

## **man pages section 5: Standards, Environments, and Macros**

Copyright © 1993, 2012, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

---

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

# Contents

---

<b>Preface</b> .....	9
<b>Introduction</b> .....	13
Intro(5) .....	14
<b>Standards, Environments, and Macros</b> .....	15
acl(5) .....	16
ad(5) .....	23
ars(5) .....	24
ascii(5) .....	27
attributes(5) .....	29
audit_binfile(5) .....	38
audit_flags(5) .....	40
audit_remote(5) .....	41
audit_syslog(5) .....	46
brands(5) .....	49
cancellation(5) .....	51
charmap(5) .....	57
condition(5) .....	61
crypt_bsdbf(5) .....	63
crypt_bsdmd5(5) .....	64
crypt_sha256(5) .....	65
crypt_sha512(5) .....	66
crypt_sunmd5(5) .....	67
crypt_unix(5) .....	68
datasets(5) .....	69
device_clean(5) .....	71

dhcp(5) .....	73
dhcp_modules(5) .....	75
environ(5) .....	76
eqnchar(5) .....	82
extendedFILE(5) .....	83
extensions(5) .....	85
fedfs(5) .....	86
filesystem(5) .....	89
fmri(5) .....	93
fnmatch(5) .....	116
formats(5) .....	120
fsattr(5) .....	125
grub(5) .....	136
gss_auth_rules(5) .....	138
hal(5) .....	139
iconv_1250(5) .....	140
iconv_1251(5) .....	146
iconv(5) .....	155
iconv_646(5) .....	159
iconv_852(5) .....	162
iconv_8859-1(5) .....	169
iconv_8859-2(5) .....	176
iconv_8859-5(5) .....	182
iconv_dhn(5) .....	190
iconv_koi8-r(5) .....	194
iconv_mac_cyr(5) .....	202
iconv_maz(5) .....	210
iconv_pc_cyr(5) .....	214
iconv_unicode(5) .....	220
ieee802.11(5) .....	225
ieee802.3(5) .....	226
ipfilter(5) .....	232
isalist(5) .....	236
kerberos(5) .....	238
krb5_auth_rules(5) .....	240
krb5envvar(5) .....	242

---

kssl(5) .....	245
labels(5) .....	246
largefile(5) .....	248
ldap(5) .....	252
lf64(5) .....	256
lfcompile(5) .....	263
lfcompile64(5) .....	266
locale(5) .....	268
locale_alias(5) .....	294
man(5) .....	311
mansun(5) .....	316
me(5) .....	320
mech_spnego(5) .....	325
mm(5) .....	327
ms(5) .....	334
mutex(5) .....	339
mwac(5) .....	341
nfssec(5) .....	342
nis(5) .....	345
nwam(5) .....	346
openssl(5) .....	349
pam_allow(5) .....	353
pam_authok_check(5) .....	355
pam_authok_get(5) .....	357
pam_authok_store(5) .....	359
pam_deny(5) .....	360
pam_dhkeys(5) .....	362
pam_dial_auth(5) .....	364
pam_krb5(5) .....	365
pam_krb5_migrate(5) .....	375
pam_ldap(5) .....	378
pam_list(5) .....	383
pam_passwd_auth(5) .....	386
pam_pkcs11(5) .....	388
pam_rhosts_auth(5) .....	392
pam_roles(5) .....	393

pam_sample(5) .....	396
pam_smbfs_login(5) .....	398
pam_smb_passwd(5) .....	400
pam_tsol_account(5) .....	402
pam_tty_tickets(5) .....	404
pam_unix_account(5) .....	406
pam_unix_auth(5) .....	408
pam_unix_cred(5) .....	411
pam_unix_session(5) .....	413
pam_user_policy(5) .....	414
pam_zfs_key(5) .....	417
pkcs11_kernel(5) .....	420
pkcs11_kms(5) .....	422
pkcs11_softtoken(5) .....	425
pkcs11_tpm(5) .....	429
pkg(5) .....	432
privileges(5) .....	450
prof(5) .....	464
rbac(5) .....	465
regex(5) .....	468
regexp(5) .....	477
resource_controls(5) .....	484
sgml(5) .....	493
smf(5) .....	497
smf_bootstrap(5) .....	505
smf_method(5) .....	508
smf_restarter(5) .....	515
smf_security(5) .....	516
smf_template(5) .....	519
solaris10(5) .....	528
solaris(5) .....	532
standards(5) .....	537
sticky(5) .....	543
suri(5) .....	544
tecla(5) .....	546
term(5) .....	568

---

threads(5) .....	572
trusted_extensions(5) .....	579
vgrindefs(5) .....	580
wbem(5) .....	583
zones(5) .....	586



# 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 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 9E describes the DDI (Device Driver Interface)/DKI (Driver/Kernel Interface), 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> <ul style="list-style-type: none"><li>[ ] Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.</li><li>. . . Ellipses. Several values can be provided for the previous argument, or the previous argument can be specified multiple times, for example, “filename . . .”.</li><li>  Separator. Only one of the arguments separated by this character can be specified at a time.</li><li>{ } Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.</li></ul>
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 <a href="#">ioctl(2)</a> 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).

---

	<p><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>.</p>
OPTIONS	<p>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.</p>
OPERANDS	<p>This section lists the command operands and describes how they affect the actions of the command.</p>
OUTPUT	<p>This section describes the output – standard output, standard error, or output files – generated by the command.</p>
RETURN VALUES	<p>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.</p>
ERRORS	<p>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.</p>
USAGE	<p>This section lists special rules, features, and commands that require in-depth explanations. The subsections listed here are used to explain built-in functionality:</p> <ul style="list-style-type: none"><li>Commands</li><li>Modifiers</li><li>Variables</li><li>Expressions</li><li>Input Grammar</li></ul>
EXAMPLES	<p>This section provides examples of usage or of how to use a command or function. Wherever possible a complete</p>

	<p>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.</p>
ENVIRONMENT VARIABLES	<p>This section lists any environment variables that the command or function affects, followed by a brief description of the effect.</p>
EXIT STATUS	<p>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.</p>
FILES	<p>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.</p>
ATTRIBUTES	<p>This section lists characteristics of commands, utilities, and device drivers by defining the attribute type and its corresponding value. See <a href="#">attributes(5)</a> for more information.</p>
SEE ALSO	<p>This section lists references to other man pages, in-house documentation, and outside publications.</p>
DIAGNOSTICS	<p>This section lists diagnostic messages with a brief explanation of the condition causing the error.</p>
WARNINGS	<p>This section lists warnings about special conditions which could seriously affect your working conditions. This is not a list of diagnostics.</p>
NOTES	<p>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.</p>
BUGS	<p>This section describes known bugs and, wherever possible, suggests workarounds.</p>

**R E F E R E N C E**

**Introduction**

**Name** Intro – introduction to miscellany

**Description** Among the topics presented in this section are:

Standards	The POSIX (IEEE) Standards and the X/Open Specifications are described on the standards 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.

**Acknowledgments** Oracle America, Inc. gratefully acknowledges The Open Group for permission to reproduce portions of its copyrighted documentation. Original documentation from The Open Group can be obtained online at <http://www.opengroup.org/bookstore/>.

The Institute of Electrical and Electronics Engineers and The Open Group, have given us permission to reprint portions of their documentation.

In the following statement, the phrase “this text” refers to portions of the system documentation.

Portions of this text are reprinted and reproduced in electronic form in the SunOS Reference Manual, from IEEE Std 1003.1, 2004 Edition, Standard for Information Technology -- Portable Operating System Interface (POSIX), The Open Group Base Specifications Issue 6, Copyright (C) 2001-2004 by the Institute of Electrical and Electronics Engineers, Inc and The Open Group. In the event of any discrepancy between these versions and the original IEEE and The Open Group Standard, the original IEEE and The Open Group Standard is the referee document. The original Standard can be obtained online at <http://www.opengroup.org/unix/online.html>.

This notice shall appear on any product containing this material.

## REFERENCE

### Standards, Environments, and Macros

**Name** acl – Access Control Lists

**Description** Access control lists (ACLs) are discretionary access control mechanisms that grant and deny access to files and directories. Two different ACL models are supported in the Solaris release: POSIX-draft ACLs and NFSv4 ACLs.

The older, POSIX-draft model is supported by the UFS file system. This model is based on a withdrawn ACL POSIX specification that was never standardized. It was subsequently withdrawn by the POSIX committee.

The other model is based on the standards of the NFSv4 working group and is an approved standard from the Internet Engineering Task Force (IETF). The ZFS file system uses the NFSv4 model, and provides richer semantics and finer grained permission capabilities than the POSIX-draft model.

POSIX-draft ACLs provide an alternative security mechanism to basic UNIX file permissions in the Solaris release. Their purpose is to further restrict access to files and directories or to extend permissions to a particular user. ACLs can be used to change the permissions for the standard owner, group and other class bits of a file's mode. ACLs can give additional users and groups access to the file. A directory can also have a special kind of ACL called a *default* ACL, which defines ACL entries to be inherited by descendants of the directory. POSIX-draft ACLs have an ACL entry called *mask*. The mask defines the maximum permissions that can be granted to additional user and group entries. Whenever a file is created or its mode is changed by `chmod(1)` or `chmod(2)`, the mask is recomputed. It is recomputed to be the group permission defined in the mode passed to `chmod(2)`.

The POSIX-draft ACL model uses the standard rwx model of traditional UNIX permissions.

An ACL is represented as follows:

```
acl_entry[ ,acl_entry] . . .
```

Each *acl\_entry* contains one ACL entry. An ACL entry is represented by two or three colon-separated( : ) fields.

*user*:[*uid*]:*perms*      If *uid* blank, it represents the file owner.

*group*:[*gid*]:*perms*      If *gid* is blank, it represents the owning group.

*other*:*perms*              Represents the file other class.

*mask*:*perms*              Defines the MAX permission to hand out.

For example to give user joe read and write permissions, the ACL entry is specified as:

```
user:joe:rw-
```

NFSv4 ACLs NFSv4 ACL model is based loosely on the Windows NT ACL model. NFSv4 ACLs provide a much richer ACL model than POSIX-draft ACLs.

The major differences between NFSv4 and POSIX-draft ACLs are as follows:

- NFSv4 ACLs provide finer grained permissions than the rwx model.
- NFSv4 ACLs allow for both ALLOW and DENY entries.
- NFSv4 ACLs provide a rich set of inheritance semantics. POSIX ACLs also have inheritance, but with the NFSv4 model you can control the following inheritance features:
  - Whether inheritance cascades to both files and directories or only to files or directories.
  - In the case of directories, you can indicate whether inheritance is applied to the directory itself, to just one level of subdirectories, or cascades to all subdirectories of the directory.
- NFSv4 ACLs provide a mechanism for hooking into a system's audit trail. Currently, Solaris does not support this mechanism.
- NFSv4 ACLs enable administrators to specify the order in which ACL entries are checked. With POSIX-draft ACLs the file system reorders ACL entries into a well defined, strict access, checking order.

POSIX-draft ACL semantics can be achieved with NFSv4 ACLs. However, only some NFSv4 ACLs can be translated to equivalent POSIX-draft ACLs.

Permissions can be specified in three different chmod ACL formats: verbose, compact, or positional. The verbose format uses words to indicate that the permissions are separated with a forward slash (/) character. Compact format uses the permission letters and positional format uses the permission letters or the hyphen (-) to identify no permissions.

The permissions for verbose mode and their abbreviated form in parentheses for compact and positional mode are described as follows:

read_data (r)	Permission to read the data of the file
list_directory (r)	Permission to list the contents of a directory.
write_data (w)	Permission to modify a file's data anywhere in the file's offset range. This includes the ability to grow the file or write to any arbitrary offset.
add_file (w)	Permission to add a new file to a directory.
append_data (p)	The ability to modify the file's data, but only starting at EOF. Currently, this permission is not supported.
add_subdirectory (p)	Permission to create a subdirectory to a directory.

read_xattr (R)	The ability to read the extended attributes of a file or do a lookup in the extended attributes directory.
write_xattr (w)	The ability to create extended attributes or write to the extended attributes directory.
execute (x)	Permission to execute a file.
read_attributes (a)	The ability to read basic attributes (non-ACLs) of a file. Basic attributes are considered to be the stat level attributes. Allowing this access mask bit means that the entity can execute <code>ls(1)</code> and <code>stat(2)</code> .
write_attributes (A)	Permission to change the times associated with a file or directory to an arbitrary value.
delete (d)	Permission to delete the file.
delete_child (D)	Permission to delete a file within a directory.
read_acl (c)	Permission to read the ACL.
write_acl (C)	Permission to write the ACL or the ability to execute <code>chmod(1)</code> or <code>setfacl(1)</code> .
write_owner (o)	Permission to change the owner or the ability to execute <code>chown(1)</code> or <code>chgrp(1)</code> .
synchronize (s)	Permission to access a file locally at the server with synchronous reads and writes. Currently, this permission is not supported.

The following inheritance flags are supported by NFSv4:

file_inherit (f)	Inherit to all newly created files in a directory.
dir_inherit (d)	Inherit to all newly created directories in a directory.
inherit_only (i)	Placed on a directory, but does not apply to the directory itself, only to newly created files and directories. This flag requires <code>file_inherit</code> and or <code>dir_inherit</code> to indicate what to inherit.
no_propagate (n)	Placed on directories and indicates that ACL entries should only be inherited one level of the tree. This flag requires <code>file_inherit</code> and or <code>dir_inherit</code> to indicate what to inherit.
successful_access (S)	Indicates if an alarm or audit record should be initiated upon successful accesses. Used with audit/alarm ACE types.
failed_access (F)	Indicates if an alarm or audit record should be initiated when access fails. Used with audit/alarm ACE types.
inherited (I)	ACE was inherited.
-	No permission granted.

An NFSv4 ACL is expressed using the following syntax:

```
acl_entry[,acl_entry]...
```

```
owner@:<perms>[:inheritance flags]:<allow|deny>
group@:<perms>[:inheritance flags]:<allow|deny>
everyone@:<perms>[:inheritance flags]:<allow|deny>
user:<username>[:inheritance flags]:<allow|deny>
group:<groupname>[:inheritance flags]:<allow|deny>
```

```
owner@    File owner
group@    Group owner
user      Permissions for a specific user
group     Permissions for a specific group
```

Permission and inheritance flags are separated by a / character.

ACL specification examples:

```
user:fred:read_data/write_data/read_attributes:file_inherit:allow
owner@:read_data:allow,group@:read_data:allow,user:tom:read_data:deny
```

Using the compact ACL format, permissions are specified by using 14 unique letters to indicate permissions.

Using the positional ACL format, permissions are specified as positional arguments similar to the `ls -V` format. The hyphen (-), which indicates that no permission is granted at that position, can be omitted and only the required letters have to be specified.

The letters above are listed in the order they would be specified in positional notation.

With these letters you can specify permissions in the following equivalent ways.

```
user:fred:rw-----R-----:file_inherit:allow
```

Or you can remove the - and scrunch it together.

```
user:fred:rwR:file_inherit:allow
```

The inheritance flags can also be specified in a more compact manner, as follows:

```
user:fred:rwR:f:allow
user:fred:rwR:f-----:allow
```

Shell-level Solaris API The Solaris command interface supports the manipulation of ACLs. The following Solaris utilities accommodate both ACL models:

<code>chmod</code>	The <code>chmod</code> utility has been enhanced to allow for the setting and deleting of ACLs. This is achieved by extending the symbolic-mode argument to support ACL manipulation. See <a href="#">chmod(1)</a> for details.
<code>compress</code>	When a file is compressed any ACL associated with the original file is preserved with the compressed file.
<code>cp</code>	By default, <code>cp</code> ignores ACLs, unless the <code>-p</code> option is specified. When <code>-p</code> is specified the owner and group id, permission modes, modification and access times, ACLs, and extended attributes if applicable are preserved.
<code>cpio</code>	ACLs are preserved when the <code>-P</code> option is specified.
<code>find</code>	<code>find</code> locates files with ACLs when the <code>-acl</code> flag is specified.
<code>ls</code>	By default <code>ls</code> does not display ACL information. When the <code>-v</code> option is specified, a file's ACL is displayed.
<code>mv</code>	When a file is moved, all attributes are carried along with the renamed file. When a file is moved across a file system boundary, the ACLs are replicated. If the ACL information cannot be replicated, the move fails and the source file is not removed.
<code>pack</code>	When a file is packed, any ACL associated with the original file is preserved with the packed file.
<code>rcp</code>	<code>rcp</code> has been enhanced to support copying. A file's ACL is only preserved when the remote host supports ACLs.
<code>tar</code>	ACLs are preserved when the <code>-p</code> option is specified.
<code>unpack</code>	When a file with an ACL is unpacked, the unpacked file retains the ACL information.

Application-level API The primary interfaces required to access file system ACLs at the programmatic level are the `acl_get()` and `acl_set()` functions. These functions support both POSIX draft ACLs and NFSv4 ACLs.

Retrieving a file's ACL 

```
int acl_get(const char *path, int flag, acl_t **aclp);
int facl_get(int fd, int flag, acl_t **aclp);
```

The `acl_get(3SEC)` and `facl_get(3SEC)` functions retrieves an ACL on a file whose name is given by `path` or referenced by the open file descriptor `fd`. The `flag` argument specifies whether a trivial ACL should be retrieved. When the `flag` argument equals `ACL_NO_TRIVIAL` then only ACLs that are not trivial are retrieved. The ACL is returned in the `aclp` argument.

Freeing ACL structure 

```
void acl_free(acl_t *aclp);
```

The `acl_free()` function frees up memory allocated for the argument `aclp`.

Setting an ACL on a file `int acl_set(const char *path, acl_t *aclp);`  
`int facl_set(int fd, acl_t *aclp);`

The `acl_set(3SEC)` and `facl_get(3SEC)` functions are used for setting an ACL on a file whose name is given by path or referenced by the open file descriptor `fd`. The `aclp` argument specifies the ACL to set. The `acl_set(3SEC)` translates an POSIX-draft ACL into a NFSv4 ACL when the target file systems supports NFSv4 ACLs. No translation is performed when trying to set an NFSv4 ACL on a POSIX-draft ACL supported file system.

Determining an ACL's trivialness `int acl_trivial(const char *path);`

The `acl_trivial()` function is used to determine whether a file has a trivial ACL. Whether an ACL is trivial depends on the type of the ACL. A POSIX draft ACL is nontrivial if it has greater than `MIN_ACL_ENTRIES`. An NFSv4/ZFS-style ACL is nontrivial if it either has entries other than `owner@`, `group@`, and `everyone@`, has inheritance flags set, or is not ordered in a manner that meets POSIX access control requirements.

Removing all ACLs from a file `int acl_strip(const char *path, uid_t uid, gid_t gid, mode_t mode);`

The `acl_strip()` function removes all ACLs from a file and replaces them with a trivial ACL based off of the passed in argument `mode`. After replacing the ACL the owner and group of the file are set to the values specified in the `uid` and `gid` parameters.

Converting ACLs to/from external representation `int acl_fromtext(const char *path, acl_t **aclp);`  
`char *acl_totext(acl_t *aclp, int flags);`

The `acl_text()` function converts an internal ACL representation pointed to by `aclp` into an external representation. See `DESCRIPTION` for details about external representation.

The `acl_fromtext()` functions converts and external representation into an internal representation. See `DESCRIPTION` for details about external representation.

**Examples** The following examples demonstrate how the API can be used to perform basic operations on ACLs.

**EXAMPLE 1** Retrieving and Setting an ACL

Use the following to retrieve an ACL and set it on another file:

```
error = acl_get("file", ACL_NO_TRIVIAL, &aclp);

if (error == 0 && aclp != NULL) {
    error = acl_set("file2", aclp)
    acl_free(aclp);
}
...
```

**EXAMPLE 2** Retrieving and Setting Any ACLs

Use the following to retrieve any ACL, including trivial ACLs, and set it on another file:

```
error = acl_get("file3", 0, &aclp);
if (error == 0) {
    error = acl_set("file4", aclp)
    acl_free(aclp);
}
...
```

**EXAMPLE 3** Determining if a File has a Trivial ACL

Use the following to determine if a file has a trivial ACL:

```
istrivial = acl_trivial("file")

if (istrivial == 0)
    printf("file %s has a trivial ACL\n", file);
else
    printf("file %s has a NON-trivial ACL\n", file);
...
```

**EXAMPLE 4** Removing all ACLs from a File

Use the following to remove all ACLs from a file, and set a new mode, owner, and group:

```
error = acl_strip("file", 10, 100, 0644);
...
```

**See Also** [chgrp\(1\)](#), [chmod\(1\)](#), [chown\(1\)](#), [cp\(1\)](#), [cpio\(1\)](#), [find\(1\)](#), [ls\(1\)](#), [mv\(1\)](#), [tar\(1\)](#), [setfacl\(1\)](#), [chmod\(2\)](#), [acl\(2\)](#), [stat\(2\)](#), [acl\\_get\(3SEC\)](#), [aclsort\(3SEC\)](#), [acl\\_fromtext\(3SEC\)](#), [acl\\_free\(3SEC\)](#), [acl\\_strip\(3SEC\)](#), [acl\\_trivial\(3SEC\)](#)

**Name** ad – Active Directory as a naming repository

**Description** Solaris clients can obtain naming information from Active Directory (AD) servers.

The Solaris system must first join an AD domain and then add the `ad` keyword to the appropriate entries in the `nsswitch.conf(4)` file. The Solaris system joins the AD domain by using the `kclient(1M)` utility. The AD name service only supports the naming databases for `passwd` and `group`.

Windows users are not able to log in. The `user_attr(4)` database has no entries for Windows users, and the `passwd(1)` command does not support the synchronization of user passwords with AD.

The Solaris AD client uses auto-discovery techniques to find AD directory servers, such as domain controllers and global catalog servers. The client also uses the LDAP v3 protocol to access naming information from AD servers. The AD server schema requires no modification because the AD client works with native AD schema. The Solaris AD client uses the `idmap(1M)` service to map between Windows security identifiers (SIDs) and Solaris user identifiers (UIDs) and group identifiers (GIDs). User names and group names are taken from the `sAMAccountName` attribute of the AD user and group objects and then tagged with the domain where the objects reside. The domain name is separated from the user name or group name by the `@` character.

The client uses the SASL/GSSAPI/KRB5 security model. The `kclient` utility is used to join the client to AD. During the join operation, `kclient` configures Kerberos v5 on the client. See `kclient(1M)`.

<b>Files</b>	<code>/etc/nsswitch.conf</code>	Configuration file for the name-service switch.
	<code>/etc/nsswitch.ad</code>	Sample configuration file for the name-service switch configured with <code>ad</code> , <code>dns</code> and <code>files</code> .
	<code>/usr/lib/nss_ad.so.1</code>	Name service switch module for AD.

**See Also** `passwd(1)`, `svcs(1)`, `idmap(1M)`, `idmapd(1M)`, `kclient(1M)`, `svcadm(1M)`, `svccfg(1M)`, `svccfg(1M)`, `nsswitch.conf(4)`, `user_attr(4)`, `smf(5)`

**Name** ars – receive Solaris audit logs from a remote server

**Synopsis** /usr/sbin/auditd

**Description** Audit Remote Server (ARS) is the counterpart of the [audit\\_remote\(5\)](#) plugin. Data sent by the plugin can be captured, processed, and stored by the server according to the its configuration.

