_unix — traditional UNIX crypt algorithm, 32

CSI — attributes of interfaces, 18

D

data storage modules for the DHCP service — dhcp_modules, 35

dhcp — Dynamic Host Configuration Protocol, 33

dhcp_modules — data storage modules for the DHCP service, 35

document production

- man — macros to format manual pages, 238

- mansun — macros to format manual pages, 242

- me — macros to format technical papers, 246

- mm — macros to format articles, theses and books, 251

- ms — macros to format articles, theses and books, 258

Dynamic Host Configuration Protocol — dhcp, 33

E

environ — user environment, 37

environment variables

- HOME, 37
- LANG, 37
- LC_COLLATE, 37
- LC_CTYPE, 37
- LC_MESSAGES, 37
- LC_MONETARY, 37

environment variables (Continued)

- LC_NUMERIC, 37
- LC_TIME, 37
- MSGVERB, 37
- NETPATH, 37
- PATH, 37
- SEV_LEVEL, 37
- TERM, 37
- TZ, 37

eqnchar — special character definitions for eqn, 42

extended file attributes — fsattr, 96

extensions — localedef extensions description file, 43

F

file format notation — formats, formats, 91

file name pattern matching — fnmatch, 63

filesystem — file system layout, 44

filesystem — file system organization, 44

- Root File System, 44

- /usr File System, 53

fnmatch — file name pattern matching, 63

fns — overview of FNS, 67

- Composite Names, 67

- FNS and Naming Systems, 68

FNS

- overview — fns, 67

- overview of FNS References — fns_references, 85

- overview over DNS implementation — fns_dns, 69

- overview over files implementation — fns_files, 71

- overview over NIS+ implementation — fns_nis+, 77

- overview over NIS (YP) implementation — fns_nis, 79

- overview over X.500 implementation — fns_x500, 88

fns — overview of FNS

- Why FNS?, 67

- XFN, 67

fns_dns — overview of FNS over DNS implementation, 69

fns_files — overview of FNS over files
 implementation, 71
 fns_files — overview of FNS over /etc files
 implementation, FNS Policies and /etc
 Files, 71
 fns_initial_context — overview of the FNS
 Initial Context, 73
 fns_nis — overview of FNS over NIS (YP)
 implementation, 79
 Federating NIS with DNS or X.500, 79
 FNS Policies and NIS, 79
 NIS Security, 79
 fns_nis+ — overview of FNS over NIS+
 implementation, 77
 FNS Policies and NIS+, 77
 fns_policies — overview of the FNS Policies, 81
 fns_references — overview of FNS
 References, 85
 Address Types, 85
 Reference Types, 85
 fns_x500 — overview of FNS over X.500
 implementation, 88
 formats — file format notation, 91
 fsattr — extended file attributes, 96

I

iconv — code set conversion tables, 122
 iconv_1250 — code set conversion tables for MS
 1250 (Windows Latin 2), 107
 iconv_1251 — code set conversion tables for MS
 1251 (Windows Cyrillic), 113
 iconv_646 — code set conversion tables for ISO
 646, 126
 iconv_852 — code set conversion tables for MS
 852 (MS-DOS Latin 2), 129
 iconv_8859-1 — code set conversion tables for
 ISO 8859-1 (Latin 1), 135
 iconv_8859-2 — code set conversion tables for
 ISO 8859-2 (Latin 2), 141
 iconv_8859-5 — code set conversion tables for
 ISO 8859-5 (Cyrillic), 147
 iconv_dhn — code set conversion tables for
 DHN (Dom Handlowy Nauki), 155
 iconv_koi8-r — code set conversion tables for
 KOI8-R, 159

iconv_mac_cyr — code set conversion tables for
 Macintosh Cyrillic, 167
 iconv_maz — code set conversion tables for
 Mazovia, 175
 iconv_pc_cyr — code set conversion tables for
 Alternative PC Cyrillic, 179
 iconv_unicode — code set conversion tables for
 Unicode, 185
 internationalized basic and extended regular
 expression matching — regex, 304
 isalist — the native instruction sets known to
 Solaris, 190
 ISO — standards and specifications supported
 by Solaris, 328

L

large file status of utilities — largefile, 192
 largefile — large file status of utilities, 192
 Large file aware utilities, 192
 Large file safe utilities, 194
 lf64 — transitional interfaces for 64-bit file
 offsets, 196
 Data Types, 196
 System Interfaces, 198
 lfcompile — large file compilation environment,
 Access to Additional Large File
 Interfaces, 203
 lfcompile64 — transitional compilation
 environment, 206
 Access to Additional Large File
 Interfaces, 206
 live_upgrade, 208
 locale — subset of a user's environment that
 depends on language and cultural
 conventions, 212
 collating-element keyword, 220
 collating-symbol keyword, 221
 Collation Order, 222
 LC_COLLATE, 219
 LC_CTYPE, 215
 LC_MESSAGES, 236
 LC_MONETARY, 225
 LC_NUMERIC, 229
 LC_TIME, 231
 LC_TIME C-language Access, 233
 LC_TIME General Information, 236

locale — subset of a user's environment that depends on language and cultural conventions (Continued)