ARS is delivered as a disabled Solaris audit component. It is necessary to configure ARS before it can be used to process a remote audit trail. ARS configuration is twofold:

- the underlying security mechanisms used for secure audit data transport has to be configured (see [audit\\_remote\(5\)](#));
- the audit remote subsystem has to be configured.

To observe and configure the ARS, use the [auditconfig\(1M\)](#) -setremote and -getremote options. The configuration is divided between the configuration of *server* and *group*. The *server* configuration allows for changing common ARS parameters, while the *group* keyword allows configuration of connection groups, the sets of hosts sharing the same local storage parameters.

Server configuration attributes *listen\_address*  
The address the server listens on. An empty *listen\_address* attribute defaults to listen on all local addresses.

*listen\_port*  
The local listening port; 0 defaults to 16162, the port associated with the “solaris-audit” Internet service name. See [services\(4\)](#).

*login\_grace\_time*  
The server disconnects after login grace time (in seconds) if the connection has not been successfully established; 0 defaults to no limit.

*max\_startups*  
The number of concurrent unauthenticated connections to the server at which the server starts refusing new connections. The value might be specified in *begin:rate:full* format to allow random early drop mode, for example 10:30:60, meaning that ARS would refuse connection attempts with a probability of *rate*/100 (30% in our example) if there are currently 10 (from the *start* field) unauthenticated connections. The probability increases linearly and all connection attempts are refused if the number of unauthenticated connections reaches *full* (60 in our example).

Group configuration attributes The *binfile\_dir*, *binfile\_fsize*, and *binfile\_minfree* attributes follow the respective *p\_\** attributes defined in [audit\\_binfile\(5\)](#). Brief descriptions follow.

*binfile\_dir*  
The directory for storing per host audit data.

*binfile\_fsize*

The maximum size of each of the stored audit trail files; 0 defaults to no limit.

*binfile\_minfree*

The minimum free space on file system with *binfile\_dir* before the *audit\_binfile* informs the administrator via *audit\_warn(1M)*; 0 defaults to no limit.

*hosts*

The hosts in the given connection group allowed to send audit data to server. A comma is a delimiter in case of multiple host entries. If *hosts* is empty, such connection group is called a wild card connection group. If a new connection cannot be classified to any other (non-wild card) connection group and there is an active wild card connection group configured, the new connection is classified to that connection group. Only one active wild card connection group can be configured.

For comprehensive configuration description and examples, see the section on Auditing in *Oracle Solaris 11.1 Administration: Security Services*.

**Examples** EXAMPLE 1 Audit Remote Server configuration

The following example describes steps to configure audit remote server to listen on specific address. One wild card and one non-wild card connection group will be created. The non-wild card connection group configuration will address remote audit data from *tic.cz.example.com* and *tac.us.example.com*, the trail will be stored in */var/audit/remote*.

```
# Print the current audit remote server configuration.
# Both server and connection groups (if any) is displayed.
auditconfig -getremote

# Set address the audit remote server will listen on.
auditconfig -setremote server "listen_address=192.168.0.1"

# Create two connection groups. Note that by default the
# connection group is created with no hosts specified
# (wild card connection group).
auditconfig -setremote group create clockhouse
auditconfig -setremote group create sink

# Add hosts to the connection group (convert the wild card
# connection group no non-wild card one). Set the storage
# directory and activate the connection group.
auditconfig -setremote group active clockhouse \
    "hosts=tic.cz.example.com,tac.us.example.com,
    binfile_dir=/var/audit/remote"

# Activate the wild card connection group.
auditconfig -setremote group active sink
```

**EXAMPLE 1** Audit Remote Server configuration (Continued)

```
# Verify the audit remote server configuration.
auditconfig -getremote

# Start or refresh the audit service.
audit -s
```

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

**See Also** [audit\(1M\)](#), [auditconfig\(1M\)](#), [auditd\(1M\)](#), [audit\\_warn\(1M\)](#), [services\(4\)](#), [attributes\(5\)](#), [audit\\_binfile\(5\)](#), [smf\(5\)](#)

See the section on Auditing in *Oracle Solaris 11.1 Administration: Security Services*.

**Notes** The audit service FMRI is `svc:/system/auditd:default`.

**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
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.

**Name** attributes, architecture, availability, CSI, stability, MT-Level, standard – attributes of interfaces

**Description** The ATTRIBUTES section of a manual page contains a table defining attribute types and their corresponding values. The following is an example of an attributes table. Not all attribute types are appropriate for all types of interfaces.

ATTRIBUTETYPE	ATTRIBUTEVALUE
Architecture	SPARC
Availability	system/kernel
CSI	Enabled
Interface Stability	Committed
MT-Level	Safe
Standard	See <a href="#">standards(5)</a> .

**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 [pkg\(1\)](#).

**Code Set Independence (CSI)** OS utilities and libraries 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, 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.

For practical reasons, however, 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 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** Oracle Solaris 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 Oracle Solaris will endeavor to ensure that these continue to work in future minor releases. Applications that depend only on Committed 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 Oracle Solaris 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.

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 Committed interfaces. Initial product releases are usually 1.0.
Minor	x.y	Compared to an x.0 or earlier release (y!=0), it is likely to contain: feature additions, compatible changes to Committed interfaces, or likely incompatible changes to Uncommitted or Volatile 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. Incompatible changes to Volatile interfaces are possible.

In the context of interface stability, update releases (occasionally referred to as patch releases) should be considered equivalent to Micro Releases.

### 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
Committed	Major (x.0)	Incompatibilities are exceptional.
Uncommitted	Minor (x.y)	Incompatibilities are common.
Volatile	Micro (x.y.z)	Incompatibilities are common.

The interface stability level classifications described on this manual page apply to both source and binary interfaces unless otherwise stated. All stability level classifications are public, with the exception of the `Private` classification. The precise stability level of a public interface (one that is documented in the manual pages) is unspecified unless explicitly stated. The stability level of an undocumented interface is implicitly `Private`.

The existence of documentation other than the documentation that is a component of the Oracle Solaris product should not be construed to imply any level of stability for interfaces provided by the Oracle Solaris product. The only source of stability level information is the reference manual pages.

### Committed

The intention of a Committed interface is to enable third parties to develop applications to these interfaces, release them, and have confidence that they will run on all releases of the product after the one in which the interface was introduced, and within the same Major release. Even at a Major release, incompatible changes are expected to be rare, and to have strong justifications.

Interfaces defined and controlled as industry standards are most often treated as Committed interfaces. In this case, the controlling body and/or public, versioned document is typically noted in a “Standard” entry in the Attributes table or elsewhere in the documentation.

Although a truly exceptional event, incompatible changes are possible in any release if the associated defect is serious enough as outlined in the Exceptions section of this document or in a Minor release by following the End of Feature process. If support of a Committed interface must be discontinued, Oracle Solaris will attempt to provide notification and the stability level will be marked Obsolete.

### Uncommitted

No commitment is made about either source or binary compatibility of these interfaces from one Minor release to the next. Even the drastic incompatible change of removal of the interface in a Minor release is possible. Uncommitted interfaces are generally not appropriate for use by release-independent products.

Incompatible changes to the interface are intended to be motivated by true improvement to the interface which may include ease of use considerations. The general expectation should be that Uncommitted interfaces are not likely to change incompatibly and if such changes occur they will be small in impact and may often have a mitigation plan.

Uncommitted interfaces generally fall into one of the following subcategories:

1. Interfaces that are experimental or transitional. They are typically used to give outside developers early access to new or rapidly changing technology, or to provide an interim solution to a problem where a more general solution is anticipated.
2. Interfaces whose specification is controlled by an outside body yet Oracle Solaris expects to make a reasonable effort to maintain compatibility with previous releases until the next Minor release at which time Oracle Solaris expects to synchronize with the external specification.
3. Interfaces whose target audience values innovation (and possibly ease of use) over stability. This attribute is often associated with administrative interfaces for higher tier components.

For Uncommitted interfaces, Oracle Solaris makes 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.

## Volatile

Volatile interfaces can change at any time and for any reason.

The Volatile interface stability level allows Oracle Solaris products to quickly track a fluid, rapidly evolving specification. In many cases, this is preferred to providing additional stability to the interface, as it may better meet the expectations of the consumer.

The most common application of this taxonomy level is to interfaces that are controlled by a body other than Oracle Solaris, but unlike specifications controlled by standards bodies or Free or Open Source Software (FOSS) communities which value interface compatibility, it can not be asserted that an incompatible change to the interface specification would be exceedingly rare. It may also be applied to FOSS controlled software where it is deemed more important to track the community with minimal latency than to provide stability to our customers.

It also common to apply the Volatile classification level to interfaces in the process of being defined by trusted or widely accepted organization. These are generically referred to as draft standards. An “IETF Internet draft” is a well understood example of a specification under development.

Volatile can also be applied to experimental interfaces.

No assertion is made regarding either source or binary compatibility of Volatile interfaces between any two releases, including patches. Applications containing these interfaces might fail to function properly in any future release.

## Not-an-Interface

The situation occasionally occurs where there exists an entity that could be inferred to be an interface, but actually is not. Common examples are output from CLIs intended only for human consumption and the exact layout of a GUI.

This classification is a convenience term to be used to clarify such situations where such confusion is identified as likely. Failure to apply this term to an entity is not an indication that the entity is some form of interface. It only indicates that the potential for confusion was not identified.

## Private

A Private interface is an interface provided by a component (or product) intended only for the use of that component. A Private interface might still be visible to or accessible by other components. Because the use of interfaces private to another component carries great stability risks, such use is explicitly not supported. Components not supplied by Oracle Solaris should not use Private interfaces.

Most Private interfaces are not documented. It is an exceptional case when a Private interface is documented. Reasons for documenting a Private interface include, but are not limited to, the intention that the interface might be reclassified to one of the public stability level classifications in the future or the fact that the interface is inordinately visible.

## Obsolete

Obsolete is a modifier that can appear in conjunction with the above classification levels. The Obsolete modifier indicates an interface that is “deprecated” and/or no longer advised for general use. An existing interface may be downgraded from some other status (such as Committed or Uncommitted) by the application of the Obsolete modifier to encourage customers to migrate from that interface before it may be removed (or incompatibly changed).

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, Oracle Solaris will attempt to provide notification before discontinuing support. Use of an Obsolete interface may produce warning messages.

## Exceptions

There are rare instances when it is in the best interest of both Oracle Solaris and the customer to break the interface stability commitment. The following list contains the common, known reasons for the interface provider to violate an interface stability commitment, but does not preclude others.

1. Security holes where the vulnerability is inherent in the interface.
2. Data corruption where the vulnerability is inherent in the interface.
3. Standards violations uncovered by a change in interpretation or enhancement of conformance tests.
4. An interface specification which isn't controlled by Oracle Solaris has been changed incompatibly and the vast majority of interface consumers expect the newer interface.
5. Not making the incompatible change would be incomprehensible to our customers. One example of this would be to have not incompatibly changed pcfs when the DOS 8.3 naming restrictions were abandoned.

Incompatible changes allowed by exception will always be delivered in the “most major” release vehicle possible. However, often the consequences of the vulnerabilities or contractual branding requirements will force delivery in a patch.

## Compatibility with Earlier Interface Classification Schemes

In releases up to and including Solaris 10, a different interface classification scheme was used. The following table summarizes the mapping between the old and new classification schemes.

Old	New	Comments
Standard	Committed	An entry in the attributes table for the Standard attribute type should appear.
Stable	Committed	Name change.
Evolving	Uncommitted	Actual commitments match.

Old	New	Comments
Unstable	Uncommitted	Name change.
External	Volatile	Name change with expansion of allowed usage.
Obsolete	(Obsolete)	Was a classification, now a modifier.

The increased importance of Free or Open Source Software motivated the name change from Stable/Unstable to Committed/Uncommitted. Stable conflicted with the common use of the term in FOSS communities.

Ambiguity in the definition of Evolving was causing difficulty in interpretation. As part of the migration to the new classification scheme, many formerly Evolving interfaces were upgraded to Committed. However, upon encountering the term Evolving, Uncommitted should be inferred.

MT-Level Libraries are classified into categories that define their ability to support multiple threads. Manual pages containing functions that are of multiple or differing levels describe this in 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 multithreaded 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 multithreaded 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 time to execute within the library. Unsafe libraries might contain functions that are Safe; however, most of the library's functions are unsafe to call. Some functions that are Unsafe have reentrant counterparts that are MT-Safe. Reentrant functions are designated by the `_r` suffix appended to the function name.

#### 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. 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 functions that can be safely called from a signal handler. A thread that is executing an Async-Signal-Safe function will not deadlock with itself if interrupted by a signal. Signals are only a problem for MT-Safe functions that acquire locks.

Async-Signal-Safe functions are also MT-Safe. Signals are disabled when locks are acquired in Async-Signal-Safe functions. These signals prevent a signal handler that might acquire the same lock from being called.

#### 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.

#### Fork-Safe

The `fork(2)` function replicates only the calling thread in the child process. The `fork1(2)` function exists for compatibility with the past and is synonymous with `fork()`. If a thread other than the one performing the fork holds a lock when `fork()` is called, the lock will still be held in the child process but there will be no lock owner since the owning thread was not replicated. A child calling a function that attempts to acquire the lock will deadlock itself.

When `fork()` is called, a Fork-Safe library arranges to have all of its internal locks held only by the thread performing the fork. This is usually accomplished with `pthread_atfork(3C)`, which is called when the library is initialized.

The `forkall(2)` function provides the capability for the rare case when a process needs to replicate all of its threads when performing a fork. No `pthread_atfork()` actions are performed when `forkall()` is called. There are dangers associated with calling `forkall()`. If some threads in a process are performing I/O operations when another thread calls `forkall()`, they will continue performing the same I/O operations in both the parent and child processes, possibly causing data corruption. For this and other race-condition reasons, the use of `forkall()` is discouraged.

In all Solaris releases prior to Solaris 10, the behavior of `fork()` depended on whether or not the application was linked with `-lpthread` (POSIX threads, see [standards\(5\)](#)). If linked with `-lpthread`, `fork()` behaved like `fork1()`; otherwise it behaved like `forkall()`. To avoid any confusion concerning the behavior of `fork()`, applications can specifically call `fork1()` or `forkall()` as appropriate.

### Cancel-Safety

If a multithreaded application uses `pthread_cancel(3C)` 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(3C)`), 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(3C)` 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 multithreaded 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_ASYNCCHRONOUS`. 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 might never be necessary to worry about asynchronous cancel safety. 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.

**Standard** Many interfaces are defined and controlled as industry standards. When this is the case, the controlling body and/or public, versioned document is noted in this section.

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 interfaces based upon a public standard. When the standard or specification allows alternative implementation choices, the manual page usually only describes the alternative implemented by Oracle Solaris. The manual page also describes any compatible extensions to the base definition of Standard interfaces provided by Oracle Solaris.

No endorsement of the referenced controlling body or document should be inferred by its presence as a “Standard” entry. The controlling body may be a very formal organization, as in ISO or ANSI, a less formal, but generally accepted organization such as IETF, or as informal as the sole contributor in the case of FOSS (Free or Open Source Software).

**See Also** [uname\(1\)](#), [Intro\(3\)](#), [standards\(5\)](#)

[pkg\(1\)](#)

**Name** audit\_binfile – generation of Solaris audit logs

**Synopsis** /usr/lib/security/audit\_binfile.so

**Description** The audit\_binfile plugin module for Solaris audit, /usr/lib/security/audit\_binfile.so, writes binary audit data to files as configured in [auditconfig\(1M\)](#); it is the default plugin for the Solaris audit daemon [auditd\(1M\)](#). Its output is described by [audit.log\(4\)](#).

The audit\_binfile plugin is loaded by auditd if the plugin is configured as an active via auditconfig. Use the auditconfig -setplugin option to change all the plugin related configuration parameters.

**Object Attributes** The following attributes specify the configuration of audit\_binfile plugin:

p\_dir

```
dir1[,dir2],... [,dirn]
```

A list of directories, where the audit files will be created. Any valid writable directory can be specified.

p\_minfree

A percentage, which indicates the amount of free space required on the target p\_dir. If free space falls below this threshold, the audit daemon [auditd\(1M\)](#) invokes the shell script [audit\\_warn\(1M\)](#). If no threshold is specified, the default is 1%.

p\_fsize

The p\_fsize attribute defines the maximum size that an audit file can become before it is automatically closed and a new audit file is opened. This is equivalent to an administrator issuing an `audit -n` command when the audit file size equals the value specified by the administrator. The default size is zero (0), which allows the file to grow without bound. The value specified must be higher than 500KB and lower than 16 exabytes (EB). The used file system might further lower the limits. The format of the p\_fsize value can be specified as an exact value in bytes or in a human-readable form with a suffix of B, K, M, G, T, P, E, Z (for bytes, kilobytes, megabytes, gigabytes, terabytes, petabytes, exabytes, or zettabytes, respectively). Suffixes of KB, MB, GB, TB, PB, EB, and ZB are also accepted.

**Examples** The following directives cause audit\_binfile.so to be loaded, specify the directories for writing audit logs, and specify the percentage of required free space per directory.

```
auditconfig -setplugin audit_binfile active \  
"p_dir=/var/audit/jedgar/eggplant,/var/audit/jedgar.aux/eggplant,  
/var/audit/global/eggplant;p_minfree=20;p_fsize=4.5GB"
```

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	MT-Safe
Interface Stability	Committed

**See Also** [auditconfig\(1M\)](#), [auditd\(1M\)](#), [audit\\_warn\(1M\)](#), [syslog.conf\(4\)](#), [attributes\(5\)](#)

*Oracle Solaris 11.1 Administration: Security Services*

**Name** audit\_flags – audit preselection flags

**Description** Audit flags specify which audit classes are to be audited for a process. Audit classes are defined in the [audit\\_class\(4\)](#) file and group together like audit events as defined in the [audit\\_event\(4\)](#) file. The default Solaris system-wide audit flags are configured as part of the audit service using [auditconfig\(1M\)](#). Additional per-user or per-role audit flags may be configured in the [user\\_attr\(4\)](#) database or in the profiles granted to the user by the *audit\_flags=always-audit-flags:never-audit-flags* keyword. The audit flags of a process are called the preselection mask. The preselection mask is set at login and role assumption time by combining the default Solaris system-wide audit flags with the per-user audit flags (default flags + *always-audit-flags*) - *never-audit-flags*.

Audit flags are specified as a character string representing the audit class names to be audited. Each flag identifies an audit class and is separated by a comma (,) from others in the string. An audit class name preceded by - means that the class should be audited for failure only; successful attempts are not audited. An audit class name preceded by + means that the class should be audited for success only; failed attempts are not audited. Without a prefix, the audit class name indicates that the class is to be audited for both successes and failures. The special string “all” indicates that all audit events are to be audited; -all indicates that all failed attempts are to be audited and +all indicates that all successful attempts are to be audited. The prefixes ^, ^- and ^+ turn off flags specified earlier in the string (^- and ^+ for failed and successful attempts respectively, ^ for both). They are typically used to reset flags. The special string no indicates no audit events are to be audited.

**Examples** **EXAMPLE 1** Preselect to audit for successful and failed “lo” (login/logout), “am” (administration) audit events and all failed audit events except for failed “fm” (file attribute modify) events.

```
lo,am,-all,^-fm
```

**EXAMPLE 2** Preselect to audit for successful and failed “lo” (login/logout), “as” (system-wide administration) and failed “fm” (file attribute modify) events.

```
lo,as,-fm
```

**See Also** [profiles\(1\)](#), [auditconfig\(1M\)](#), [auditd\(1M\)](#), [usermod\(1M\)](#), [audit\\_class\(4\)](#), [audit\\_event\(4\)](#), [prof\\_attr\(4\)](#), [user\\_attr\(4\)](#)

*Oracle Solaris 11.1 Administration: Security Services*

- 
- Name** audit\_remote – send Solaris audit logs to a remote server
- Synopsis** /usr/lib/security/audit\_remote.so
- Description** The audit\_remote plugin module for Solaris audit, /usr/lib/security/audit\_remote.so, sends binary audit records ([audit.log\(4\)](#)) to audit servers as they are configured with [auditconfig\(1M\)](#).
- The audit\_remote plugin is loaded by [auditd\(1M\)](#) if the plugin is configured as an active via [auditconfig](#). Use the `auditconfig -setplugin` option to change all the plugin related configuration parameters.
- The Solaris audit service daemon's audit remote service, [ars\(5\)](#), may be configured with [auditconfig](#) to receive the binary audit records send by `audit_remote`.
- Object Attributes** The following attributes specify the configuration of `audit_remote` plugin:
- p\_hosts**
- ```
host1[:[port1][:mech1]][,host2[:[port2][:mech2]],... \
    hostn[:[portn][:mechn]]]
```
- A list of audit hosts/servers. Audit records are sent to the first available host. If a host is unreachable or a timeout occurs while sending data, the next host in the list is tried. If connection to all hosts fails, the list is tried again from the beginning.
- The *host* part of a `p_hosts` entry can be in any form acceptable to [getipnodebyname\(3SOCKET\)](#).
- The *port* part of a `p_hosts` entry is the port on host that is contacted to initiate an audit server connection. If not specified, the port number is that assigned to the `solaris-audit` service. See [getservbyname\(3XNET\)](#).
- The *mech* part of a `p_host` entry is the GSS-API mechanism name ([mech\(4\)](#)). If not specified, the local host's default mechanism is used. The recommended mechanism is `kerberos_v5`.
- p\_retries**
- The number of retries for connecting to and sending data to a server.
- The default value is 3.
- p\_timeout**
- The number of seconds in which a connection/sending data timeouts.
- The default value is 5 seconds.
- qsize**
- The maximum number of outstanding audit records to keep.
- The default is the value of the kernel queue control high water mark. See [auditconfig\(1M\)](#).

If set to 0, the default value is the value of the kernel queue control high water mark. See [auditconfig\(1M\)](#).

**GSSSESSION** The `audit_remote` plugin is a TCP client that authenticates configured audit servers using the GSS-API ([libgss\(3LIB\)](#)). Binary Solaris Audit records are sent with integrity and confidentiality protection as per-message tokens generated by [gss\\_wrap\(3GSS\)](#).

The plugin initiates a TCP connection to an audit server (*host:port:mech*) and establishes a GSS security context (with [gss\\_init\\_sec\\_context\(3GSS\)](#)), with appropriate security mechanism ([mech\(4\)](#)).

If no port is specified, the service name `solaris-audit` is looked up to obtain a TCP port number. If no mechanism is specified, the `GSS_C_NO_OID` is used as a `mech_type` parameter of [gss\\_init\\_sec\\_context\(3GSS\)](#), and causes the underlying GSS-API to use the local default mechanism.

[gss\\_init\\_sec\\_context\(3GSS\)](#) uses `GSS_C_NO_CREDENTIAL` as the initiator credential handle and a target name of the form `audit@<host_fqdn>`. The server is expected to use [gss\\_accept\\_sec\\_context\(3GSS\)](#) to complete the context establishment.

Once the security context is established, the client (`audit_remote` plugin) calls [gss\\_wrap\(3GSS\)](#) to achieve the confidentiality of the transferred payload - the audit records. The server is expected to use [gss\\_unwrap\(3GSS\)](#) to unwrap the received data and [gss\\_get\\_mic\(3GSS\)](#) to obtain the MIC (Message Integrity Code) to be later sent back to the plugin as a message retrieval acknowledgment.

For example, if the `kerberos_v5` mechanism is configured as `GSS_API` mechanism on the client and both sides agree on using this mechanism, the client side has to be eligible to non-interactively gain session keys for the `audit/<host_fqdn>@<REALM>` principal from the Kerberos KDC/TGS. At the same time the identity running the audit server application has to have the long term keys associated with the `audit/<host_fqdn>@<REALM>` principal stored in the keytab file ([krb5.conf\(4\)](#)) to be able to decrypt the session keys.

The `audit_remote` plugin initiates a connection to first server in the `p_hosts` list. If the connection fails or audit record sends are not responded to in `p_timeout` seconds, after `p_retries` attempts the plugin tries to connect to the next server. If the connection to the last server fails, the plugin retries to connect to the first host in the list. [audit\\_warn\(1M\)](#) is executed at every unsuccessful attempt to connect to the server or send timeout with the plugin option `plugin audit_remote.so retry <count> <error>.<error>` is connection `<host:port> <the network error>`. An `EPROTO` network error indicates that the client plugin did not get a successful protocol version handshake.

**PROTOCOL DESCRIPTION** All protocol messages are preceded by the 4 octets of the size of the data to follow. This size is in network byte order.

The protocol begins with version negotiation followed by a GSS-API security context token exchange. On error the connection is closed (and any output token optionally sent).

The version negotiation takes place in the clear with the plugin sending an octet array of the comma (,) separated list of versions supported. The current version number is the characters 01. The receiver is expected to respond with the version that they accept (in the current case that is the characters 01). A mismatch is considered an error and the connection is closed.

The version octet array sent by the plugin and the version characters accepted by the receiver are concatenated together to make up the application data field of the channel bindings of the GSS security context establishment.

```
<plugin version characters> || <server accepted version characters>
"||" represents concatenation
```

Subsequent tokens contain a 64 bit sequence number in network byte order and a single audit record (`audit.log(4)`); the client uses confidentiality protection. wrap (64 bit sequence number || audit record)

The server acknowledges the receipt (and is then responsible for any data loss) with the received 64 bit sequence number and a MIC token of the unwrapped 64 bit sequence number and audit record. MIC verification on the client side acknowledges the audit record can be freed and not saved for possible retransmission.

```
64 bit sequence number || mic (64 bit sequence number || audit record)
```

Secure remote audit client/server communication flow:

- 1) Client <--> Server - TCP handshake
- 2) Client <--> Server - protocol version negotiation:
  - a) Client --> Server - send data size - uint32\_t value (2)
  - b) Client --> Server - send clear text message of the versions supported comma separated, e.g., "01,02,03" for versions 1 and 2 and 3. The only version supported at present is "01"
  - c) Client <-- Server - send data size - uint32\_t value (2)
  - d) Client <-- Server - send clear text version selected ("01")
 

```
:no version match; close connection; try next host
```
- 3) Security context initiation:
  - a) Client - Construct channel bindings:
 

```
initiator address type (GSS_C_AF_NULLADDR)
acceptor address type (GSS_C_AF_NULLADDR)
application data value (4 octets "0101")
```
  - b) Client --> Server - send token (data) size - uint32\_t value
  - c) Client --> Server - GSS-API per-context token
  - d) Client <-- Server - send token (data) size
  - e) Client <-- Server - GSS-API per-context token
 

```
:repeat a-e until security context is initialized; if unsuccessful,
```

```
        close connection; try next host

4) Client - transmit thread, when audit record to be sent:
  a) Client --> Server - send data size
  b) Client --> Server - GSS-API per-message token
        wrap (sequence number || audit record)
        :repeat a-b while less than max (qsize) outstanding records

5) Client - receive thread:
  a) Client <-- Server - receive data size - uint32_t value
  b) Client <-- Server - receive sequence number - uint64_t value
  c) Client <-- Server - receive MIC
  d) Client          - MIC verification - OK
  e) Client          - remove particular audit record
                    pointed by the sequence number from the
                    retransmit buffer
:repeat a-e, on error close connection; try next host;
retransmit unacknowledged audit records

6) Server - receive thread:
  a) Client --> Server - receive data size
  b) Client --> Server - GSS-API receive, unwrap, store
        per-message token

7) Server - transmit thread:
  a) Server - MIC generation - message integrity code
        mic (sequence number || audit record)
  b) Client <-- Server - send data size
  c) Client < -- Server - send sequence number
  d) Client <-- Server - send MIC
```

### Examples EXAMPLE 1 Loading audit\_remote.so and Specifying the Remote Audit Servers

The following directives cause `audit_remote.so` to be loaded and specify the remote audit servers to where the audit records are sent. The `kerberos_v5` security mechanism is defined to be used when communicating with the servers.

```
auditconfig -setplugin audit_remote active \  
    "p_timeout=90;p_retries=2;  
    p_hosts=eggplant.eng.sun.com::kerberos_v5,  
    purple.ebay.sun.com:4592:kerberos_v5"
```

### EXAMPLE 2 Using the Configuration of Usage Default Security Mechanism

The following example shows the configuration of usage of default security mechanism. It also shows use of default port on one of the configured servers:

```
auditconfig -setplugin audit_remote active \  
    "p_timeout=10;p_retries=2;  
    p_hosts=jedger.eng.sun.com, jbadams.ebay.sun.com:4592"
```

**EXAMPLE 3** Internal plugin queue size settings

Certain conditions (for instance, high or bursty audit data traffic in combination with slow communication lines between server and client) may cause the number of outstanding audit records queued by the `audit_remote` plugin to reach the maximum number configured. The following example shows how to set the queue size parameter.

```
auditconfig -setplugin audit_remote "" 1000
```

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level            | MT-Safe         |
| Interface Stability | See below.      |

The plugin configuration parameters are Committed. The client/server protocol (version "01") is Contracted Project Private. See [audit.log\(4\)](#) for the audit record format and content stability.

**See Also** [auditd\(1M\)](#), [auditconfig\(1M\)](#), [audit\\_warn\(1M\)](#), [getipnodebyname\(3SOCKET\)](#), [getservbyname\(3XNET\)](#), [gss\\_accept\\_sec\\_context\(3GSS\)](#), [gss\\_get\\_mic\(3GSS\)](#), [gss\\_init\\_sec\\_context\(3GSS\)](#), [gss\\_wrap\(3GSS\)](#), [gss\\_unwrap\(3GSS\)](#), [libgss\(3LIB\)](#), [libsocket\(3LIB\)](#), [audit.log\(4\)](#), [krb5.conf\(4\)](#), [mech\(4\)](#), [ars\(5\)](#), [attributes\(5\)](#), [kerberos\(5\)](#), [tcp\(7P\)](#)

*Oracle Solaris 11.1 Administration: Security Services*

**Notes** `audit_remote` authenticates itself to the remote audit service by way of GSS-API ([libgss\(3LIB\)](#)). Default gss credentials are used as provided by the gss implementation mechanism, such as Kerberos.

The `solaris-audit` service port assigned by IANA is 16162.

**Name** audit\_syslog – realtime conversion of Solaris audit data to syslog messages

**Synopsis** /usr/lib/security/audit\_syslog.so

**Description** The audit\_syslog plugin module for Solaris audit, /usr/lib/security/audit\_syslog.so, provides realtime conversion of Solaris audit data to syslog-formatted (text) data and sends it to a syslog daemon as configured in [syslog.conf\(4\)](#). The plugin's path is specified with the [auditconfig\(1M\)](#) utility.

Messages to syslog are written if the plugin is configured as an active via [auditconfig](#). Use the `auditconfig -setplugin` option to change all the plugin related configuration parameters. Syslog messages are generated with the facility code of LOG\_AUDIT (audit in [syslog.conf\(4\)](#)) and severity of LOG\_NOTICE. Audit syslog messages contain data selected from the tokens described for the binary audit log. (See [audit.log\(4\)](#)). As with all syslog messages, each line in a syslog file consists of two parts, a syslog header and a message.

The syslog header contains the date and time the message was generated, the host name from which it was sent, `auditd` to indicate that it was generated by the audit daemon, an ID field used internally by `syslogd`, and `audit.notice` indicating the syslog facility and severity values. The syslog header ends with the characters `]`, that is, a closing square bracket and a space.

The message part starts with the event type from the header token. All subsequent data appears only if contained in the original audit record and there is room in the 1024-byte maximum length syslog line. In the following example, the backslash (`\`) indicates a continuation; actual syslog messages are contained on one line:

```
Oct 31 11:38:08 smothers auditd: [ID 917521 audit.notice] chdir(2) ok\
session 401 by joeuser as root:other from myultra obj /export/home
```

In the preceding example, `chdir(2)` is the event type. Following this field is additional data, described below. This data is omitted if it is not contained in the source audit record.

|                                            |                                                                                                                                                                                                            |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>ok or failed</code>                  | Comes from the return or exit token.                                                                                                                                                                       |
| <code>session &lt;#&gt;</code>             | <code>&lt;#&gt;</code> is the session ID from the subject token.                                                                                                                                           |
| <code>by &lt;name&gt;</code>               | <code>&lt;name&gt;</code> is the audit ID from the subject token.                                                                                                                                          |
| <code>as &lt;name&gt;:&lt;group&gt;</code> | <code>&lt;name&gt;</code> is the effective user ID and <code>&lt;group&gt;</code> is the effective group ID from the subject token.                                                                        |
| <code>in &lt;zone name&gt;</code>          | The zone name. This field is generated only if the zonename audit policy is set.                                                                                                                           |
| <code>from &lt;terminal&gt;</code>         | <code>&lt;terminal&gt;</code> is the text machine address from the subject token.                                                                                                                          |
| <code>obj &lt;path&gt;</code>              | <code>&lt;path&gt;</code> is the path from the path token. The path can be truncated from the left if necessary to fit it on the line. Truncation is indicated by leading ellipsis ( <code>. . .</code> ). |

`proc_uid <owner>`      `<owner>` is the effective user ID of the process owner.

`proc_auid <owner>`      `<owner>` is the audit ID of the process owner.

The following are example syslog messages:

```
Nov  4  8:27:07 smothers auditd: [ID 175219 audit.notice]
\system booted
```

```
Nov  4  9:28:17 smothers auditd: [ID 752191 audit.notice] \
login - rlogin ok session 401 by joeuser as joeuser:staff from myultra
```

```
Nov  4 10:29:27 smothers auditd: [ID 521917 audit.notice] \
access(2) ok session 255 by janeuser as janeuser:staff from \
129.146.89.30 obj /etc/passwd
```

**Object Attributes** The `p_flag` attribute is used to further filter audit data being sent to the syslog daemon beyond the classes specified through the `flags` and `naflags` (see [auditconfig\(1M\)](#)) and through the user-specific lines of [user\\_attr\(4\)](#). The parameter is a comma-separated list; each item represents an audit class (see [audit\\_class\(4\)](#)) and is specified using the syntax described in [audit\\_flags\(5\)](#). The default (empty `p_flags` listed) is that no audit records are generated.

**Examples** **EXAMPLE 1** One Use of the plugin Line

In the specification shown below, the `plugin` (in conjunction with setting `flags` and `naflags`) is used to allow class records for `lo` but allows class records for `am` for failures only. Omission of the `fm` class records results in no `fm` class records being output. The `pc` parameter has no effect because you cannot add classes to those defined by means of `flags` and `naflags` and by [user\\_attr\(4\)](#). You can only remove them.

```
auditconfig -setflags lo,am,fm
auditconfig -setnaflags lo
auditconfig -setplugin audit_syslog active "p_flags=lo,-am,pc"
```

**EXAMPLE 2** Use of all

In the specification shown below, with one exception, `all` allows all flags defined by means of `flags` and `naflags` (and [user\\_attr\(4\)](#)). The exception the `am` metaclass, which is equivalent to `ss`, `as`, `ua`, which is modified to output all `ua` events but only failure events for `ss` and `as`.

```
auditconfig -setflags lo,am
auditconfig -setnaflags lo
auditconfig -setplugin audit_syslog active "p_flags=all,^+ss,^+as"
```

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| MT Level            | MT-Safe         |
| Interface Stability | See below.      |

The message format and message content are Uncommitted. The configuration parameters are Committed.

**See Also** [auditconfig\(1M\)](#), [auditd\(1M\)](#), [audit\\_class\(4\)](#), [syslog.conf\(4\)](#), [user\\_attr\(4\)](#), [attributes\(5\)](#), [audit\\_flags\(5\)](#)

*Oracle Solaris 11.1 Administration: Security Services*

**Notes** Activating the `audit_syslog` plugin requires that `/etc/syslog.conf` is configured to store `syslog` messages of facility `audit` and severity `notice` or above in a file intended for Solaris audit records. An example of such a line in `syslog.conf` is:

```
audit.notice          /var/audit/audit.log
```

Messages from `syslog` are sent to remote `syslog` servers by means of UDP, which does not guarantee delivery or ensure the correct order of arrival of messages.

If the parameters specified for the `plugin` line result in no classes being preselected, an error is reported by means of a `syslog` alert with the `LOG_DAEMON` facility code.

The time field in the `syslog` header is generated by [syslog\(3C\)](#) and only approximates the time given in the binary audit log. Normally the time field shows the same whole second or at most a few seconds difference.

- 
- Name** brands – alternate operating environments for non-global zones
- Description** The branded zone (BrandZ) framework extends the Solaris Zones infrastructure described in [zones\(5\)](#) to include the creation of brands, which provide non-global zones that contain non-native operating environments.
- The term “brand” can refer to a wide range of operating environments. All brand management is performed as extensions to the current zones structure.
- Every zone is configured with an associated brand. The brand type is used to determine which scripts are executed when a zone is installed and booted. In addition, a zone's brand is used to properly identify the correct application type at application launch time. The default brand is determined by the installed distribution in the global zone.
- A branded zone will support exactly one brand of non-native binary, which means that a branded zone provides a single operating environment. Once a zone has been assigned a brand, that brand cannot be changed or removed.
- BrandZ extends the zones tools in the following ways:
- A brand is an attribute of a zone, set at zone create time.
  - The `zoncfg` tool (see [zoncfg\(1M\)](#)) is used to set a zone's brand type and configure the zone.
  - The `zoneadm` tool (see [zoneadm\(1M\)](#)) is used to report a zone's brand type and administer the zone.
- Brands** For Oracle Solaris 11, the available brands are:
- `solaris`  
See [solaris\(5\)](#).
  - `solaris10`  
See [solaris10\(5\)](#).
  - `labeled`  
Implements trusted extensions, see [txzonemgr\(1M\)](#).
- Device Support** The devices supported by each zone are documented in the man pages and other documentation for that brand. The zones infrastructure detects any attempt to add an unsupported device and issues a warning to the administrator. If the administrator chooses to add an unsupported device despite that warning, that device might or might not work as expected. The configuration will be untested and unsupported.
- Attributes** See [attributes\(5\)](#) for a description of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Availability        | system/zones    |
| Interface Stability | Committed       |

**See Also** `mdb(1)`, `zlogin(1)`, `zonename(1)`, `dtrace(1M)`, `in.rlogind(1M)`, `sshd(1M)`, `txzonemgr(1M)`, `zoneadm(1M)`, `zonecfg(1M)`, `kill(2)`, `prionctl(2)`, `getzoneid(3C)`, `ucred_get(3C)`, `getzoneid(3C)`, `proc(4)`, `attributes(5)`, `solaris(5)`, `privileges(5)`, `solaris(5)`, `solaris10(5)`, `zones(5)`, `crgetzoneid(9F)`

**Name** cancellation – overview of concepts related to POSIX thread cancellation

| Description | FUNCTION                 | ACTION                                              |
|-------------|--------------------------|-----------------------------------------------------|
|             | pthread_cancel()         | Cancels thread execution.                           |
|             | pthread_setcancelstate() | Sets the cancellation <i>state</i> of a thread.     |
|             | pthread_setcanceltype()  | Sets the cancellation <i>type</i> of a thread.      |
|             | pthread_testcancel()     | Creates a cancellation point in the calling thread. |
|             | pthread_cleanup_push()   | Pushes a cleanup handler routine.                   |
|             | pthread_cleanup_pop()    | Pops a cleanup handler routine.                     |

**Cancellation** Thread cancellation allows a thread to terminate the execution of any application thread in the process. Cancellation is useful when further operations of one or more threads are undesirable or unnecessary.

An example of a situation that could benefit from using cancellation is an asynchronously-generated cancel condition such as a user requesting to close or exit some running operation. Another example is the completion of a task undertaken by a number of threads, such as solving a maze. While many threads search for the solution, one of the threads might solve the puzzle while the others continue to operate. Since they are serving no purpose at that point, they should all be canceled.

**Planning Steps** Planning and programming for most cancellations follow this pattern:

1. Identify which threads you want to cancel, and insert `pthread_cancel(3C)` statements.
2. Identify system-defined cancellation points where a thread that might be canceled could have changed system or program state that should be restored. See the `Cancellation Points` for a list.
3. When a thread changes the system or program state just before a cancellation point, and should restore that state before the thread is canceled, place a cleanup handler before the cancellation point with `pthread_cleanup_push(3C)`. Wherever a thread restores the changed state, pop the cleanup handler from the cleanup stack with `pthread_cleanup_pop(3C)`.
4. Know whether the threads you are canceling call into cancel-unsafe libraries, and disable cancellation with `pthread_setcancelstate(3C)` before the call into the library. See `Cancellation State and Cancel-Safe`.
5. To cancel a thread in a procedure that contains no cancellation points, insert your own cancellation points with `pthread_testcancel(3C)`. This function creates cancellation points by testing for pending cancellations and performing those cancellations if they are found. Push and pop cleanup handlers around the cancellation point, if necessary (see Step 3, above).

**Cancellation Points** The system defines certain points at which cancellation can occur (cancellation points), and you can create additional cancellation points in your application with `pthread_testcancel()`.

The following cancellation points are defined by the system (system-defined cancellation points): `creat(2)`, `aiosuspend(3C)`, `close(2)`, `creat(2)`, `getmsg(2)`, `getpmsg(2)`, `lockf(3C)`, `mq_receive(3C)`, `mq_send(3C)`, `msgrcv(2)`, `msgsnd(2)`, `msync(3C)`, `nanosleep(3C)`, `open(2)`, `pause(2)`, `poll(2)`, `pread(2)`, `pthread_cond_timedwait(3C)`, `pthread_cond_wait(3C)`, `pthread_join(3C)`, `pthread_testcancel(3C)`, `putmsg(2)`, `putpmsg(2)`, `pwrite(2)`, `read(2)`, `readv(2)`, `select(3C)`, `sem_wait(3C)`, `sigpause(3C)`, `sigwaitinfo(3C)`, `sigsuspend(2)`, `sigtimedwait(3C)`, `sigwait(2)`, `sleep(3C)`, `sync(2)`, `system(3C)`, `tcdrain(3C)`, `usleep(3C)`, `wait(3C)`, `waitid(2)`, `wait3(3C)`, `waitpid(3C)`, `write(2)`, `writev(2)`, and `fcntl(2)`, when specifying `F_SETLKW` as the command.

When cancellation is asynchronous, cancellation can occur at any time (before, during, or after the execution of the function defined as the cancellation point). When cancellation is deferred (the default case), cancellation occurs only within the scope of a function defined as a cancellation point (after the function is called and before the function returns). See **Cancellation Type** for more information about deferred and asynchronous cancellation.

Choosing where to place cancellation points and understanding how cancellation affects your program depend upon your understanding of both your application and of cancellation mechanics.

Typically, any call that might require a long wait should be a cancellation point. Operations need to check for pending cancellation requests when the operation is about to block indefinitely. This includes threads waiting in `pthread_cond_wait()` and `pthread_cond_timedwait()`, threads waiting for the termination of another thread in `pthread_join()`, and threads blocked on `sigwait()`.

A mutex is explicitly not a cancellation point and should be held for only the minimal essential time.

Most of the dangers in performing cancellations deal with properly restoring invariants and freeing shared resources. For example, a carelessly canceled thread might leave a mutex in a locked state, leading to a deadlock. Or it might leave a region of memory allocated with no way to identify it and therefore no way to free it.

**Cleanup Handlers** When a thread is canceled, it should release resources and clean up the state that is shared with other threads. So, whenever a thread that might be canceled changes the state of the system or of the program, be sure to push a cleanup handler with `pthread_cleanup_push(3C)` before the cancellation point.

When a thread is canceled, all the currently-stacked cleanup handlers are executed in last-in-first-out (LIFO) order. Each handler is run in the scope in which it was pushed. When the last cleanup handler returns, the thread-specific data destructor functions are called. Thread execution terminates when the last destructor function returns.

When, in the normal course of the program, an uncanceled thread restores state that it had previously changed, be sure to pop the cleanup handler (that you had set up where the change took place) using `pthread_cleanup_pop(3C)`. That way, if the thread is canceled later, only currently-changed state will be restored by the handlers that are left in the stack.

The `pthread_cleanup_push()` and `pthread_cleanup_pop()` functions can be implemented as macros. The application must ensure that they appear as statements, and in pairs within the same lexical scope (that is, the `pthread_cleanup_push()` macro can be thought to expand to a token list whose first token is '{' with `pthread_cleanup_pop()` expanding to a token list whose last token is the corresponding '}').

The effect of the use of `return`, `break`, `continue`, and `goto` to prematurely leave a code block described by a pair of `pthread_cleanup_push()` and `pthread_cleanup_pop()` function calls is undefined.

**Cancellation State** Most programmers will use only the default cancellation state of `PTHREAD_CANCEL_ENABLE`, but can choose to change the state by using `pthread_setcancelstate(3C)`, which determines whether a thread is cancelable at all. With the default *state* of `PTHREAD_CANCEL_ENABLE`, cancellation is enabled and the thread is cancelable at points determined by its cancellation *type*. See Cancellation Type.

If the *state* is `PTHREAD_CANCEL_DISABLE`, cancellation is disabled, the thread is not cancelable at any point, and all cancellation requests to it are held pending.

You might want to disable cancellation before a call to a cancel-unsafe library, restoring the old cancel state when the call returns from the library. See `Cancel-Safe` for explanations of cancel safety.

**Cancellation Type** A thread's cancellation type is set with `pthread_setcanceltype(3C)`, and determines whether the thread can be canceled anywhere in its execution or only at cancellation points.

With the default *type* of `PTHREAD_CANCEL_DEFERRED`, the thread is cancelable only at cancellation points, and then only when cancellation is enabled.