- Locale Definition, 212
- order_end keyword, 225
- order_start keyword, 221

localedef extensions description file — extensions, 43

M

macros

- to format articles, theses and books — mm, 251

- to format articles, theses and books — ms, 258

- to format Manual pages — man, 238

- to format Manual pages — mansun, 242

- to format technical papers — me, 246

man — macros to format manual pages, 238

mansun — macros to format manual pages, 242

manual pages

- macros to format manual pages — man, 238

- Sun macros to format manual pages — mansun, 242

mark files for special treatment — sticky, 333

me — macros to format technical papers, 246

mm — macros to format articles, theses and books, 251

ms — macros to format articles, theses and books, 258

MT-Level — attributes of interfaces, 18

N

native instruction sets known to Solaris — isalist, 190

NFS and sticky bits — sticky, 333

nfssec — overview of NFS security modes, 263

O

overview of FNS — fns, 67

overview of FNS over DNS implementation — fns_dns, 69

overview of FNS over files implementation — fns_files, 71

overview of FNS over NIS+ implementation — fns_nis+, 77

overview of FNS over NIS (YP) implementation — fns_nis, 79

overview of FNS over X.500 implementation — fns_x500, 88

overview of FNS References — fns_references, 85

overview of NFS security modes — nfssec, 263

overview of the FNS Initial Context — fns_initial_context, 73

overview of the FNS Policies — fns_policies, 81

P

PAM account management module for UNIX — pam_unix_account, 296

pam_authok_check — authentication and password management module, 265

pam_authok_get — authentication and password management module, 267

pam_authok_store — password management module, 269

pam_dial_auth — authentication management for dialups, 272

pam_krb5 — authentication, account, session and password management for Kerberos V5, 273

pam_krb5 — authentication, account, session, and password management PAM modules for Kerberos V5, 273

pam_passwd_auth — authentication module for password, 283

pam_projects — account management PAM module for projects, 285

pam_rhosts_auth — authentication management using ruserok(), 286

pam_sample — sample module for PAM, 289

pam_unix — authentication, account, session and password management for UNIX, 293

pam_unix_account — PAM account management module for UNIX, 296

pam_unix_session — session management PAM module for UNIX, 299
password management module —
 pam_authok_store, 269
POSIX — standards and specifications supported by Solaris, 328
POSIX.1 — standards and specifications supported by Solaris, 328
POSIX.2 — standards and specifications supported by Solaris, 328
profiling utilities, profile within a function —
 prof, 300

R

rbac — role-based access control, 301
regex — internationalized basic and extended regular expression matching, 304
regular expression compile and match routines —
 — advance, 313
 — compile, 313
 — regex, 313
 — step, 313
role-based access control — rbac, 301

S

ftp — authentication system, 320
session management PAM module for UNIX —
 pam_unix_session, 299
sgml — Standard Generalized Markup Language, 322
sgml — Standard Generalized Markup language, RefEntry, 322
sgml — Standard Generalized Markup Language
 RefMeta, 322
 RefNameDiv, 323
 RefSect1, 324
 RefSect2, 324
 RefSynopsisDiv, 324
shell environment, conventional names for terminals — term, 334
solbook — Standard Generalized Markup Language, 322

special character definitions for eqn —
 eqnchar, 42
stability — attributes of interfaces, 18
Standard Generalized Markup Language —
 — sgml, 322
 — solbook, 322
standards — standards and specifications supported by Solaris, 328
standards and specifications supported by Solaris — ANSI, 328
standards and specifications supported by Solaris — C++, 328
standards and specifications supported by Solaris — C, 328
standards and specifications supported by Solaris — ISO, 328
standards and specifications supported by Solaris — POSIX.1, 328
standards and specifications supported by Solaris — POSIX.2, 328
standards and specifications supported by Solaris — POSIX, 328
standards and specifications supported by Solaris — standards, 328
standards and specifications supported by Solaris — SUS, 328
standards and specifications supported by Solaris — SUSv2, 328
standards and specifications supported by Solaris — SVID3, 328
standards and specifications supported by Solaris — SVID, 328
standards and specifications supported by Solaris — XNS4, 328
standards and specifications supported by Solaris — XNS5, 328
standards and specifications supported by Solaris — XNS, 328
standards and specifications supported by Solaris — XPG3, 328
standards and specifications supported by Solaris — XPG4, 328
standards and specifications supported by Solaris — XPG4v2, 328
standards and specifications supported by Solaris — XPG, 328
sticky — mark files for special treatment, 333

subset of a user's environment that depends on language and cultural conventions — locale, 212
SUS — standards and specifications supported by Solaris, 328
SUSv2 — standards and specifications supported by Solaris, 328
SVID — standards and specifications supported by Solaris, 328
SVID3 — standards and specifications supported by Solaris, 328

T

term — conventional names for terminals, 334
terminals, conventional names — term, 334
traditional UNIX crypt algorithm — crypt_unix, 32
transitional compilation environment — lfcompile64, 206
transitional interfaces for 64-bit file offsets — lf64, 196

U

unicode, code set conversion tables — iconv_unicode, 185
user environment, — environ, 37

V

vgrindefs — vgrind language definitions, 338

W

wbem — Web-Based Enterprise Management, 341
Web-Based Enterprise Management — wbem, 341

X

XNS — standards and specifications supported by Solaris, 328
XNS4 — standards and specifications supported by Solaris, 328
XNS5 — standards and specifications supported by Solaris, 328
XPG — standards and specifications supported by Solaris, 328
XPG3 — standards and specifications supported by Solaris, 328
XPG4 — standards and specifications supported by Solaris, 328
XPG4v2 — standards and specifications supported by Solaris, 328