If the *type* is `PTHREAD_CANCEL_ASYNCHRONOUS`, the thread is cancelable at any point in its execution (assuming, of course, that cancellation is enabled). Try to limit regions of asynchronous cancellation to sequences with no external dependencies that could result in dangling resources or unresolved state conditions. Using asynchronous cancellation is discouraged because of the danger involved in trying to guarantee correct cleanup handling at absolutely every point in the program.

| Cancellation Type/State Table |       |                   |
|-------------------------------|-------|-------------------|
| Type                          | State |                   |
|                               |       | Enabled (Default) |

| Cancellation Type/State Table |                                                                                                            |                                                                                                                                                         |
|-------------------------------|------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| Deferred (Default)            | Cancellation occurs when the target thread reaches a cancellation point and a cancel is pending. (Default) | All cancellation requests to the target thread are held pending.                                                                                        |
| Asynchronous                  | Receipt of a <code>pthread_cancel()</code> call causes immediate cancellation.                             | All cancellation requests to the target thread are held pending; as soon as cancellation is re-enabled, pending cancellations are executed immediately. |

**Cancel-Safe** With the arrival of POSIX cancellation, the Cancel-Safe level has been added to the list of MT-Safety levels. See [attributes\(5\)](#). An application or library is Cancel-Safe whenever it has arranged for cleanup handlers to restore system or program state wherever cancellation can occur. The application or library is specifically Deferred-Cancel-Safe when it is Cancel-Safe for threads whose cancellation type is `PTHREAD_CANCEL_DEFERRED`. See [Cancellation State](#). It is specifically Asynchronous-Cancel-Safe when it is Cancel-Safe for threads whose cancellation type is `PTHREAD_CANCEL_ASYNCHRONOUS`.

It is easier to arrange for deferred cancel safety, as this requires system and program state protection only around cancellation points. In general, expect that most applications and libraries are not Asynchronous-Cancel-Safe.

**POSIX Threads Only** The cancellation functions described in this manual page are available for POSIX threads, only (the Solaris threads interfaces do not provide cancellation functions).

### Examples **EXAMPLE 1** Cancellation example

The following short C++ example shows the pushing/popping of cancellation handlers, the disabling/enabling of cancellation, the use of `pthread_testcancel()`, and so on. The `free_res()` cancellation handler in this example is a dummy function that simply prints a message, but that would free resources in a real application. The function `f2()` is called from the main thread, and goes deep into its call stack by calling itself recursively.

Before `f2()` starts running, the newly created thread has probably posted a cancellation on the main thread since the main thread calls `thr_yield()` right after creating `thread2`. Because cancellation was initially disabled in the main thread, through a call to `pthread_setcancelstate()`, the call to `f2()` from `main()` continues and constructs `X` at each recursive call, even though the main thread has a pending cancellation.

When `f2()` is called for the fifty-first time (when `"i == 50"`), `f2()` enables cancellation by calling `pthread_setcancelstate()`. It then establishes a cancellation point for itself by calling `pthread_testcancel()`. (Because a cancellation is pending, a call to a cancellation point such as `read(2)` or `write(2)` would also cancel the caller here.)

**EXAMPLE 1** Cancellation example *(Continued)*

After the `main()` thread is canceled at the fifty-first iteration, all the cleanup handlers that were pushed are called in sequence; this is indicated by the calls to `free_res()` and the calls to the destructor for `X`. At each level, the C++ runtime calls the destructor for `X` and then the cancellation handler, `free_res()`. The print messages from `free_res()` and `X`'s destructor show the sequence of calls.

At the end, the main thread is joined by `thread2`. Because the main thread was canceled, its return status from `pthread_join()` is `PTHREAD_CANCELED`. After the status is printed, `thread2` returns, killing the process (since it is the last thread in the process).

```
#include <pthread.h>
#include <sched.h>
extern "C" void thr_yield(void);

extern "C" void printf(...);

struct X {
    int x;
    X(int i){x = i; printf("X(%d) constructed.\n", i);}
    ~X(){ printf("X(%d) destroyed.\n", x);}
};

void
free_res(void *i)
{
    printf("Freeing '%d'\n",i);
}

char* f2(int i)
{
    try {
        X dummy(i);
        pthread_cleanup_push(free_res, (void *)i);
        if (i == 50) {
            pthread_setcancelstate(PTHREAD_CANCEL_ENABLE, NULL);
            pthread_testcancel();
        }
        f2(i+1);
        pthread_cleanup_pop(0);
    }
    catch (int) {
        printf("Error: In handler.\n");
    }
    return "f2";
}
```

**EXAMPLE 1** Cancellation example *(Continued)*

```
void *
thread2(void *tid)
{
    void *sts;

    printf("I am new thread :%d\n", pthread_self());

    pthread_cancel((pthread_t)tid);

    pthread_join((pthread_t)tid, &sts);

    printf("main thread cancelled due to %d\n", sts);

    return (sts);
}

main()
{
    pthread_setcancelstate(PTHREAD_CANCEL_DISABLE, NULL);
    pthread_create(NULL, NULL, thread2, (void *)pthread_self());
    thr_yield();
    printf("Returned from %s\n", f2(0));
}
```

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTETYPE | ATTRIBUTEVALUE |
|---------------|----------------|
| MT-Level      | MT-Safe        |

**See Also** [read\(2\)](#), [sigwait\(2\)](#), [write\(2\)](#), [Intro\(3\)](#), [condition\(5\)](#), [pthread\\_cleanup\\_pop\(3C\)](#), [pthread\\_cleanup\\_push\(3C\)](#), [pthread\\_exit\(3C\)](#), [pthread\\_join\(3C\)](#), [pthread\\_setcancelstate\(3C\)](#), [pthread\\_setcanceltype\(3C\)](#), [pthread\\_testcancel\(3C\)](#), [setjmp\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

**Name** charmap – character set description file

**Description** A character set description file or *charmap* 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.

The character set description file provides:

- 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 `localedef(1)` 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** 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.

|       |       |       |       |       |       |
|-------|-------|-------|-------|-------|-------|
| <ACK> | <DC2> | <ENQ> | <FS>  | <IS4> | <SOH> |
| <BEL> | <DC3> | <EOT> | <GS>  | <LF>  | <STX> |
| <BS>  | <DC4> | <ESC> | <HT>  | <NAK> | <SUB> |
| <CAN> | <DEL> | <ETB> | <IS1> | <RS>  | <SYN> |
| <CR>  | <DLE> | <ETX> | <IS2> | <SI>  | <US>  |
| <DC1> | <EM>  | <FF>  | <IS3> | <SO>  | <VT>  |

**Declarations** 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.

- <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.



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>
```

```
"%c%0" , <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. Bytes are treated as unsigned octets and carry is propagated between the bytes as necessary to represent the range. However, because this causes a null byte in the second or subsequent bytes of a character, such a declaration should not be specified. For example, the line

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

is interpreted as:

```
<j0101>                \d129\d254
<j0102>                \d129\d255
<j0103>                \d130\d00
<j0104>                \d130\d01
```

The expanded declaration of the symbol `<j0103>` in the above example is an invalid specification, because it contains a null byte in the second byte of a character.

The comment is optional.

**Width Specification** The following declarations can follow the character set mapping definitions (after the “END CHARMAP” statement). Each consists of the keyword shown in the following list, starting in column 1, followed by the value(s) to be associated to the keyword, as defined below.

```
WIDTH                A non-negative integer value defining the column width for the printable
                    character in the coded character set mapping definitions. Coded
```

character set character values are defined using symbolic character names followed by column width values. Defining a character with more than one WIDTH produces undefined results. The END WIDTH keyword is used to terminate the WIDTH definitions. Specifying the width of a non-printable character in a WIDTH declaration produces undefined results.

**WIDTH\_DEFAULT** A non-negative integer value defining the default column width for any printable character not listed by one of the WIDTH keywords. If no WIDTH\_DEFAULT keyword is included in the charmap, the default character width is 1.

Example:

After the “END CHARMAP” statement, a syntax for a width definition would be:

```
WIDTH
<A>          1
<B>          1
<C>...<Z>    1
...
<fool>...<foon> 2
...
END WIDTH
```

In this example, the numerical code point values represented by the symbols <A> and <B> are assigned a width of 1. The code point values <C> to <Z> inclusive, that is, <C>, <D>, <E>, and so on, are also assigned a width of 1. Using <A> . . . <Z> would have required fewer lines, but the alternative was shown to demonstrate flexibility. The keyword WIDTH\_DEFAULT could have been added as appropriate.

**See Also** [locale\(1\)](#), [localedef\(1\)](#), [nl\\_langinfo\(3C\)](#), [extensions\(5\)](#), [locale\(5\)](#)

---

|                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>           | condition – concepts related to condition variables                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>    | <p>Occasionally, a thread running within a mutex needs to wait for an event, in which case it blocks or sleeps. When a thread is waiting for another thread to communicate its disposition, it uses a condition variable in conjunction with a mutex. Although a mutex is exclusive and the code it protects is sharable (at certain moments), condition variables enable the synchronization of differing events that share a mutex, but not necessarily data. Several condition variables may be used by threads to signal each other when a task is complete, which then allows the next waiting thread to take ownership of the mutex.</p> <p>A condition variable enables threads to atomically block and test the condition under the protection of a mutual exclusion lock (mutex) until the condition is satisfied. If the condition is false, a thread blocks on a condition variable and atomically releases the mutex that is waiting for the condition to change. If another thread changes the condition, it may wake up waiting threads by signaling the associated condition variable. The waiting threads, upon awakening, reacquire the mutex and re-evaluate the condition.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Initialize</b>     | <p>Condition variables and mutexes should be global. Condition variables that are allocated in writable memory can synchronize threads among processes if they are shared by the cooperating processes (see <a href="#">mmap(2)</a>) and are initialized for this purpose.</p> <p>The scope of a condition variable is either intra-process or inter-process. This is dependent upon whether the argument is passed implicitly or explicitly to the initialization of that condition variable. A condition variable does not need to be explicitly initialized. A condition variable is initialized with all zeros, by default, and its scope is set to within the calling process. For inter-process synchronization, a condition variable must be initialized once, and only once, before use.</p> <p>A condition variable must not be simultaneously initialized by multiple threads or re-initialized while in use by other threads.</p> <p>Condition variables attributes may be set to the default or customized at initialization. POSIX threads even allow the default values to be customized. Establishing these attributes varies depending upon whether POSIX or Solaris threads are used. Similar to the distinctions between POSIX and Solaris thread creation, POSIX condition variables implement the default, intra-process, unless an attribute object is modified for inter-process prior to the initialization of the condition variable. Solaris condition variables also implement as the default, intra-process; however, they set this attribute according to the argument, <i>type</i>, passed to their initialization function.</p> |
| <b>Condition Wait</b> | <p>The condition wait interface allows a thread to wait for a condition and atomically release the associated mutex that it needs to hold to check the condition. The thread waits for another thread to make the condition true and that thread's resulting call to signal and wakeup the waiting thread.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**Condition Signaling** A condition signal allows a thread to unblock the next thread waiting on the condition variable, whereas, a condition broadcast allows a thread to unblock all threads waiting on the condition variable.

**Destroy** The condition destroy functions destroy any state, but not the space, associated with the condition variable.

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| MT-Level       | MT-Safe         |

**See Also** [fork\(2\)](#), [mmap\(2\)](#), [setitimer\(2\)](#), [shmop\(2\)](#), [cond\\_broadcast\(3C\)](#), [cond\\_destroy\(3C\)](#), [cond\\_init\(3C\)](#), [cond\\_signal\(3C\)](#), [cond\\_timedwait\(3C\)](#), [cond\\_wait\(3C\)](#), [pthread\\_cond\\_broadcast\(3C\)](#), [pthread\\_cond\\_destroy\(3C\)](#), [pthread\\_cond\\_init\(3C\)](#), [pthread\\_cond\\_signal\(3C\)](#), [pthread\\_cond\\_timedwait\(3C\)](#), [pthread\\_cond\\_wait\(3C\)](#), [pthread\\_condattr\\_init\(3C\)](#), [signal\(3C\)](#), [attributes\(5\)](#), [mutex\(5\)](#), [standards\(5\)](#)

**Notes** If more than one thread is blocked on a condition variable, the order in which threads are unblocked is determined by the scheduling policy.

USYNC\_THREAD does not support multiple mappings to the same logical synch object. If you need to `mmap()` a synch object to different locations within the same address space, then the synch object should be initialized as a shared object `USYNC_PROCESS` for Solaris, and `PTHREAD_PROCESS_PRIVATE` for POSIX.

**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.

The maximum password length for crypt\_bsdbf is 72 characters.

The default number of rounds is  $\log(4)$ . This can be increased by updating `/etc/security/crypt.conf`, for example, for  $\log(12)$  rounds:

```
2a crypt_bsdbf.so.1 12
```

**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\_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.

The maximum password length for crypt\_bsdmd5 is 255 characters.

**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\_sha256 – password hashing module using SHA–256 message hash algorithm

**Synopsis** /usr/lib/security/\$ISA/crypt\_sha256.so

**Description** The crypt\_sha256 module is a one-way password hashing module for use with [crypt\(3C\)](#) that uses the SHA–256 message hash algorithm. The algorithm identifier for [crypt.conf\(4\)](#) and [policy.conf\(4\)](#) is 5.

This module is designed to make it difficult to crack passwords that use brute force attacks based on high speed SHA–256 implementations that use code inlining, unrolled loops, and table lookup.

The maximum password length for crypt\_sha256 is 255 characters.

The following options can be passed to the module by means of [crypt.conf\(4\)](#):

rounds=<positive\_number>

Specifies the number of rounds of SHA-256 to use in generation of the salt; the default number of rounds is 5000. Negative values have no effect and are ignored. The minimum number of rounds cannot be below 1000.

The number of additional rounds is stored in the salt string returned by [crypt\\_gensalt\(3C\)](#). For example:

```
$5,rounds=6000$nlxmTtpz$
```

When [crypt\\_gensalt\(3C\)](#) is being used to generate a new salt, if the number of additional rounds configured in [crypt.conf\(4\)](#) is greater than that in the old salt, the value from [crypt.conf\(4\)](#) is used instead. This allows for migration to stronger (but more time-consuming) salts on password change.

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Interface Stability | Committed       |
| 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\_sha512 – password hashing module using SHA-512 message hash algorithm

**Synopsis** /usr/lib/security/\$ISA/crypt\_sha512.so

**Description** The crypt\_sha512 module is a one-way password hashing module for use with [crypt\(3C\)](#) that uses the SHA-512 message hash algorithm. The algorithm identifier for [crypt.conf\(4\)](#) and [policy.conf\(4\)](#) is 6.

This module is designed to make it difficult to crack passwords that use brute force attacks based on high speed SHA-512 implementations that use code inlining, unrolled loops, and table lookup.

The maximum password length for crypt\_sha512 is 255 characters.

The following options can be passed to the module by means of [crypt.conf\(4\)](#):

`rounds=<positive_number>`

Specifies the number of rounds of SHA-512 to use in generation of the salt; the default number of rounds is 5000. Negative values have no effect and are ignored. The minimum number of rounds cannot be below 1000.

The number of additional rounds is stored in the salt string returned by [crypt\\_gensalt\(3C\)](#). For example:

```
$6,rounds=6000$nlxmTtpz$
```

When [crypt\\_gensalt\(3C\)](#) is being used to generate a new salt, if the number of additional rounds configured in [crypt.conf\(4\)](#) is greater than that in the old salt, the value from [crypt.conf\(4\)](#) is used instead. This allows for migration to stronger (but more time-consuming) salts on password change.

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Interface Stability | Committed       |
| 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** The crypt\_sunmd5 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 md5.

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.

The maximum password length for crypt\_sunmd5 is 255 characters.

The following options can be passed to the module by means of [crypt.conf\(4\)](#):

`rounds=<positive_number>` Specifies the number of additional rounds of MD5 to use in generation of the salt; the default number of rounds is 4096. Negative values have no effect and are ignored, that is, the number of rounds cannot be lowered below 4096.

The number of additional rounds is stored in the salt string returned by [crypt\\_gensalt\(3C\)](#). For example:

```
$md5, rounds=1000$n1xmTtpz$
```

When [crypt\\_gensalt\(3C\)](#) is being used to generate a new salt, if the number of additional rounds configured in [crypt.conf\(4\)](#) is greater than that in the old salt, the value from [crypt.conf\(4\)](#) is used instead. This allows for migration to stronger (but more time-consuming) salts on password change.

**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\_unix – traditional UNIX crypt algorithm

**Description** The crypt\_unix algorithm is the traditional UNIX crypt algorithm. It is not considered sufficiently secure for current systems and is provided for backwards compatibility. The [crypt\\_sunmd5\(5\)](#), [crypt\\_bsdmd5\(5\)](#), or [crypt\\_bsdbf\(5\)](#) algorithm should be used instead.

The algorithm identifier for [policy.conf\(4\)](#) is `__unix__`. There is no entry in [crypt.conf\(4\)](#) for this algorithm.

The crypt\_unix algorithm is internal to libc and provides the string encoding function used by [crypt\(3C\)](#) when the first character of the salt is not a "\$".

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 `crypt()` are used with this algorithm; the rest are silently ignored. The salt is a two-character string chosen from the set `[a-zA-Z0-9./]`. This string is used to perturb the hashing algorithm in one of 4096 different ways.

The maximum password length for crypt\_unix is 8 characters.

**Usage** The return value of the crypt\_unix algorithm might not be portable among standard-conforming systems. See [standards\(5\)](#).

**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\\_bsdbf\(5\)](#), [crypt\\_bsdmd5\(5\)](#), [crypt\\_sunmd5\(5\)](#), [standards\(5\)](#)

**Name** datasets – Oracle Solaris ZFS dataset organization

**Description** Oracle Solaris is installed to a hierarchical set of ZFS datasets with well-defined names. The ZFS datasets, composed of both file systems and volumes, are stored on a bootable ZFS storage pool in the global zone. Snapshots are also ZFS datasets, but will not be discussed here.

The datasets are created automatically by one of the installation tools: the automated installer, the text installer, or the GUI installer. For non-global zones, the datasets are created by the zone brand during the zone installation and might differ from the hierarchy described here, which applies to the `solaris(5)` brand and the global zone.

**rpool Dataset** `rpool` is the default name of the top-level dataset in the bootable ZFS storage pool and is mounted at `/rpool`. Bootable ZFS storage pools have certain restrictions over normal ZFS storage pools. These restrictions are discussed in the “Oracle Solaris Administration: ZFS File Systems” section of the Oracle Solaris 11 Information Library.

Administrators can create datasets anywhere beneath the top-level dataset, but those created within the `'ROOT/be-<name>'` dataset are treated specially by `beadm(1M)` and are discussed later.

**ROOT Dataset** This is a container dataset which is not mounted by Solaris during operation. It collects sets of related datasets, each set forming a ZFS Boot Environment (BE). A BE is created on installation. Additional BEs are created by `pkg(1)` and can be modified, created, or destroyed by `beadm`.

In the global zone, the ROOT dataset must be stored within a bootable ZFS storage pool, called `'rpool'` by default. For non-global zones, the ROOT dataset is stored in a dataset with the leaf name `'rpool'` beneath the zonename of each zone.

The list below shows the datasets that are created beneath the ROOT dataset. This list uses the default storage pool name.

`rpool/ROOT/solaris`

This is the default dataset that is created for the first boot environment installed on the system. It contains the root file system and is mounted on `'/'`.

The name of the leaf component of this dataset corresponds to the name of the BE that is associated with it. Any datasets created by the administrator as children of the `'ROOT/<be-name>'` dataset will be associated with that BE, and will be created, cloned, destroyed and have snapshots taken of them as part of the BE by `beadm`. New datasets created beneath `'ROOT/<be-name>'` should have the `zfs` property `'canmount=noauto'` set so that the datasets are mounted by `beadm` when appropriate.

All packaged content that is part of the image in a boot environment must fully reside within the datasets belonging to the boot environment.

`rpool/ROOT/solaris/var`

This is a dataset associated with each boot environment that contains the `/var` file system for the BE.

- Dump Dataset** `rpool/dump` is a ZFS volume created during installation and is used as the default dump device for the system crash dump facility. This dataset is shared across all BEs. The crash dump facility can be managed with [dumpadm\(1M\)](#) and the dump device can be resized using [zfs\(1M\)](#). The `rpool/dump` dataset is not present on non-global zones. Other devices can be used instead of `rpool/dump`.
- Swap Dataset** `rpool/swap` is a ZFS zvol created during installation and is used as the default swap area for the system. This dataset is shared across all BEs. Use [swap\(1M\)](#) to add or remove swap volumes. Use `zfs` to resize swap volumes. The `rpool/swap` dataset is not present on non-global zones. Other swap volumes can be used instead, as well as `rpool/swap`.
- `rpool/VARSHARE` `rpool/VARSHARE` is a dataset created on installation or upgrade, and is not part of any one BE. Instead, it is shared across all BEs and is mounted automatically by that BE during boot. It contains unpackaged content from `/var`. For example, audit information is stored in `/var/audit`, core files are stored in `/var/cores`, saved crash dumps are stored in `/var/crash`, and system mail is stored in `/var/mail`.
- `rpool/export/home` `rpool/export/home` is a dataset created on installation which can be used for user home directories. It is the default location for user home directories if the `-d` option is not passed to [useradd\(1M\)](#) to specify a different directory. If the user home directory does not already exist and the parent directory is the mount point of a ZFS dataset (for example, `rpool/export/home`), then a child of that dataset will be created and mounted at the specified location. These datasets are shared across all BEs.
- See Also** [beadm\(1M\)](#), [dumpadm\(1M\)](#), [swap\(1M\)](#), [useradd\(1M\)](#), [zfs\(1M\)](#), [zoneadm\(1M\)](#), [zpool\(1M\)](#), [brands\(5\)](#), [filesystem\(5\)](#), [pkg\(5\)](#), [solaris\(5\)](#)

**Name** device\_clean – device clean programs

**Description** Each allocatable device has a device clean program associated with it. Device clean programs are invoked by `deallocate(1)` to clean device states, registers, and any residual information in the device before the device is allocated to a user. Such cleaning is required by the object reuse policy.

Use `list_devices(1)` to obtain the names and types of allocatable devices as well as the cleaning program and the authorizations that are associated with each device.

On a system configured with Trusted Extensions, device clean programs are also invoked by `allocate(1)`, in which case the program can optionally mount appropriate media for the caller.

The following device clean programs reside in `/etc/security/lib`.

|                          |                |
|--------------------------|----------------|
| <code>audio_clean</code> | audio devices  |
| <code>st_clean</code>    | tape devices   |
| <code>sr_clean</code>    | CD-ROM devices |

On a system configured with Trusted Extensions, the following additional cleaning programs and wrappers are available.

|                                  |                                                                                                                                       |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <code>disk_clean</code>          | CD-ROM and other removable media devices. This program mounts the device during the execution of <code>allocate</code> , if required. |
| <code>audio_clean_wrapper</code> | wrapper to make <code>audio_clean</code> work with CDE                                                                                |
| <code>wdwrapper</code>           | wrapper to make other cleaning programs work with CDE                                                                                 |
| <code>wdwmsg</code>              | CDE dialog boxes for cleaning programs                                                                                                |

Administrators can create device clean programs for their sites. These programs must adhere to the syntax described below.

```
/etc/security/lib/device-clean-program [-i | -f | -s | -I] \
-m mode -u user-name -z zone-name -p zone-path device-name
```

where:

|                    |                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------|
| <i>device-name</i> | The name of the device that is to be cleaned. Use <code>list_devices</code> to obtain the list of allocatable devices. |
| <code>-i</code>    | Invoke boot-time initialization.                                                                                       |
| <code>-f</code>    | Force cleanup by the administrator.                                                                                    |
| <code>-s</code>    | Invoke standard cleanup by the user.                                                                                   |
| <code>-I</code>    | Same as <code>-i</code> , with no error or warning.                                                                    |

The following options are supported only when the system is configured with Trusted Extensions.

- m *mode* Specify the mode in which the clean program is invoked. Valid values are `allocate` and `deallocate`. The default mode is `allocate`.
- u *user-name* Specify the name of user who executes the device clean program. The default user is the caller.
- z *zone-name* Specify the name of the zone in which the device is to be allocated or deallocated. The default zone is the global zone.
- p *zone-path* Establish the root path of the zone that is specified by *zone-name*. Default is `/`.

**Exit Status** The following exit values are returned:

- 0 Successful completion.
- 1 An error. Caller can place device in error state.
- 2 A system error. Caller can place device in error state.

On a system configured with Trusted Extensions, the following additional exit values are returned:

- 3 Mounting of device failed. Caller shall not place device in error state.
- 4 Mounting of device succeeded.

**Files** `/etc/security/lib/*` device clean programs

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Availability        | system/core-os  |
| Interface Stability | See below.      |

The Invocation is Uncommitted. The Output is Not-an-interface.

**See Also** [allocate\(1\)](#), [deallocate\(1\)](#), [list\\_devices\(1\)](#), [attributes\(5\)](#)

*Oracle Solaris 11.1 Administration: Security Services*

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |  |                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                     |                                                                                                                                                                                                   |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>                | dhcp – Dynamic Host Configuration Protocol                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |  |                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                     |                                                                                                                                                                                                   |
| <b>Description</b>         | <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 Oracle Solaris DHCP implementation.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |  |                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                     |                                                                                                                                                                                                   |
| Oracle Solaris DHCP Client | <p>The Oracle Solaris DHCP client is implemented as background daemon, <a href="#">dhcpcagent(1M)</a>.</p> <p>For IPv4, 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.</p> <p>For IPv6, this daemon is started automatically when commanded by <code>in.ndpd</code> (based on IPv6 Routing Advertisement messages). No <code>/etc/dhcp.interface</code> file is necessary, but such a file can be used to specify an interface as “primary,” provided that IPv4 DHCP is also in use.</p> <p>Network parameters needed for system configuration during bootup are extracted from the information received by the daemon through the use of the <a href="#">dhcpcinfo(1)</a> command. The daemon's default behavior can be altered by changing the tunables in the <code>/etc/default/dhcpcagent</code> file. The daemon is controlled by the <a href="#">ifconfig(1M)</a> utility. Check the status of the daemon using the <a href="#">netstat(1M)</a> and <a href="#">ifconfig(1M)</a> commands.</p> |  |                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                     |                                                                                                                                                                                                   |
| Oracle Solaris DHCP Server | <p>The Oracle Solaris DHCP server is implemented as a background daemon, <a href="#">in.dhcpd(1M)</a>. This daemon can deliver network configuration information to either BOOTP or DHCP clients. The Oracle Solaris DHCP service can be managed using the <a href="#">dhcpcmgr(1M)</a> GUI or the command line utilities <a href="#">dhcpcconfig(1M)</a>, <a href="#">dhtadm(1M)</a>, and <a href="#">pntadm(1M)</a>.</p> <p>The Oracle Solaris DHCP server is obsolete. Use the ISC DHCP server instead. See <code>usr/share/man/man5/isc-dhcp.5</code></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                     |                                                                                                                                                                                                   |
| DHCP Configuration Tables  | <p>The Oracle Solaris DHCP server stores client configuration information in the following two types of tables:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;">dhcptab tables</td> <td>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 <a href="#">dhcptab(4)</a> can be viewed and modified using the <a href="#">dhtadm(1M)</a> command or <a href="#">dhcpcmgr(1M)</a> graphical utility. See <a href="#">dhcptab(4)</a> for more information about the syntax of <code>dhcptab</code> records. See <a href="#">dhcp_inittab(4)</a> for more information about the DHCP options and symbols.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">DHCP network tables</td> <td>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</td> </tr> </table>                                                                    |  | dhcptab tables | 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 <a href="#">dhcptab(4)</a> can be viewed and modified using the <a href="#">dhtadm(1M)</a> command or <a href="#">dhcpcmgr(1M)</a> graphical utility. See <a href="#">dhcptab(4)</a> for more information about the syntax of <code>dhcptab</code> records. See <a href="#">dhcp_inittab(4)</a> for more information about the DHCP options and symbols. | DHCP network tables | 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 |
| dhcptab tables             | 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 <a href="#">dhcptab(4)</a> can be viewed and modified using the <a href="#">dhtadm(1M)</a> command or <a href="#">dhcpcmgr(1M)</a> graphical utility. See <a href="#">dhcptab(4)</a> for more information about the syntax of <code>dhcptab</code> records. See <a href="#">dhcp_inittab(4)</a> for more information about the DHCP options and symbols.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |  |                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                     |                                                                                                                                                                                                   |
| DHCP network tables        | 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 `pntadm` command or `dhcpmgr` graphical utility. See `dhcp_network(4)` for more information about network tables.

**See Also** `dhcinfo(1)`, `dhcagent(1M)`, `dhcpcconfig(1M)`, `dhcpmgr(1M)`, `dhtadm(1M)`, `ifconfig(1M)`, `in.dhcpd(1M)`, `in.ndpd(1M)`, `netstat(1M)`, `pntadm(1M)`, `syslog(3C)`, `dhcp_network(4)`, `dhcptab(4)`, `dhcpsvc.conf(4)`, `dhcp_inittab(4)`, `ndpd.conf(4)`, `dhcp_modules(5)`

`usr/share/man/man5/isc-dhcp.5`

*Oracle 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.

Lemon, T. and B. Sommerfeld. *RFC 4361, Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)*. Nominum and Sun Microsystems. February 2006.

Droms, R. *RFC 3315, Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*. Cisco Systems. July 2003.

**Name** dhcp\_modules – data storage modules for the DHCP service

**Description** This man page describes the characteristics of data storage modules (public modules) for use by the Solaris Dynamic Host Configuration Protocol (DHCP) service.

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 Oracle and `ufs` file systems.

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 ([dhcpgmr\(1M\)](#), [dhcpconfig\(1M\)](#), [dhtadm\(1M\)](#), and [pntadm\(1M\)](#)).

Public modules may be provided by Sun Microsystems, Inc or by third parties.

The Solaris DHCP service management architecture provides a mechanism for plugging in public module-specific administration functionality into the [dhcpgmr\(1M\)](#) and [dhcpconfig\(1M\)](#) 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.

The Solaris DHCP service bundles three modules with the service, which are described below. There are three `dhcpsvc.conf(4)` DHCP service configuration parameters pertaining to public modules: `RESOURCE`, `PATH`, and `RESOURCE_CONFIG`. See [dhcpsvc.conf\(4\)](#) for more information about these parameters.

**SUNWfiles** 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.

This module's data may be shared between DHCP servers through the use of NFS.

**SUNWbinfiles** 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 SUNWfiles.

This module's data cannot be shared between DHCP servers.

**See Also** [crontab\(1\)](#), [dhcpconfig\(1M\)](#), [dhcpgmr\(1M\)](#), [dhtadm\(1M\)](#), [pntadm\(1M\)](#), [dhcpsvc.conf\(4\)](#), [dhcp\(5\)](#)

*Oracle Solaris DHCP Service Developer's Guide*

**Name** environ – user environment

**Description** When a process begins execution, one of the exec family of functions makes available an array of strings called the environment; see [exec\(2\)](#). By convention, these strings have the form *variable=value*, for example, `PATH=/sbin:/usr/sbin`. These environmental variables provide a way to make information about a program's environment available to programs.

A name can be placed in the environment by the `export` command and *name=value* arguments in [sh\(1\)](#), or by one of the exec functions. It is unwise to conflict with certain shell variables such as `MAIL`, `PS1`, `PS2`, and `IFS` that are frequently exported by `.profile` files; see [profile\(4\)](#).

The following environmental variables can be used by applications and are expected to be set in the target run-time environment.

#### HOME

The name of the user's login directory, set by [login\(1\)](#) from the password file; see [passwd\(4\)](#).

#### LANG

The string used to specify internationalization information that allows users to work with different national conventions. The [setlocale\(3C\)](#) function checks the `LANG` environment variable when it is called with "" as the `locale` argument. `LANG` is used as the default locale if the corresponding environment variable for a particular category is unset or null. If, however, `LC_ALL` is set to a valid, non-empty value, its contents are used to override both the `LANG` and the other `LC_*` variables. For example, when invoked as `setlocale(LC_CTYPE, "")`, `setlocale()` queries the `LC_CTYPE` environment variable first to see if it is set and non-null. If `LC_CTYPE` is not set or null, then `setlocale()` checks the `LANG` environment variable to see if it is set and non-null. If both `LANG` and `LC_CTYPE` are unset or `NULL`, the default C locale is used to set the `LC_CTYPE` category.

Most commands invoke `setlocale(LC_ALL, "")` prior to any other processing. This allows the command to be used with different national conventions by setting the appropriate environment variables.

The following environment variables correspond to each category of [setlocale\(3C\)](#):

#### LC\_ALL

If set to a valid, non-empty string value, override the values of `LANG` and all the other `LC_*` variables.

#### 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 [localedef\(1\)](#) command. This environment variable affects [strcoll\(3C\)](#) and [strxfrm\(3C\)](#).

### 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 can 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.

The extended interpretation of %L in Solaris includes the support for accepted locale name aliases as described in `gettext(1)`, `gettext(3C)`, `catopen(3C)`, `setlocale(3C)`, and `locale_alias(5)`.

NLSPATH is 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)`.

#### 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 can 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. The value of TZ has one of the two formats (spaces inserted for clarity):

:characters

or

std offset dst offset, rule

If TZ is of the first format (that is, if the first character is a colon (:)), or if TZ is not of the second format, then TZ designates a path to a timezone database file relative to `/usr/share/lib/zoneinfo/`, ignoring a leading colon if one exists.

Otherwise, TZ is of the second form, which when expanded is as follows:

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

#### *std* and *dst*

Indicate no less than three, nor more than {TZNAME\_MAX}, bytes that are the designation for the standard (*std*) or the alternative (*dst*, such as Daylight Savings Time) timezone. Only *std* is required; if *dst* is missing, then the alternative time does not apply in this timezone. Each of these fields can occur in either of two formats, quoted or unquoted:

- In the quoted form, the first character is the less-than ('<') character and the last character is the greater-than ('>') character. All characters between these quoting characters are alphanumeric characters from the portable character set in the current locale, the plus-sign ('+') character, or the minus-sign ('-') character. The *std* and *dst* fields in this case do not include the quoting characters.

- In the unquoted form, all characters in these fields are alphabetic characters from the portable character set in the current locale.

The interpretation of these fields is unspecified if either field is less than three bytes (except for the case when *dst* is missing), more than {TZNAME\_MAX} bytes, or if they contain characters other than those specified.

#### *offset*

Indicate 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 can 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 can 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 can cause unpredictable behavior. If preceded by a -, the time zone is east of the Prime Meridian. Otherwise, it is west of the Prime Meridian (which can 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 occurs. 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  $d^{\text{th}}$  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 can occur in either the fourth or the fifth week). Week 1 is the first week in which the  $d^{\text{th}}$  day occurs. Day zero is Sunday.

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

The *time* has the same format as *offset* except that no leading sign (“-” or “+”) is allowed. If *time* is not specified, the default value 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)`, `locale_alias(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 characters listed in the following table.

|                 |                       |                  |                   |                |                |
|-----------------|-----------------------|------------------|-------------------|----------------|----------------|
| <i>ciplus</i>   | $\oplus$              | $  $             | $\parallel$       | <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>&gt;wig</i>  | $\gtrsim$             | <i>ppd</i>       | $\perp$           | <i>prop</i>    | $\propto$      |
| <i>&lt;wig</i>  | $\lesssim$            | <i>&lt;-&gt;</i> | $\leftrightarrow$ | <i>empty</i>   | $\emptyset$    |
| <i>=wig</i>     | $\doteq$              | <i>&lt;=&gt;</i> | $\Leftrightarrow$ | <i>member</i>  | $\in$          |
| <i>star</i>     | $*$                   | <i> &lt;</i>     | $\leftarrow$      | <i>nomem</i>   | $\notin$       |
| <i>bigstar</i>  | $\ast$                | <i> &gt;</i>     | $\rightarrow$     | <i>cup</i>     | $\cup$         |
| <i>=dot</i>     | $\dot{=}$             | <i>ang</i>       | $\sphericalangle$ | <i>cap</i>     | $\cap$         |
| <i>orsign</i>   | $\vee$                | <i>rang</i>      | $\sphericalcap$   | <i>incl</i>    | $\sqsubset$    |
| <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>     | $\nabla$              | <i>quarter</i>   | $\frac{1}{4}$     | <i>!subset</i> | $\subsetneq$   |
| <i>oppE</i>     | $\equiv$              | <i>3quarter</i>  | $\frac{3}{4}$     | <i>!supset</i> | $\supsetneq$   |
| <i>angstrom</i> | $\text{\AA}$          | <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   | text/doctools   |

**See Also** [eqn\(1\)](#), [nroff\(1\)](#), [troff\(1\)](#), [attributes\(5\)](#)

**Name** extendedFILE – enable extended FILE facility usage

**Synopsis** `$ ulimit -n N_file_descriptors`  
`$ LD_PRELOAD_32=/usr/lib/extendedFILE.so.1 application [arg...]`

**Description** The extendedFILE.so.1 is not a library but an enabler of the extended FILE facility.

The extended FILE facility allows 32-bit processes to use any valid file descriptor with the standard I/O (see [stdio\(3C\)](#)) C library functions. Historically, 32-bit applications have been limited to using the first 256 numerical file descriptors for use with standard I/O streams. By using the extended FILE facility this limitation is lifted. Any valid file descriptor can be used with standard I/O. See the NOTES section of [enable\\_extended\\_FILE\\_stdio\(3C\)](#).

The extended FILE facility is enabled from the shell level before an application is launched. The file descriptor limit must also be raised. The syntax for raising the file descriptor limit is

```
$ ulimit -n max_file_descriptors
$ LD_PRELOAD_32=/usr/lib/extendedFILE.so.1 application [arg...]
```

where *max\_file\_descriptors* is the maximum number of file descriptors desired. See [limit\(1\)](#). The maximum value is the same as the maximum value for [open\(2\)](#).

**Environment Variables** The following environment variables control the behavior of the extended FILE facility.

|                                  |                                                                                                                                                                                                                                                                                                                                        |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>_STDIO_BADFD</code>        | This variable takes an integer representing the lowest file descriptor, which will be made unallocatable. This action provides a protection mechanism so that applications that abuse interfaces do not experience silent data corruption. The value must be between 3 and 255 inclusive.                                              |
| <code>_STDIO_BADFD_SIGNAL</code> | This variable takes an integer or string representing any valid signal. See <a href="#">signal.h(3HEAD)</a> for valid values or strings. This environment variable causes the specified signal to be sent to the application if certain exceptional cases are detected during the use of this facility. The default signal is SIGABRT. |

**Examples** **EXAMPLE 1** Limit the number of file descriptors and FILE standard I/O structures.

The following example limits the number of file descriptors and FILE standard I/O structures to 1000.

```
$ ulimit -n 1000
$ LD_PRELOAD_32=/usr/lib/extendedFILE.so.1 application [arg...]
```

**EXAMPLE 2** Enable the extended FILE facility.

The following example enables the extended FILE facility. See [enable\\_extended\\_FILE\\_stdio\(3C\)](#) for more examples.

**EXAMPLE 2** Enable the extended FILE facility. *(Continued)*

```
$ ulimit -n 1000
$ _STDIO_BADFD=100 _STDIO_BADFD_SIGNAL=SIGABRT \
  LD_PRELOAD_32=/usr/lib/extendedFILE.so.1 \
  application [arg ...]
```

**EXAMPLE 3** Set up the extended FILE environment and start the application.

The following shell script first sets up the proper extended FILE environment and then starts the application:

```
#!/bin/sh
if [ $# = 0 ]; then
    echo "usage: $0 application [arguments...]"
    exit 1
fi
ulimit -n 1000
# _STDIO_BADFD=196; export _STDIO_BADFD
# _STDIO_BADFD_SIGNAL=SIGABRT; export _STDIO_BADFD_SIGNAL
LD_PRELOAD_32=/usr/lib/extendedFILE.so.1; export LD_PRELOAD_32
"$@"
```

**Files** /usr/lib/extendedFILE.so.1 enabling library

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE         |
|---------------------|-------------------------|
| Availability        | system/library (32-bit) |
| Interface Stability | Committed               |
| MT-Level            | Safe                    |

**See Also** [limit\(1\)](#), [open\(2\)](#), [enable\\_extended\\_FILE\\_stdio\(3C\)](#), [fdopen\(3C\)](#), [fopen\(3C\)](#), [popen\(3C\)](#), [signal.h\(3HEAD\)](#), [stdio\(3C\)](#), [attributes\(5\)](#)

**Warnings** The following displayed message

```
Application violated extended FILE safety mechanism.
Please read the man page for extendedFILE.
Aborting
```

is an indication that your application is modifying the internal file descriptor field of the FILE structure from standard I/O. Continued use of this extended FILE facility could harm your data. Do not use the extended FILE facility with your application.

**Name** extensions – localedef extensions description file

**Description** A localedef extensions description file or *extensions* file defines various extensions for the [localedef\(1\)](#) command.

The localedef extensions description file provides:

- EUC code set width information via the `cswidth` keyword: `cswidth bc1 : sw1, bc2 : sw2, bc3 : sw3` where `bc1`, `bc2`, and `bc3` indicate the number of bytes (byte count) per character for EUC codesets 1, 2, and 3, respectively. `sw1`, `sw2`, and `sw3` indicate screen width for EUC codesets 1, 2, and 3, respectively.
- Other extensions which will be documented in a future release.

**See Also** [locale\(1\)](#), [localedef\(1\)](#), [environ\(5\)](#), [locale\(5\)](#)

**Name** fedfs – overview of Federated Filesystem (FedFS)

**Description** The NFSv4 Working Group of the Internet Engineering Task Force (IETF) has been working to standardize a set of protocols which together permit construction and maintenance of a federated filesystem, where many different file servers together share a single namespace. These protocols make it possible to create a multi-vendor global namespace.

Solaris currently has support for two of the three FedFS protocols:

DNS SRV records

The Solaris automounter ([automount\(1M\)](#)) supports the `-fedfs` map, which by default will mount `/nfs4/dns_domain` by looking up a DNS SRV record for the domain. This map is enabled by default in `/etc/auto_master`.

LDAP-based referrals

The Solaris NFS server will follow FedFS referrals by looking up location data in LDAP, guided by the schema specified by the FedFS specs.

**Examples** EXAMPLE 1 Using the `-fedfs` Automounter Map

To mount a path like `/nfs4/cthon.org`, your NFS client should be set up to use DNS such that the domain's DNS server should answer a query for the `_nfs-domainroot._tcp` SRV record like this:

```
$ nslookup '-q=srv' _nfs-domainroot._tcp.cthon.org
Server:          1.2.3.4
Address:         1.2.3.4#53

_nfs-domainroot._tcp.cthon.org service = 0 0 2049 root-a.cthon.org.
_nfs-domainroot._tcp.cthon.org service = 1 1 2049 root-n.cthon.org.

$ cd /nfs4/cthon.org

$ df .
/nfs4/cthon.org (root-a.cthon.org,root-b.cthon.org:/domainroot/\
cthon.org):120379963 blocks 120379963 files
```

EXAMPLE 2 Using `nsdbparams` to Set Up Communications with an NSDB

The following example sets up communications with an NSDB called `nsdb.cthon.org` and makes it the default NSDB:

```
# nsdbparams update -D cn=Manager,dc=cthon,dc=org -w cthon.org \
nsdb.cthon.org
# nsdbparams show nikon.us.example.com
nikon.us.example.com:389
  default bind DN: cn=Manager,dc=cthon,dc=org
  default bind PW: cthon.org
  default NCE: dc=cthon,dc=org
  sectype: FEDFS_SEC_NONE
```

EXAMPLE 2 Using nsdbparams to Set Up Communications with an NSDB (Continued)

```
# nsdbparams set nsdb.cthon.org
# nsdbparams get
default nsdb: nsdb.cthon.org
default port: 389
```

EXAMPLE 3 Using the NSDB Tools

The following commands illustrate the use of the NSDB tools available in Oracle Solaris.

```
# nsdb-nces
Host: nsdb.cthon.org:389
  namingContext 'dc=cthon,dc=org' is a FedFS NCE, DIT starts at ''

# nsdb-list
SDB: nsdb.cthon.org:389, dc=cthon,dc=org
  FSN UUID: 7cc0bf04-5459-11e1-8083-80093d11d889
    FSL UUID: 7cc33c02-5459-11e1-8084-00093d11d889 = filer-a:/tmp
  FSN UUID: db48f160-5858-11e1-b459-80093d11d889
    FSL UUID: db4998c2-5858-11e1-b45a-00093d11d889 = filer-j:/tmp

# nsdb-resolve-fsn 7cc0bf04-5459-11e1-8083-80093d11d889
For FSN UUID 7cc0bf04-5459-11e1-8083-80093d11d889
  FSL UUID: 7cc33c02-5459-11e1-8084-00093d11d889
  Location: filer-a:/tmp
```

EXAMPLE 4 Using nfsref to Create a Referral

The following sequence of commands illustrates the use of nfsref to create a referral.

```
# nsdb-list
NSDB: nsdb.cthon.org:389, dc=cthon,dc=org
  FSN UUID: 7cc0bf04-5459-11e1-8083-80093d11d889
    FSL UUID: 7cc33c02-5459-11e1-8084-00093d11d889 = filer-a:/tmp
  FSN UUID: db48f160-5858-11e1-b459-80093d11d889
    FSL UUID: db4998c2-5858-11e1-b45a-00093d11d889 = filer-j:/tmp

# nfsref -t nfs-fedfs add /root/tools filer-k:/tools filer-x:/tools
Enter password for cn=Manager,dc=cthon,dc=org:
Created reparse point /root/tools

# nsdb-list
NSDB: nsdb.cthon.org:389, dc=cthon,dc=org
  FSN UUID: 7cc0bf04-5459-11e1-8083-80093d11d889
    FSL UUID: 7cc33c02-5459-11e1-8084-00093d11d889 = filer-a:/tmp
  FSN UUID: db48f160-5858-11e1-b459-80093d11d889
    FSL UUID: db4998c2-5858-11e1-b45a-00093d11d889 = filer-j:/tmp
  FSN UUID: 004b2382-9663-11e1-8c79-80093d11d888
    FSL UUID: 004da2b0-9663-11e1-8c7a-00093d11d888 = filer-k:/tools
```

EXAMPLE 4 Using nfsref to Create a Referral (Continued)

FSL UUID: 004e9bac-9663-11e1-8c7b-00093d11d888 = filer-x:/tools

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE         |
|----------------|-------------------------|
| Availability   | service/file-system/nfs |

**See Also** [automount\(1M\)](#), [nfsref\(1M\)](#), [nsdb-list\(1M\)](#), [nsdbparams\(1M\)](#), [attributes\(5\)](#)

RFC 5716: *Requirements for Federated File Systems*

**Name** filesystem – file system organization

**Synopsis** /

**Description** The file system is a hierarchical structure of descendent file systems and directories that are used to organize system-related components and binaries as well as non-system-related components and home directories in the Oracle Solaris OS. By default, the root file system is installed within a ZFS root pool and specifically, is a ZFS file system with separate directories of system-related components, such as `etc`, `usr`, and `var`, that must be available for the system to function correctly. After a system is installed, the root of the Solaris file system is mounted, which means files and directories are accessible.

All subdirectories of the root file system that are part of the Oracle Solaris OS, with the exception of `/var`, must be contained in the same file system as the root file system. In addition, all packaged software must reside in the root pool, with the exception of the swap and dump devices. A default swap device and dump device are created automatically as ZFS volumes in the root pool when a system is installed.

The following file system content descriptions make use of platform, platform-dependent, platform-independent, and platform-specific terms. Platform refers to a system's Instruction Set Architecture or processor type as returned by the `uname -i` command.

Platform-dependent refers to a file that is installed on all platforms and whose contents vary depending on the platform. Like a platform-dependent file, a platform-independent file is installed on all platforms. However, the contents of the latter type remains the same on all platforms. An example of a platform-dependent file is compiled, executable program. An example of a platform-independent file is a standard configuration file, such as `/etc/hosts`. Unlike a platform-dependent or a platform-independent file, the platform-specific file is installed only on a subset of supported platforms. Most platform-specific files are gathered under `/platform` and `/usr/platform`.

**Root File System** The root file system contains files and directories that are critical for system operation, such as the kernel, the device drivers, and the programs used to boot the system. These components are described below. The root (`/`) directory also contains mount point directories where local and remote file systems can be attached to the file system hierarchy.

`/`

Root directory of the entire file system name space. This is a special file system that is mounted by the kernel at system boot time.

`/boot`

Directory that contains files and executables that are needed for booting the system.

`/bin`

Symbolic link to the `/usr/bin` directory that contains system executables and scripts.

`/dev`

Directory that contains special device files. Typically, device files are built to match the kernel and hardware configuration of the system.

**/devices**

Mount point directory for the `devfs` file system that manages the device name space.

**/etc**

Directory that contains platform-dependent administrative and configuration files and databases that are not shared among systems. This directory defines the system's identity. An approved installation location for bundled Solaris software.

**/export/home or /home**

Directory or file system mount point for user home directories, which store user files. By default, the `/home` directory is an automounted file system.

**/kernel**

Directory 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 Oracle Solaris software and for add-on system software.

**/lib**

Directory that contains core system libraries. Historically, this directory contained essential library components for system startup.

**/media**

Directory for accessing removable media that is automatically mounted.

**/mnt**

Default temporary mount point directory for file systems. This empty directory is used to temporarily mount a file system.

**/net**

Temporary mount point directory for file systems that are mounted by the automounter.

**/opt**

Directory for unbundled application packages.

**/platform**

Directory of platform-specific objects that need to reside in the root file system. It contains a series of directories, one per supported platform. The semantics of the series of directories is equivalent to `/` (root).

**/proc**

Mount point directory for the process file system.

**/root**

Home directory for the root user.

**/rpool**

Mount point directory for the ZFS boot-related components. By default, the root pool is named `rpool` during installation.

- `/sbin`  
Symbolic link to the `/usr/sbin` directory.
- `/system`  
Mount point directory for the contract (CTFS) and object (OBJFS) file systems.
- `/tmp`  
Directory that contains temporary files that are removed during a boot operation.
- `/usr`  
Directory that contains platform-dependent and platform-independent binaries and files. The `/usr/share` subdirectory contains platform-independent files. The rest of the `/usr` directory contains platform-dependent files.
- `/usr/bin`  
Directory that contains platform-dependent, user-invoked executables. These are commands that users expect to be run as part of their normal `$PATH`. An approved installation location for bundled Oracle Solaris software. The analogous location for unbundled system software or for applications is `/opt/package/bname/bin`.
- `/usr/lib`  
Directory that augments the contents of `/lib` with additional system libraries, and other supporting files that are required by programs at runtime.
- `/usr/java`  
Directory that contains Java files and executables.
- `/usr/sbin`  
Contains essential executables used in the booting process and in manual system recovery. Historically, this directory was needed to recover the system before the `/usr` file system was mounted. In this Oracle Solaris release, `/usr` is a directory, not a separate file system and is available when the root file system is mounted.
- `/var`  
Directory or file system that contains varying files that are unique to a system but can grow to an arbitrary or variable size. An example is a log file. An approved installation location for bundled Oracle Solaris software.
- `/var/share`  
Directory that is shared across several boot environments, and is resident on a separate dataset beneath the `zpool` containing boot environments. Packages should not deliver content here, since package metadata in a given boot environment may not reflect the current content of `/var/share`.
- Packages may deliver directories to `/var/.migrate`, which will be created automatically beneath `/var/share` during boot. See `pkg(5)` and the IPS Developer's Guide for a description of how to share data across boot environments.

By default, the following directories are shared:

```
/var/share/audit
/var/share/cores
/var/share/crash
/var/share/mail
```

Symlinks are delivered to `/var` to point to each shared directory.

If existing datasets attempt to mount on one of those symlinks (for example, users with an existing dataset that normally mounts on `/var/mail`), then those datasets will be mounted beneath `/var/share`, since filesystem mounts traverse symlinks. The dataset will still be accessible through the original mountpoint.

`/var/tmp`

Directory that contains files that vary in size or presence during normal system operations. The content of this directory is not removed during a boot operation. It is possible to change the default behavior for `/var/tmp` to clear all of the files except editor temporary files by setting the `clean_vartmp` property value of the `rmtmpfiles` service. This is done with the following commands:

```
# svccfg -s svc:/system/rmtmpfiles setprop\
    options/clean_vartmp = "true"
# svcadm refresh svc:/system/rmtmpfiles:default
```

The `solaris.smf.value.rmtmpfiles` authorization is required to modify this property.

**See Also** [isainfo\(1\)](#), [svcs\(1\)](#), [uname\(1\)](#), [automount\(1M\)](#), [automountd\(1M\)](#), [boot\(1M\)](#), [init\(1M\)](#), [kernel\(1M\)](#), [mount\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [zfs\(1M\)](#), [zpool\(1M\)](#), [mount\(2\)](#), [Intro\(4\)](#), [proc\(4\)](#), [ctfs\(7FS\)](#), [devfs\(7FS\)](#), [objfs\(7FS\)](#),

**Name** fmri – Fault Managed Resource Identifier

**Description** Oracle Solaris subsystems such as the Service Management Facility (see [smf\(5\)](#)), the Fault Management Daemon (see [fmd\(1M\)](#), [fmadm\(1M\)](#)), and the Image Packaging System (see [pkg\(5\)](#)) identify each element they manage using a *fault managed resource identifier*, abbreviated as FMRI.

FMRI are used to identify both hardware resources and software resources and abstractions. For hardware resources the FMRI are usually an implementation detail, and when a component is diagnosed as faulty it is the FRU (field replaceable unit) location label which is used to identify it. If no FRU label is available, the FMRI is used. For software abstractions such as an SMF service or a [pkg\(5\)](#) software package the FMRI (often in string form) is the one and only means of naming the managed resource (abbreviations aside).

All FMRI include an indication of the FMRI scheme adhered to, and the version of that FMRI scheme in use. Once the scheme and version are known, the remainder of the FMRI is interpreted as specified in the formal definition of that FMRI scheme version. Schemes exists for `svc`, `pkg`, `cpu`, `hc` (hardware component), `dev` (device), and a number of others.

The formal definition of a given version of a particular FMRI scheme is specified as an unordered series of FMRI members specified as (member-name, data-type, data-semantics) tuples. An instance of an FMRI is a series of (name, type, value) tuples wherein the name and type match the FMRI specification and the value provided is a value of the given type and derived and interpreted according to the semantics described in the formal FMRI definition for that scheme and version. In Oracle Solaris C library APIs such as [libfmevent\(3LIB\)](#), FMRI are represented as `nvLists` using the [libnvpair\(3LIB\)](#) library.

It is often more convenient to represent a given FMRI instance as a string, such as in command line interaction or in command output. The specification of each FMRI scheme version includes rules for formatting an FMRI of that scheme and version as a string instead of an name-value pair list, and given an FMRI string those same rules can be applied in reverse to recover the more-formal FMRI form. In some cases the string form of an FMRI can include a reduced amount of information as compared to the formal form, but still provide enough information to identify the resource for administrative purposes. The string form of an FMRI usually does not include the FMRI version number, and so the version is understood to be that whose formal FMRI-to-string formatting rules the given FMRI string matches. The string form of an FMRI is not unlike a URI as outlined in [RFC 2396](#) which we use as a syntactical model. In particular, we reserve the same set of characters `:`, `;`, `=`, `@`, `/`, `&`, `'`, `$` and `,` and require the same escaped encoding conventions should these characters appear in an FMRI member value that is rendered in the string.

**FMRI Members** The formal definition of a given version of a particular FMRI scheme is specified as an unordered series of:

```
(member-name, data-type, data-semantics)
```

tuples, some of which are required and some of which are optional. The formal definition also specifies how to render the FMRI in string form.

**member-name**

This associates a name with the FMRI member, with characters drawn from the following alphabet:

A-Z  
a-z  
'\_' (underscore)  
'-' (hyphen)  
'.' (period)

**data-type**

This specifies the data type for this member. Legal primitive datatypes are listed in the following table, and for all primitive types we also define an array of that type as a legal datatype. The primitive integer and double datatypes are defined as in C, and we add boolean value and string primitive datatypes:

| Primitive Datatype | Description                     |
|--------------------|---------------------------------|
| boolean_value      | Values can be 0 or 1            |
| string             | ASCII string                    |
| [u]int8            | (Un)signed 8-bit integer        |
| [u]int16           | (Un)signed 16-bit integer       |
| [u]int32           | (Un)signed 32-bit integer       |
| [u]int64           | (Un)signed 64-bit integer       |
| double             | Double precision floating point |

For each primitive datatype we form an aggregate datatype that is an ordered array of the primitive type:

| Primitive Datatype  | Description                        |
|---------------------|------------------------------------|
| boolean_value array | Array of boolean_value             |
| string array        | Array of ASCII strings             |
| [u]int8 array       | Array of (un)signed 8-bit integer  |
| [u]int16 array      | Array of (un)signed 16-bit integer |
| [u]int32 array      | Array of (un)signed 32-bit integer |
| [u]int64 array      | Array of (un)signed 64-bit integer |
| double array        | Array of double                    |

A further aggregate datatype is defined - the `nvlist` (name- value list). An `nvlist` is an unordered list of (`list-member-name`, `list-member-datatype`) in which `list-member-name` names the member using the same alphabet as for `member-name` above, and `list-member-datatype` is any primitive or aggregate datatype including `nvlist` itself. We term the members of an `nvlist`: `nvpairs` (name-value pairs).

A special type of `nvlist` is an `fmri` which is an `nvlist` that conforms to one of the FMRI scheme versions defined below.

| Synthetic Datatype  | Description                            |
|---------------------|----------------------------------------|
| <code>nvlist</code> | Unordered list of <code>nvpairs</code> |

```

nvlst array      Array of nvlst
fmri             An nvlst that specifies an FMRI
                 scheme and version and includes other
                 members so as to be a valid FMRI
                 in that scheme version.

```

#### data-`semantics`

The data semantics for a given member in an FMRI definition specify how the values for that member should be interpreted, that is, their meaning. The data semantics can constrain the set of legal values that the member can assume to a subset of those representable by the member data type.

**Stabilities** Some FMRI schemes, such as `svc` for SMF, are administrator-facing while others are more of a private implementation detail. Correspondingly, the definition of some FMRI schemes constitute a `Committed` interface, while the definition of others form a `Private` interface. See [attributes\(5\)](#). As with all `Private` interfaces the definition of `Private` FMRI components is subject to possibly incompatible change at any time, and you should not consume such interfaces nor rely on the presence or assumed semantics of private FMRI members.

In defining an FMRI scheme below we list a series of (`member-name`, `data-type`, `data-semantics`), specify which members are required and which are optional, and the rules for rendering an FMRI in string form. In terms of stability assignments, there are two stability levels associated with each member:

- Stability of the (`member-name`, `data-type`) pair
- Stability of the corresponding data-semantics

The default in both cases is `Private`. When the stability of both aspects of a member are identical (the common case) we write the stability level just once. When they differ then we write the member name and data type stability first, followed by a comma and the data semantics stability level. Today the only stability levels employed in FMRI are `Committed`, `Private`, and `Not An Interface` but others are allowed. The semantics of a `Committed` FMRI component do not change incompatibly for given scheme and version thereof.

#### FMRI Scheme and Version Identification

All FMRI definitions are required to include members for the scheme name and major version thereof. The name/datatype and data semantics stabilities are `Committed`, thus all FMRI definitions are required to begin as follows:

| Member Name          | Data Type           | Stability              |
|----------------------|---------------------|------------------------|
| -----                | -----               | -----                  |
| <code>scheme</code>  | <code>string</code> | <code>Committed</code> |
| <code>version</code> | <code>uint8</code>  | <code>Committed</code> |

```

scheme
  Names this FMRI scheme

version
  Scheme major version

```

Together these select the set of rules (as below) by which the given FMRI should be interpreted. A small number of scheme names are defined to cope with identifying resources within differing subsystems, and each scheme is versioned starting at version number 0.

When using the informal string representation of an FMRI, the scheme is always written but the scheme version in use is usually not included. The version in use is understood to be the latest for which the given FMRI string matches the FMRI-to-string conversion rules.

FMRI Authority A further required, although implicitly present in some cases, FMRI member is authority:

| Member Name | Data Type | Stability           |
|-------------|-----------|---------------------|
| [authority] | nvlist    | Committed, (Varies) |

Every FMRI includes authority information, either explicitly with the `authority nvlist` if present in the FMRI or implicitly that of the local fault management domain if not present. Authority information has two flavors:

- High-level authority information, such as to identify the platform and hostname within which a particular resource that is the subject of an FMRI is located
- Leaf-level authority information, providing identity information for the specific resource that is the subject of the FMRI

For example, an FMRI identifying a disk can have high-level authority information detailing the platform (host name, product name, chassis serial) in which the disk resides, and leaf-level authority information detailing the individual disk part number, serial, and so forth.

It is an unfortunate fact that some FMRI schemes have evolved to include leaf-level authority information outside of the `authority nvlist` as standalone FMRI members, and written as part of the path component in string form.

In their string representation, both high-level and leaf-level, FMRI authority information is prefaced with a colon (:), for example:

```
hc://:chassis-serial=12345/<string>
```

“:chassis-serial=12345” is a string represented authority member.

There are two versions of the high level authority (version 0 and version 1). Following are the descriptions for each.

The leaf-level authority members are described in the scheme specific sections below.

#### Version 0

| Member Name  | Data Type | Stability                   |
|--------------|-----------|-----------------------------|
| [server-id]  | string    | Committed, Not an Interface |
| [chassis-id] | string    | Committed, Not an Interface |

```
[product-sn] string    Committed, Not an Interface
[product-id] string    Committed, Not an Interface
[domain-id]  string    Committed, Not an Interface
[host-id]    string    Committed, Not an Interface
```

[server-id]

The hostname (uname -n) string for the entity on which the fault manager is running. Not invariant.

[chassis-id]

This is a serial number identifying the chassis within which the resource resides. Some large systems can include multiple distinct chassis components.

[product-sn]

In some cases chassis-id alone does not fully identify the system, for example, for service entitlement keyed by a product serial number and where the product can comprise multiple chassis elements. Invariant.

[product-id]

The product name string, such as Sun-Fire-X4600. Invariant.

[domain-id]

Identifies a particular hardware domain in the case of hardware that supports this feature. Not invariant.

[host-id]

The `hostid(1)` string in Oracle Solaris); not generally used. Not necessarily invariant.

### Version 1

| Member Name       | Data Type | Stability                   |
|-------------------|-----------|-----------------------------|
| [system-mfg]      | string    | Committed, Not an Interface |
| [system-name]     | string    | Committed, Not an Interface |
| [system-part]     | string    | Committed, Not an Interface |
| [system-serial]   | string    | Committed, Not an Interface |
| [sys-comp-mfg]    | string    | Committed, Not an Interface |
| [sys-comp-name]   | string    | Committed, Not an Interface |
| [sys-comp-part]   | string    | Committed, Not an Interface |
| [sys-comp-serial] | string    | Committed, Not an Interface |
| [chassis-mfg]     | string    | Committed, Not an Interface |
| [chassis-name]    | string    | Committed, Not an Interface |
| [chassis-part]    | string    | Committed, Not an Interface |
| [chassis-serial]  | string    | Committed, Not an Interface |
| [chassis-alias]   | string    | Committed, Not an Interface |
| [server-name]     | string    | Committed, Not an Interface |
| [domain-name]     | string    | Committed, Not an Interface |
| [host-id]         | string    | Committed, Not an Interface |

**[system-mfg]**

The manufacturer of the top-level product. In a complex product the system-\* members contain the service entitlement data. In a simple standalone rack system the system-\*, sys-comp-\* and chassis-\* data is duplicated. Invariant.

**[system-name]**

The model name of the top-level product. Invariant.

**[system-part]**

The part number of the top-level product. Invariant.

**[system-serial]**

The serial number of the top-level product. Invariant.

**[sys-comp-mfg]**

The manufacturer of the component system-level product. Invariant.

**[sys-comp-name]**

The model name of the component system-level product. Invariant.

**[sys-comp-part]**

The part number of the component system-level product. Invariant.

**[sys-comp-serial]**

The serial number of the component system-level product. Invariant.

**[chassis-mfg]**

The manufacturer of the chassis within which the resource resides. Some large systems can include multiple distinct chassis components. Invariant.

**[chassis-name]**

The model name of the chassis. Invariant.

**[chassis-part]**

The part number of the chassis. Invariant.

**[chassis-serial]**

The serial number of the chassis. Invariant.

In some cases a single set of \*-name/\*-part/\*-serial alone does not fully identify the system, for example, for service entitlement keyed by the product serial number and where the product can comprise multiple chassis elements. Invariant.

**[chassis-alias]**

This is the Chassis Receptacle Occupant alias value assigned to this chassis via 'fmadm add-alias'. Not invariant.

**[server-name]**

The hostname (uname -n) string for the entity on which the fault manager is running. Not invariant.

[domain-name]

Identifies a particular hardware domain in the case of hardware that supports this feature.  
Not invariant.

[host-id]

The `hostid(1)` string in Oracle Solaris); not generally used. Not necessarily invariant.

Logical vs. Universal  
FMRI Schemes

A logical FMRI scheme defines FMRI's that can only meaningfully be interpreted within the fault management domain (typically an Oracle Solaris instance) in which they were generated. Identical FMRI's of a logical scheme that are native to distinct fault management domains do not necessarily identify the same actual resource. For example a `cpu id 1` using the `cpu` scheme must be interpreted in the Oracle Solaris instance that generated the FMRI (that is, its native context), and many other Oracle Solaris instances also have a logical processor id number 1 but all such FMRI's do not identify the same actual processor. The native fault management domain is identified by its high-level authority information which is typically not captured in an FMRI instance but instead understood to be that of the local Oracle Solaris instance.

A universal FMRI scheme identifies resources in a universally unique manner, and two identical FMRI's in a universal scheme identify the same actual resource wherever they are interpreted. Such schemes are used when ambiguity must be avoided, such as in identifying hardware components that are faulted.

FMRI schemes do not include a member indicating whether the scheme is logical or universal.

FMRI Schemes and  
Versions

The following FMRI scheme versions are defined in the sections that follow:

| Scheme           | Version(s) | Universal? | Description                      |
|------------------|------------|------------|----------------------------------|
| <code>cpu</code> | 0, 1       | No         | Logical processor identification |
| <code>dev</code> | 0          | No         | Device resources                 |
| <code>fmd</code> | 0, 1       | No         | Fault Management Daemon modules  |
| <code>hc</code>  | 0, 1       | Yes        | Hardware Components              |
| <code>mem</code> | 0          | No         | Memory modules and cache         |
| <code>mod</code> | 0          | No         | Kernel modules                   |
| <code>pkg</code> | 0, 1       | Yes        | Software packages                |
| <code>svc</code> | 0          | No         | SMF services & service instances |
| <code>sw</code>  | 0          | No         | Software objects                 |
| <code>zfs</code> | 0          | No         | zfs filesystem resources         |

Of these, the `svc` and `pkg` schemes are the most-likely to be encountered in day-to-day Oracle Solaris use, while the remainder are usually hidden behind command line interaction with the fault management subsystem when problems have been diagnosed.

### SCHEME `cpu` VERSION 0

The `cpu` scheme is used to identify the logical Oracle Solaris processor resource that can be affected by a problem that has been diagnosed, and which can be offlined in response to that problem diagnosis. It does not necessarily identify the physical resource involved, and should not be used to identify parts for replacement.

The `cpu` scheme is a logical FMRI scheme; an FMRI in the `cpu` scheme is meaningful only within the fault management domain identified by the authority information.

| Member Name           | Data Type | Stability              |
|-----------------------|-----------|------------------------|
| <code>scheme</code>   | string    | Committed, value "cpu" |
| <code>version</code>  | uint8     | Committed, value 0     |
| <code>cpuid</code>    | uint32    | Committed              |
| <code>[serial]</code> | uint64    | Private                |

#### `cpuid`

The Oracle Solaris processor instance id of this logical execution unit, as listed by `psradm(1M)` on the Oracle Solaris instance identified by the authority information.

#### `[serial]`

The semantics of this Private FMRI member are not documented.

String form:

```
cpu:///cpuid=<cpuid>[/serial=<hex-serial>]
```

<hex-serial> is the *serial* member written in hexadecimal with no 0x prefix.

### SCHEME `cpu` VERSION 1

Version 1 of the `cpu` FMRI scheme has the same use as version 0, with some private FMRI member changes and additions.

| Member Name               | Data Type | Stability              |
|---------------------------|-----------|------------------------|
| <code>scheme</code>       | string    | Committed, value "cpu" |
| <code>version</code>      | uint8     | Committed, value 1     |
| <code>cpuid</code>        | uint32    | Committed              |
| <code>[serial]</code>     | string    | Private                |
| <code>[cpumask]</code>    | uint8     | Private                |
| <code>[cpufru]</code>     | string    | Private                |
| <code>[cacheindex]</code> | uint32    | Private                |
| <code>[cacheway]</code>   | uint32    | Private                |
| <code>[cachebit]</code>   | uint16    | Private                |
| <code>[cachetype]</code>  | uint8     | Private                |

#### `cpuid`

As per `cpu` scheme version 0

#### `[serial]`

```
[cpumask]
[cpufru]
[cacheindex]
[cacheway]
[cachebit]
[cachetype]
```

The semantics of these Private FMRI members are not documented.

String form:

```
cpu:///cpuid=<cpuid>[/serial=<serial>]
      [/cacheindex=.../cacheway=.../
      cachebit=.../cachetype=...]
```

Members `cpumask` and `cpufru` do not appear in the string form.

### SCHEME dev VERSION 0

The `dev` scheme is used to identify devices. It is not an invariant scheme. A `dev` scheme FMRI should be interpreted only in the fault management domain identified within the authority information.

| Member Name                     | Data Type | Stability              |
|---------------------------------|-----------|------------------------|
| -----                           | -----     | -----                  |
| <code>scheme</code>             | string    | Committed, value "dev" |
| <code>version</code>            | uint8     | Committed, value 0     |
| <code>device-path</code>        | string    | Committed              |
| <code>[devid]</code>            | string    | Private                |
| <code>[target-port-l0id]</code> | string    | Private                |

#### device-path

A filesystem path within the Oracle Solaris `/devices` tree that identifies a device node, such as `pci@0,0/pci108e,c884@2,1`. Beyond this interpretation that this member identifies a `/devices` path, the internal structure of such paths, that is, the structure and node naming of the Oracle Solaris `/devices` tree) is Not An Interface.

`[devid]`

`[target-port-l0id]`

The semantics of these Private FMRI members are not documented.

String form:

```
dev:///[:devid=<devid>]
      [:target-port-l0id=<target-port-l0id>]
      <device-path>
```

### SCHEME fmd VERSION 0

This scheme is used to identify Oracle Solaris Fault Manager (`fmd(1M)`) modules.

High-level authority:

| Member Name  | Data Type | Stability                   |
|--------------|-----------|-----------------------------|
| [server-id]  | string    | Committed, Not an Interface |
| [chassis-id] | string    | Committed, Not an Interface |
| [product-sn] | string    | Committed, Not an Interface |
| [product-id] | string    | Committed, Not an Interface |
| [domain-id]  | string    | Committed, Not an Interface |

FMRI does not include any leaf-level authority information within the authority nvlist, instead it is featured as top-level members of the FMRI.

| Member Name | Data Type | Stability                   |
|-------------|-----------|-----------------------------|
| scheme      | string    | Committed, value "fmd"      |
| version     | uint8     | Committed, value 0          |
| [authority] | nvlist    | Committed, See above        |
| mod-name    | string    | Committed, Not an Interface |
| mod-version | string    | Committed, Not an Interface |

mod-name

The fmd module name, an opaque string.

mod-version

The fmd module version, an opaque string.

String form:

fmd:///module/mod-name>

### SCHEME fmd VERSION 1

This scheme is used to identify Oracle Solaris Fault Manager ([fmd\(1M\)](#)) modules.

High-level authority:

| Member Name       | Data Type | Stability                   |
|-------------------|-----------|-----------------------------|
| [system-mfg]      | string    | Committed, Not an Interface |
| [system-name]     | string    | Committed, Not an Interface |
| [system-part]     | string    | Committed, Not an Interface |
| [system-serial]   | string    | Committed, Not an Interface |
| [sys-comp-mfg]    | string    | Committed, Not an Interface |
| [sys-comp-name]   | string    | Committed, Not an Interface |
| [sys-comp-part]   | string    | Committed, Not an Interface |
| [sys-comp-serial] | string    | Committed, Not an Interface |
| [server-name]     | string    | Committed, Not an Interface |
| [domain-name]     | string    | Committed, Not an Interface |
| [host-id]         | string    | Committed, Not an Interface |

As with version 0 the FMRI does not contain leaf-level authority information within the authority nvlist. Instead it is featured as top-level members of the FMRI.

| Member Name | Data Type | Stability                   |
|-------------|-----------|-----------------------------|
| -----       | -----     | -----                       |
| scheme      | string    | Committed, value "fmd"      |
| version     | uint8     | Committed, value 1          |
| [authority] | nvlist    | Committed, See above        |
| mod-name    | string    | Committed, Not an Interface |
| mod-version | string    | Committed, Not an Interface |

mod-name

The fmd module name, an opaque string.

mod-version

The fmd module version, an opaque string.

String form:

```
fmd:///module/<mod-name>
```

### SCHEME legacy-hc VERSION 0

The hc scheme is used to identify hardware components. In most cases these FMRI should remain as an internal implementation detail, and where a hardware component is diagnosed as faulty it is identified (for example, in [fmadm\(1M\)](#)) by its FRU label or location label.

There is only one version of the hc scheme in use (version 0) but we differentiate two minor versions, distinguished by the presence or absence of a member named component. An hc scheme FMRI that includes a component member is an instance of the obsoleted legacy hc scheme.

Obsolete legacy hc variant:

| Member Name | Data Type | Stability             |
|-------------|-----------|-----------------------|
| -----       | -----     | -----                 |
| scheme      | string    | Committed, value "hc" |
| version     | uint8     | Committed, value 0    |
| component   | string    | Committed             |

component

The FRU label for the resource.

String form for legacy hc:

```
hc:///component=<component>
```

### SCHEME hc VERSION 0

The authority `nvlist` is the high-level authority information concerned with identifying the fault management domain in which a given FMRI was generated. Thus it includes members for the hostname, product name, chassis serial number and so on. Some of these members are invariant (such as platform serial number) and serve uniquely to identify some element, while others (such as hostname) are a soft identification which could change over time (albeit infrequently).

| Member Name  | Data Type | Stability                   |
|--------------|-----------|-----------------------------|
| [server-id]  | string    | Committed, Not an Interface |
| [chassis-id] | string    | Committed, Not an Interface |
| [product-sn] | string    | Committed, Not an Interface |
| [product-id] | string    | Committed, Not an Interface |
| [domain-id]  | string    | Committed, Not an Interface |
| [host-id]    | string    | Committed, Not an Interface |

While all members are listed as optional, as many as make sense for a given FMRI should be included. If the authority `nvlist` is absent in an FMRI then the authority is understood to be the local fault management domain; if an `hc` scheme FMRI is transmitted outside of its native domain it should include explicit high-level authority information.

The Not an Interface stability for the data semantics of each means that the value can be treated as an opaque string only and any internal structure to the string is not an interface. For example a serial number can be matched using `string compare`, but the length of the serial number string and the meaning of individual characters therein is not an interface.

String form:

These authority members, those that are present in a given FMRI, can be written in string form as an unordered series of colon-separated `name=value` pairs. We'll define this as `authstr`, to be used in the specification of the string form for individual FMRI schemes below which import the default authority information.

```
authstr = server-id=<server-id>
          [:chassis-id=<chassis-id>]
          [:product-sn=<product-sn>]
          [:product-id=<product-id>]
          [:domain-id=<domain-id>]
          [:host-id=<host-id>]
          [:devid=<devid>]
```

There is no leaf-level authority/identity information included in the authority `nvlist` - instead such information (serial and part numbers) feature as top-level members of the FMRI.

| Member Name | Data Type    | Stability                   |
|-------------|--------------|-----------------------------|
| scheme      | string       | Committed, value "hc"       |
| version     | uint8        | Committed, value 0          |
| [authority] | nvlist       | Committed, See above        |
| [serial]    | string       | Committed, Not-an-Interface |
| [devid]     | string       | Private                     |
| [part]      | string       | Committed, Not-an-Interface |
| [revision]  | string       | Committed, Not-an-Interface |
| [hc-root]   | string       | Private                     |
| hc-list-sz  | uint32       | Private                     |
| hc-list     | nvlist array | Private                     |

```
[facility]    nvlist    Private
[hc-specific] nvlist    Private
```

```
[serial]
```

The serial number string of the hardware component identified in this FMRI. This is an opaque string whose internal structure is not an interface.

```
[devid]
```

The semantics of these Private FMRI members are not documented.

```
[part]
```

The part number string of the hardware component identified in this FMRI. This is an opaque string whose internal structure is not an interface.

```
[revision]
```

The revision of the part number, as a string. Again this is an opaque string.

```
[hc-root]
```

The semantics of this Private FMRI member are not documented.

```
hc-list-sz
```

The semantics of this Private FMRI member are not documented.

```
hc-list
```

The semantics of this Private FMRI member are not documented, however we enumerate the structure of the content (subject to incompatible change). Each member of the (ordered) `nvlist` array has the following structure:

| Member Name | Data Type | Stability |
|-------------|-----------|-----------|
| hc-name     | string    | Private   |
| hc-id       | string    | Private   |

```
[facility]
```

The semantics of this Private FMRI member are not documented, however we enumerate the structure of the content (subject to incompatible change). The `facility nvlist` has members as follows:

| Member Name   | Data Type | Stability |
|---------------|-----------|-----------|
| facility-type | string    | Private   |
| facility-name | string    | Private   |

```
[hc-specific]
```

The semantics of this Private FMRI member are not documented.

String form:

```
authstr is defined above
hcauthstr = [:serial=<serial>]
           [:devid=<devid>]
           [:part=<part>]
```

```

[:revision=<revision>]

hc-path =<hc-list[0].hc-name>=<hc-list[0].hc-id>
        /<hc-list[1].hc-name>=<hc-list[1].hc-id>
        /<hc-list[2].hc-name>=<hc-list[2].hc-id>
        ...
hcspecstr = private string formed from hc-specific nvlist,
           of the form "/name=value"

facstr = <facility.facility-type>=<facility.facility-name>

hc://[authstr][hcauthstr]/[hc-root/]<hc-path>
                                   [<hcspecstr>]
                                   [?<facstr>]

```

The high-level authority information from the authority nvlist appears adjacent to the leaf-level authority information of serial and part number, all colon separated. The components of this authority string can appear in any order, with high-level and leaf authority elements interleaved. Parsing such a string is clumsy and relies on knowing in advance which components belong to which class of authority information.

### SCHEME hc VERSION 1

Version 1 of the hc FMRI scheme has the same use as version 0, with some private FMRI member changes and additions.

| Member Name      | Data Type | Stability                   |
|------------------|-----------|-----------------------------|
| [chassis-mfg]    | string    | Committed, Not an Interface |
| [chassis-name]   | string    | Committed, Not an Interface |
| [chassis-part]   | string    | Committed, Not an Interface |
| [chassis-serial] | string    | Committed, Not an Interface |
| [chassis-alias]  | string    | Committed, Not an Interface |

String form:

The same semantics apply to the hc FMRI scheme version 1 string form as version 0.

```

authstr = [:chassis-mfg=<chassis-mfg>]
          [:chassis-name=<chassis-name>]
          [:chassis-part=<chassis-part>]
          [:chassis-serial=<chassis-serial>]
          [:chassis-alias=<chassis-alias>]

```

As with hc scheme version 0 there is no leaf-level authority/identity information included in the authority nvlist with version 1 - instead such information (for example, serial/revision/part numbers) feature as top-level members of the FMRI.

| Member Name | Data Type | Stability             |
|-------------|-----------|-----------------------|
| scheme      | string    | Committed, value "hc" |

---

```

version      uint8      Committed, value 1
[authority]  nvlist     Committed, See above
[fru-mfg]    string    Committed, Not-an-Interface
[fru-serial] string    Committed, Not-an-Interface
[fru-part]   string    Committed, Not-an-Interface
[fru-revision] string    Committed, Not-an-Interface
[devid]      string    Private
[hc-root]    string    Private
hc-list-sz   uint32     Private
hc-list      nvlist array Private
[facility]    nvlist     Private
[hc-specific] nvlist     Private

```

**[fru-mfg]**

The manufacturer string of the hardware component identified in this FMRI. This is an opaque string whose internal structure is not an interface.

**[fru-serial]**

The serial number string of the hardware component identified in this FMRI. This is an opaque string whose internal structure is not an interface.

**[fru-part]**

The part number string of the hardware component identified in this FMRI. This is an opaque string whose internal structure is not an interface.

**[fru-revision]**

The revision of the part number, as a string. Again this is an opaque string.

**[devid]**

The semantics of these Private FMRI members are not documented.

**[hc-root]**

The semantics of this Private FMRI member are not documented.

**hc-list-sz**

The semantics of this Private FMRI member are not documented.

**hc-list**

The semantics of this Private FMRI member are not documented, however we enumerate the structure of the content (subject to incompatible change). Each member of the (ordered) nvlist array has the following structure:

| Member Name | Data Type | Stability |
|-------------|-----------|-----------|
| -----       |           |           |
| hc-name     | string    | Private   |
| hc-id       | string    | Private   |

**[facility]**

The semantics of this Private FMRI member are not documented, however we enumerate the structure of the content (subject to incompatible change). The facility nvlist has members as follows:

| Member Name   | Data Type | Stability |
|---------------|-----------|-----------|
| facility-type | string    | Private   |
| facility-name | string    | Private   |

[hc-specific]

The semantics of this Private FMRI member are not documented.

String form:

```

authstr is defined above
hcauthstr = [:fru-mfg=<fru-mfg>
             [:fru-serial=<fru-serial>]
             [:devid=<devid>]
             [:fru-part=<fru-part>]
             [:fru-revision=<fru-revision>]

hc-path =<hc-list[0].hc-name>=<hc-list[0].hc-id>
         /<hc-list[1].hc-name>=<hc-list[1].hc-id>
         /<hc-list[2].hc-name>=<hc-list[2].hc-id>
         ...
hcspecstr = private string formed from hc-specific nvlist,
           of the form "/name=value"

facstr = <facility.facility-type>=<facility.facility-name>

hc://[authstr][hcauthstr]/[hc-root/]<hc-path>
                                     [<hcspecstr>]
                                     [?!<facstr>]

```

The high-level authority information from the authority nvlist appears adjacent to the leaf-level authority information of serial and part number, all colon separated. The components of this authority string can appear in any order, with high-level and leaf authority elements interleaved. Parsing such a string is clumsy and relies on knowing in advance which components belong to which class of authority information.

### SCHEME mem VERSION 0

This scheme is used to identify memory modules and memory pages, although more recently these have been identified in the hc scheme. The mem scheme has also been used to identify cpu cache resources.

| Member Name | Data Type    | Stability              |
|-------------|--------------|------------------------|
| scheme      | string       | Committed, value "mem" |
| version     | uint8        | Committed, value 0     |
| unum        | string       | Committed, Private     |
| [serial]    | string array | Private                |
| [physaddr]  | uint64       | Private                |
| [offset]    | uint64       | Private                |

**unum**

This string identifies the memory module, generally as labelled on the platform silkscreens or service label. The presence of this member of type string is a Committed interface, but the internal structure of the unum string is a Private interface. Therefore, you can use the opaque unum string to identify a memory module, but the interpretation of the components of a unum string is platform dependent and Private.

A unum can identify a group (or bank) of memory modules. These should instead have been identified by multiple FMRI, one for each memory module, but this is an historical anomaly from early implementations.

**[serial]**

A string array of the serial number(s) of the memory module(s) named by the unum member. The internal structure of a serial number string is not an interface so you can use only the opaque string value in matching serial number and not apply any further interpretation.

**[physaddr]****[offset]**

The semantics of these Private FMRI members are not documented.

**String form:**

```
mem:///unum=<unum>
      [/offset=<hex-offset> | /physaddr=<hex-physaddr>]
```

**SCHEME mod VERSION 0**

This scheme is used to identify Oracle Solaris kernel modules.

| Member Name | Data Type | Stability              |
|-------------|-----------|------------------------|
| scheme      | string    | Committed, value "fmd" |
| version     | uint8     | Committed, value 0     |
| mod-name    | string    | Committed              |
| [mod-desc]  | string    | Committed              |
| [mod-pkg]   | fmri      | (fmri)                 |

**mod-name**

The module name, as in per [modinfo\(1M\)](#).

**[mod-desc]**

The module description, again as per [modinfo\(1M\)](#).

**[mod-pkg]**

A pkg scheme FMRI of the package which delivers this module.

**String form:**

```
mod:///mod-name=<mod-name>
```

**SCHEME pkg VERSION 0**

This version of the pkg scheme corresponds to legacy SVR4 packaging as used with [pkgadd\(1M\)](#), [pkginfo\(1\)](#), [pkgrm\(1M\)](#), and related utilities. This scheme is universal if package name and package version conventions are adhered to.

| Member Name | Data Type | Stability              |
|-------------|-----------|------------------------|
| -----       |           |                        |
| scheme      | string    | Committed, value "pkg" |
| version     | int8      | Committed, value 0     |
| pkg-name    | string    | Committed              |
| pkg-basedir | string    | Committed              |
| pkg-inst    | string    | Committed              |
| pkg-version | string    | Committed              |

#### pkg-name

The package name.

#### pkg-basedir

BASEDIR as per [pkginfo\(1\)](#)

#### pkg-inst

PKGINST as per [pkginfo\(1\)](#)

#### pkg-version

VERSION as per [pkginfo\(1\)](#)

String form:

```
pkg:///<pkg-name>/:version=<pkg-version>
```

### SCHEME pkg VERSION 1

This version of the pkg scheme corresponds to the [pkg\(5\)](#) Image Packaging System. Such FMRI's are used in string form in the [pkg\(1\)](#) command line and its output, and the command line permits abbreviations of the FMRI string.

FMRI's in the pkg scheme version 1 are universal: the same FMRI interpreted in two distinct contexts (such as in distinct Oracle Solaris instances) identify the same actual package (or copies thereof).

The [pkg\(5\)](#) command line uses the string form of version 1 pkg FMRI's, and permits abbreviations of such strings.

| Member Name   | Data Type | Stability              |
|---------------|-----------|------------------------|
| -----         |           |                        |
| scheme        | string    | Committed, value "pkg" |
| version       | uint8     | Committed, value 1     |
| [authority]   | nvlist    | Committed              |
| pkg-name      | string    | Committed              |
| [pkg-version] | nvlist    | Committed              |

#### [authority]

The authority nvlist specifies the package publisher, per pkg publisher output:

| Member Name | Data Type | Stability |
|-------------|-----------|-----------|
| publisher   | string    | Committed |

publisher

The publisher of this package, such as Oracle Solaris.

pkg-name

The name (stem) of the package

[pkg-version]

| Member Name | Data Type | Stability |
|-------------|-----------|-----------|
| release     | string    | Committed |
| [built-on]  | string    | Committed |
| [branch]    | string    | Committed |
| [timestamp] | string    | Committed |

release

The primary version of the component, written as a dot sequence string (an arbitrary-length, dot-separated series of non-zero-leading decimal digits).

built-on

The operating system version on which the component was built, again specified as a dot sequence.

branch

The vendor branch version, as a dot sequence.

timestamp

An ISO 8601 date string: *yyyymmddThhmmssZ* in which the T and Z are literals.

String form:

pkgauthstr = <authority.publisher>

pkg://[pkgauthstr]/<pkg-name>

[@<release>[,<built-on>][-<branch>][:<timestamp>]

## SCHEME svc VERSION 0

The svc scheme describes SMF (see [smf\(5\)](#)) services and service instances. In SMF command lines such as [svcadm\(1M\)](#), [svccfg\(1M\)](#), [svccprop\(1\)](#), [svcs\(1\)](#) and their output such FMRI are always written in their string form, and the SMF command lines permit abbreviations.

The svc scheme FMRI is a logical scheme, and, as such an FMRI should only be interpreted in the fault management domain (Oracle Solaris instance) in which it was generated.

| Member Name | Data Type | Stability              |
|-------------|-----------|------------------------|
| scheme      | string    | Committed, value "svc" |

|                |        |                    |
|----------------|--------|--------------------|
| version        | uint8  | Committed, value 0 |
| svc-name       | string | Committed          |
| [svc-instance] | string | Committed          |
| [contract-id]  | string | Committed          |
| [svc-scope]    | string | Committed          |

pkg://[pkgauthstr]/<pkg-name>

svc-name

The service name, such as network/smtp. See [smf\(5\)](#).

[svc-instance]

Used when identifying an instance of a service. Often default, but can also be things like the sendmail instance of service network/smtp.

[contract-id]

This identifies the numeric (although written as a decimal string) contract id (see [smf\(5\)](#) and [contract\(4\)](#)) used to manage a particular running instance of a service.

[svc-scope]

As per [smf\(5\)](#), all service and instance objects are contained in a scope that represents a collection of configuration information. The configuration of the local Oracle Solaris instance is called the localhost scope, and is the only currently supported scope.

String form:

```
svc://[<svc-scope>]/<svc-name>
      [:<svc-instance>[@<contract-id>]]
```

SMF subsystem commands such as [svcadm\(1M\)](#) permit (or even require) abbreviations of the string form above. The indication of svc scheme at the start of the string is considered optional and implied by the SMF command line utility. When the svc scheme identifier is used, the authority component is also abbreviated to `svc:/<svc-name>` instead of the more-formal `svc:///<svc-name>`. For example, in SMF we write

```
svc:/network/smtp:sendmail
```

instead of the more-formal

```
svc:///network/smtp:sendmail
```

or even the more formal

```
svc://localhost/network/smtp:sendmail
```

Furthermore, SMF permits further abbreviation if it identifies a unique service or instance. For example the instance name `sendmail` is unique and one can use `svcs sendmail`, `svcadm restart sendmail`, and so forth.

Such abbreviations are a convention of the SMF subsystem and not part of the formal FMRI definition. When a svc scheme FMRI is received in `nvlist` form (for example, in a [libfmevent\(3LIB\)](#) subscription it appears in the more formal form when rendered as a

string.

## SCHEME `sw` VERSION 0

This FMRI scheme is used to identify software resources such as executables, library objects, and core files. This is a logical FMRI scheme.

| Member Name | Data Type | Stability             |
|-------------|-----------|-----------------------|
| scheme      | string    | Committed, value "sw" |
| version     | uint8     | Committed, value 0    |
| object      | nvlist    | See belowPrivate      |
| [site]      | nvlist    | See belowPrivate      |
| [context]   | nvlist    | See belowPrivate      |

The `object`, `site`, and `context` members are all `Private` and so subject to incompatible change, but their content is documented below for informational purposes.

### object

This required member identifies the software object. At this time only objects resident within the filesystem are catered for.

| Member | Type   | Description                                     |
|--------|--------|-------------------------------------------------|
| path   | string | Filesystem path to object                       |
| [root] | string | If present, real path to chroot root directory  |
| [pkg]  | fmri   | "pkg" scheme fmri of package that delivers path |

### site

This optional member identifies a site within the object.

| Member   | Type   | Description                                    |
|----------|--------|------------------------------------------------|
| [token]  | string | Vendor and subsystem unique publisher token id |
| [module] | string | Source module information                      |
| [file]   | string | Source filename (translation unit)             |
| [func]   | string | Source function                                |
| [line]   | int64  | Source file line number                        |

### context

This optional member communicates runtime information.

| Member      | Type   | Description            |
|-------------|--------|------------------------|
| [origin]    | string | "userland" or "kernel" |
| [execname]  | string | Executable name        |
| [pid]       | uint64 | Process id             |
| [thread-id] | uint64 | Thread id              |

|                    |              |                                   |
|--------------------|--------------|-----------------------------------|
| [os-instance-uuid] | string       | Solaris instance UUID             |
| [zone]             | string       | Zone name, if not the global zone |
| [ctid]             | uint64       | Contract id                       |
| [stack]            | string array | Symbolic stack trace              |

String form:

```
sw://[<authority>]/
    [:root=<object.root>]
    :path=<object.path>
    [#<fragment-identifier>]
```

Slash characters / in the root and object path are not escaped.

<fragment-identifier> is one of:

```
:token=<site.token>
```

or

```
:file=<site.file>[:func=<site.func>][:line=<site.line>]
```

### SCHEME zfs VERSION 0

| Member Name | Data Type | Stability              |
|-------------|-----------|------------------------|
| -----       | -----     | -----                  |
| scheme      | string    | Committed, value "zfs" |
| version     | uint8     | Committed, value 0     |
| [pool-name] | string    | Committed              |
| pool        | uint64    | Committed              |
| [vdev]      | uint64    | Private                |
| [vdev-name] | string    | Private                |

[pool-name]

The pool name, as per `zpool list -o name`

pool

The pool GUID as per `zpool list -o guid`

[vdev]

The semantics of this Private FMRI member are not documented

[vdev-name]

The semantics of this Private FMRI member are not documented.

String form:

```
hex-vdev = 'vdev' in hexadecimal with no '0x' prefix.
hex-pool = 'pool' in hexadecimal with no '0x' prefix.
zfs://[pool_name=<pool-name>/]pool=<hex-pool>
zfs://[pool_name=<zfs://[pool_name=<hex-pool>
[/vdev=<hex-vdev>][:vdev_name=<vdev-name>]
```

The <vdev-name> is written without escaping any / characters.

**See Also** [hostid\(1\)](#), [fmadm\(1M\)](#), [fmd\(1M\)](#), [fmdump\(1M\)](#), [pkgadd\(1M\)](#), [pkginfo\(1\)](#), [pkgrm\(1M\)](#), [psradm\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [svcprop\(1\)](#), [svcs\(1\)](#), [libfmevent\(3LIB\)](#), [libnvpair\(3LIB\)](#), [contract\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

[pkg\(5\)](#)

*RFC 2396*

**Name** fnmatch – file name pattern matching

**Description** 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 [regex\(5\)](#) manual page.

Patterns Matching a Single Character The following *patterns matching a single character* match a single character: *ordinary characters*, *special pattern characters* and *pattern bracket expressions*. The pattern bracket expression will also match a single collating element.

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.

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

- ? A question-mark is a pattern that will match any character.
- \* An asterisk is a pattern that will match multiple characters, as described in [Patterns Matching Multiple Characters](#), below.
- [ The open bracket will introduce a pattern bracket expression.

The description of basic regular expression bracket expressions on the [regex\(5\)](#) manual page also applies to the pattern bracket expression, except that the exclamation-mark character ( ! ) replaces the circumflex character ( ^ ) in its role in a *non-matching list* in the regular expression notation. A bracket expression starting with an unquoted circumflex character produces unspecified results.

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 `[^\!]` to match either character.

When pattern matching is used where shell quote removal is not performed (such as in the argument to the `find -name` primary when `find` is being called using one of the `exec` functions, or in the *pattern* argument to the [fnmatch\(3C\)](#) 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 `\\` 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.

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:

`a[bc]` matches the strings `ab` and `ac`.

`a*d` matches the strings `ad`, `abd` and `abcd`, but not the string `abc`.

`a*d*` matches the strings `ad`, `abcd`, `abcdef`, `aaaad` and `adddd`.

`*a*d` matches the strings `ad`, `abcd`, `efabcd`, `aaaad` and `adddd`.

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 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** 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 `printf(3C)` 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 `scanf(3C)` function to read the input file.

**Format** The description of an individual record is as follows:

```
"<format>", [<arg1>, <arg2>, . . . , <argn>]
```

The format is a character string that contains three types of objects defined below:

*characters* Characters that are not *escape sequences* or *conversion specifications*, as described below, are copied to the output.

*escape sequences* Represent non-graphic characters.

*conversion specifications* Specifies the output format of each argument. (See below.)

The following characters have the following special meaning in the format string:

» (An empty character position.) One or more blank characters.

/\ Exactly one space character.

The notation for spaces allows some flexibility for application output. Note that an empty character position in format represents one or more blank characters on the output (not *white space*, which can include newline characters). Therefore, another utility that reads that output as its input must be prepared to parse the data using `scanf(3C)`, `awk(1)`, and so forth. The character is used when exactly one space character is output.

**Escape Sequences** The following table lists escape sequences and associated actions on display devices capable of the action.

| 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. |
| \f       | form-feed | Moves the printing position to the initial printing position of the next logical page.                                     |

| Sequence | Character       | Terminal Action                                                                                                                                                      |
|----------|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \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:

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>flags</i>                 | Zero or more <i>flags</i> , in any order, that modify the meaning of the conversion specification.                                                                                                                                                                                                                                                                                                                                                                                      |
| <i>field width</i>           | An optional string of decimal digits to specify a minimum <i>field width</i> . 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).                                                                                                                                                                                              |
| <i>precision</i>             | 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. |
| <i>conversion characters</i> | 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 -).                                                                                                                                          |
| <space> | 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 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.

The *conversion characters* and their meanings are:

*d,i,o,u,x,X* 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 `[-]dddd`. The x conversion uses the numbers and letters 0123456789abcdef and the X conversion uses the numbers and letters 0123456789ABCDEF. The *precision* 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.

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:

```
"%d\n", <foo>
```

the implementation could use a `printf()` call such as:

```
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\(3C\)](#)

**Name** fsattr – extended file attributes

**Description** 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.

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.

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. Implementations 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.

The following list details the operations that are rejected in the current implementation:

|                      |                                                                                                                                                                                                                  |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>link</code>    | 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. |
| <code>rename</code>  | 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 <code>link</code> restriction above.              |
| <code>mkdir</code>   |                                                                                                                                                                                                                  |
| <code>symlink</code> |                                                                                                                                                                                                                  |
| <code>mknode</code>  | 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):

- `cp` By default, `cp` ignores attributes and copies only file data. This is intended to maintain the semantics implied by `cp` currently, where attributes (such as owner and mode) are not copied unless the `-p` option is specified. With the `-@` (or `-p`) option, `cp` attempts to copy all attributes along with the file data.
- `cpio` The `-@` option informs `cpio` to archive attributes, but by default `cpio` ignores extended attributes. See [Extended Archive Formats](#) below for a description of the new archive records.
- `du` File sizes computed include the space allocated for any extended attributes present.
- `find` By default, `find` ignores attributes. The `-xattr` expression provides support for searches involving attribute space. It returns true if extended attributes are present on the current file.
- `fsck` The `fsck` 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 `fsck` will strip them from the files and place them in `lost+found`. 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 `fsck` utility should be run to stabilize the file system before using it in an attribute-aware environment.
- `fsdb` This `fsdb` utility is able to find the inode for the “hidden” extended attribute directory.
- `ls` The `ls -@` 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 [pathconf\(2\)](#) variable `XATTR_EXISTS` is set to true. See the `pathconf()` section below. The `-@` option uses the same general output format as the `-l` option.

- mv** 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 `cp -p` 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.
- pax** The `-@` option informs `pax` to archive attributes, but by default `pax` ignores extended attributes. The `pax(1)` utility is a generic replacement for both `tar(1)` and `cpio(1)` and is able to produce either output format in its archive. See Extended Archive Formats below for a description of the new archive records.
- tar** In the default case, `tar` does not attempt to place attributes in the archive. If the `-@` option is specified, however, `tar` 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 `tar` archive files (the same is true for `pax` and `cpio` archives) similar to the way ACLs records were defined. See Extended Archive Formats 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, including [faccessat\(2\)](#), [fchownat\(2\)](#), [fstatat\(2\)](#), [futimesat\(2\)](#), [renameat\(2\)](#), and [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 AT\_SYMLINK\_NOFOLLOW is set in the *flag* argument, the function will not automatically traverse a symbolic link at the position of the path. If AT\_TRIGGER is set in the *flag* argument and the vnode is a trigger mount point, the mount is performed and the function returns the attributes of the root of the mounted filesystem. The `fstatat()` function is a multipurpose 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)`.

### Determine accessibility of a file

```
int faccessat(int fd, const char *path, int amode, int flag);
```

The `faccessat()` function checks the file named by the pathname pointed to by the *path* argument for accessibility according to the bit pattern contained in *amode*, using the real user ID in place of the effective user ID and the real group ID in place of the effective group ID. This allows a `setuid` process to verify that the user running it would have had permission to access this file.

If *path* specifies a relative path, the file whose accessibility is to be determined is located relative to the directory associated with the file descriptor *fd* instead of the current working directory. If *path* specifies an absolute path, the *fd* argument is ignored.

If `faccessat()` is passed in the `fd` parameter the special value `AT_FDCWD`, defined in `<fcntl.h>`, the current working directory is used and the behavior is identical to a call to `access(2)`.

### 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               |
| 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                      |
|------------|--------------------|----------------------------------|
| c_magic    | 6                  | 70707                            |
| c_dev      | 6                  | First half of unique file ID     |
| 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.

| 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\(1M\)](#), [access\(2\)](#), [chown\(2\)](#), [link\(2\)](#), [open\(2\)](#), [pathconf\(2\)](#), [rename\(2\)](#), [stat\(2\)](#), [unlink\(2\)](#), [utimes\(2\)](#), [attpopen\(3C\)](#), [standards\(5\)](#)

**Name** grub – GRand Unified Bootloader 2 software on Oracle Solaris

**Description** The current release of the Oracle Solaris operating system is shipped with the GRUB (GRand Unified Bootloader) 2 software. GRUB is developed and supported by the Free Software Foundation.

The overview for the GRUB Manual, accessible at [www.gnu.org](http://www.gnu.org), describes GRUB:

Briefly, a boot loader is the first software program that runs when a computer starts. It is responsible for loading and transferring control to an operating system kernel software (such as Linux or GNU Mach). The kernel, in turn, initializes the rest of the operating system (for example, a GNU [Ed. note: or Solaris] system).

GNU GRUB is a very powerful boot loader that can load a wide variety of free as well as proprietary operating systems with chain-loading. GRUB is designed to address the complexity of booting a personal computer; both the program and this manual are tightly bound to that computer platform, although porting to other platforms may be addressed in the future. Note: Oracle has ported GRUB to the Solaris operating system.

One of the important features in GRUB is flexibility; GRUB understands filesystems and kernel executable formats, so you can load an arbitrary operating system the way you like, without recording the physical position of your kernel on the disk. Thus you can load the kernel just by specifying its file name and the drive and partition where the kernel resides.

Among Solaris machines, GRUB is supported on x86 platforms. The GRUB software that is shipped with Solaris adds one utility not present in the open-source distribution:

**bootadm(1M)** Enables you to manage the boot archive and make changes to the GRUB menu.

Beyond this Solaris-specific utility, the GRUB software is described in the GRUB manual, a PDF version of which is available from the Oracle web site.

Most administrators will not need to manually execute the GRUB utilities (which can be found in `/usr/lib/grub2/bios` and `/usr/lib/grub2/uefi64` for systems with BIOS firmware and 64-bit UEFI firmware, respectively), as **bootadm(1M)** provides an interface to modifying the GRUB menu.

Administrators should not manually edit the GRUB configuration file, `grub.cfg`, which can be found in the `/boot/grub` subdirectory of the top-level ZFS dataset for the system's root pool (that is, `/rpool/boot/grub`). This file is automatically regenerated when an administrator changes the boot configuration with **bootadm(1M)**. Advanced administrators who want to directly create GRUB menu entries must edit the supplemental GRUB configuration file named `custom.cfg`, which is stored in the same directory as the `grub.cfg` file. Note that on a freshly-installed system, this file will not be present, so the administrator will need to create it.

**See Also** `boot(1M)`, `bootadm(1M)`, `installgrub(1M)`

*Installing Oracle Solaris 11.1 Systems*

*Oracle Solaris Administration: Common Tasks*

<http://www.gnu.org/software/grub>

**Notes** Previous versions of Solaris have documented the use of the `installgrub(1M)` command for installing the GRUB boot loader. This command is deprecated and is present ONLY for convenience, disaster recovery, and downgrading back to a Solaris boot environment in which GRUB Legacy is the system boot loader. Please consult the “Booting and Shutting Down Oracle Solaris” documentation for further details.

**Name** gss\_auth\_rules – overview of GSS authorization

**Description** The establishment of the veracity of a user's credentials requires both authentication (Is this an authentic user?) and authorization (Is this authentic user, in fact, authorized?).

When a user makes use of Generic Security Services (GSS) versions of the ftp or ssh clients to connect to a server, the user is not necessarily authorized, even if his claimed GSS identity is authenticated. Authentication merely establishes that the user is who he says he is to the GSS mechanism's authentication system. Authorization is then required: it determines whether the GSS identity is permitted to access the specified Solaris user account.

The GSS authorization rules are as follows:

- If the mechanism of the connection has a set of authorization rules, then use those rules. For example, if the mechanism is Kerberos, then use the [krb5\\_auth\\_rules\(5\)](#), so that authorization is consistent between raw Kerberos applications and GSS/Kerberos applications.
- If the mechanism of the connection does not have a set of authorization rules, then authorization is successful if the remote user's gssname matches the local user's gssname exactly, as compared by [gss\\_compare\\_name\(3GSS\)](#).

**Files** /etc/passwd System account file. This information may also be in a directory service. See [passwd\(4\)](#).

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Interface Stability | Committed       |

**See Also** [ftp\(1\)](#), [ssh\(1\)](#), [gsscred\(1M\)](#), [gss\\_compare\\_name\(3GSS\)](#), [passwd\(4\)](#), [attributes\(5\)](#), [krb5\\_auth\\_rules\(5\)](#)

**Name** hal – overview of hardware abstraction layer

**Description** The Hardware Abstraction Layer (HAL) provides a view of the various hardware attached to a system. This view is updated dynamically as hardware configuration changes by means of hotplug or other mechanisms. HAL represents a piece of hardware as a device object. A device object is identified by a unique identifier and carries a set of key/value pairs, referred to as device properties. Some properties are derived from the actual hardware, some are merged from device information files (. fdi files), and some are related to the actual device configuration.

HAL provides an easy-to-use API through D-Bus. D-Bus is an IPC framework that, among other features, provides a system-wide message-bus that allows applications to talk to one another. Specifically, D-Bus provides asynchronous notification such that HAL can notify other peers on the message-bus when devices are added and removed, as well as when properties on a device are changing.

In the Solaris operating system, HAL is supported by a daemon, [hald\(1M\)](#), and a set of utilities that enable the adding and removing of devices and the modification of their properties.

**See Also** [hald\(1M\)](#), [hal-device\(1M\)](#), [hal-fdi-validate\(1M\)](#), [hal-find\(1M\)](#), [hal-find-by-capability\(1M\)](#), [hal-find-by-property\(1M\)](#), [hal-get-property\(1M\)](#), [hal-set-property\(1M\)](#), [fdi\(4\)](#)

<http://freedesktop.org>.

**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.

| 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    |

| Conversions Performed |        |         |        |
|-----------------------|--------|---------|--------|
| MS 1250               | MS 852 | MS 1250 | MS 852 |
| 255                   | 360    | 344     | 204    |
| 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.

| 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     |

| Conversions Performed |         |         |         |
|-----------------------|---------|---------|---------|
| MS 1250               | Mazovia | MS 1250 | Mazovia |
| 267                   | 372     | 357-360 | 40      |
| 270                   | 40      | 361     | 244     |
| 271                   | 206     | 362     | 40      |
| 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 |

| Conversions Performed |     |         |     |
|-----------------------|-----|---------|-----|
| MS 1250               | DHN | MS 1250 | DHN |
| 244                   | 40  | 335-336 | 40  |
| 245                   | 200 | 337     | 341 |
| 246-252               | 40  | 340     | 40  |
| 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        |

| 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        |

| Conversions Performed |            |         |            |
|-----------------------|------------|---------|------------|
| MS 1251               | ISO 8859-5 | MS 1251 | ISO 8859-5 |
| 272                   | 364        | 361     | 341        |
| 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    |

| Conversions Performed |        |         |        |
|-----------------------|--------|---------|--------|
| MS 1251               | KO18-R | MS 1251 | KO18-R |
| 213                   | 40     | 321     | 363    |
| 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    |

| Conversions Performed |        |         |        |
|-----------------------|--------|---------|--------|
| MS 1251               | KOI8-R | MS 1251 | KOI8-R |
| 262                   | 266    | 354     | 315    |
| 263                   | 246    | 355     | 316    |
| 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         |

| Conversions Performed |             |         |             |
|-----------------------|-------------|---------|-------------|
| MS 1251               | PC Cyrillic | MS 1251 | PC Cyrillic |
| 230                   | 260         | 336     | 236         |
| 231-247               | 40          | 337     | 237         |
| 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         |

| Conversions Performed |             |         |             |
|-----------------------|-------------|---------|-------------|
| MS 1251               | PC Cyrillic | MS 1251 | PC Cyrillic |
| 323                   | 223         | 371     | 351         |
| 324                   | 224         | 372     | 352         |
| 325                   | 225         | 373     | 353         |
| 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          |

| Conversions Performed |              |         |              |
|-----------------------|--------------|---------|--------------|
| MS 1251               | Mac Cyrillic | MS 1251 | Mac Cyrillic |
| 221                   | 324          | 302     | 202          |
| 222                   | 325          | 303     | 203          |
| 223                   | 322          | 304     | 204          |
| 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          |

| Conversions Performed |              |         |              |
|-----------------------|--------------|---------|--------------|
| MS 1251               | Mac Cyrillic | MS 1251 | Mac Cyrillic |
| 255                   | 55           | 335     | 235          |
| 256                   | 250          | 336     | 236          |
| 257                   | 272          | 337     | 237          |
| 355                   | 316          |         |              |

**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 – 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   |                  |

| Index of Conversion Code Tables |            |                  |
|---------------------------------|------------|------------------|
| 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        |                  |
| 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 |                  |

| Index of Conversion Code Tables |              |                   |
|---------------------------------|--------------|-------------------|
| 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       |                   |
| 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\(5\)](#)

**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        |           |            |

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        |

| 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        |

| Conversions Performed |            |        |            |
|-----------------------|------------|--------|------------|
| MS 852                | ISO 8859-2 | MS 852 | ISO 8859-2 |
| 250                   | 312        | 364    | 242        |
| 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     |

| Conversions Performed |         |         |         |
|-----------------------|---------|---------|---------|
| MS 852                | MS 1250 | MS 852  | MS 1250 |
| 214                   | 356     | 323     | 313     |
| 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     |

| Conversions Performed |         |        |         |
|-----------------------|---------|--------|---------|
| MS 852                | MS 1250 | MS 852 | MS 1250 |
| 247                   | 236     | 361    | 275     |
| 250                   | 312     | 362    | 262     |
| 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      |

| Conversions Performed |         |         |         |
|-----------------------|---------|---------|---------|
| MS 852                | Mazovia | MS 852  | Mazovia |
| 235                   | 234     | 343     | 245     |
| 236-243               | 40      | 344     | 244     |
| 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:

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

Converted to Underscore '\_' (137)

```

-----
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

```

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.

| 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)

-----  
 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       |

| Conversions Performed |           |            |           |
|-----------------------|-----------|------------|-----------|
| ISO 8859-1            | ISO 646it | ISO 8859-1 | ISO 646it |
| 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  
 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           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\_8859-2 – 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.

| 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    |

| Conversions Performed |        |            |        |
|-----------------------|--------|------------|--------|
| ISO 8859-2            | MS 852 | ISO 8859-2 | MS 852 |
| 271                   | 347    | 351        | 202    |
| 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     |

| Conversions Performed |         |            |         |
|-----------------------|---------|------------|---------|
| ISO 8859-2            | Mazovia | ISO 8859-2 | Mazovia |
| 177-237               | 40      | 324-325    | 40      |
| 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     |

| Conversions Performed |         |            |         |
|-----------------------|---------|------------|---------|
| ISO 8859-2            | Mazovia | ISO 8859-2 | Mazovia |
| 312                   | 220     | 370-373    | 40      |
| 313-320               | 40      | 374        | 201     |
| 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  |

| Conversions Performed |     |            |     |
|-----------------------|-----|------------|-----|
| ISO 8859-2            | DHN | ISO 8859-2 | DHN |
| 274                   | 220 | 363        | 216 |
| 275-276               | 40  | 364        | 223 |
| 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    |

| 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    |

| Conversions Performed |        |            |        |
|-----------------------|--------|------------|--------|
| ISO 8859-5            | KOI8-R | ISO 8859-5 | KOI8-R |
| 312                   | 377    | 372        | 252    |
| 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         |

| Conversions Performed |             |            |             |
|-----------------------|-------------|------------|-------------|
| ISO 8859-5            | PC Cyrillic | ISO 8859-5 | PC Cyrillic |
| 273                   | 213         | 330        | 250         |
| 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     |

| Conversions Performed |         |            |         |
|-----------------------|---------|------------|---------|
| ISO 8859-5            | MS 1251 | ISO 8859-5 | MS 1251 |
| 252                   | 214     | 332        | 352     |
| 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     |

| Conversions Performed |         |            |         |
|-----------------------|---------|------------|---------|
| ISO 8859-5            | MS 1251 | ISO 8859-5 | MS 1251 |
| 305                   | 325     | 365        | 276     |
| 306                   | 326     | 366        | 263     |
| 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          |

| Conversions Performed |              |            |              |
|-----------------------|--------------|------------|--------------|
| ISO 8859-5            | Mac Cyrillic | ISO 8859-5 | Mac Cyrillic |
| 252                   | 276          | 333        | 353          |
| 253                   | 40           | 334        | 354          |
| 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          |

| Conversions Performed |              |            |              |
|-----------------------|--------------|------------|--------------|
| ISO 8859-5            | Mac Cyrillic | ISO 8859-5 | Mac Cyrillic |
| 305                   | 225          | 366        | 264          |
| 306                   | 226          | 367        | 273          |
| 307                   | 227          | 370        | 300          |
| 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         |

| 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     |

| 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        |

| 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        |

| Conversions Performed |            |        |            |
|-----------------------|------------|--------|------------|
| KOI8-R                | ISO 8859-5 | KOI8-R | ISO 8859-5 |
| 312                   | 331        | 372    | 267        |
| 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         |

| Conversions Performed |             |        |             |
|-----------------------|-------------|--------|-------------|
| KOI8-R                | PC Cyrillic | KOI8-R | PC Cyrillic |
| 307                   | 243         | 354    | 213         |
| 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     |

| Conversions Performed |         |        |         |
|-----------------------|---------|--------|---------|
| KO18-R                | MS 1251 | KO18-R | MS 1251 |
| 242                   | 203     | 322    | 360     |
| 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     |

| Conversions Performed |         |        |         |
|-----------------------|---------|--------|---------|
| KOI8-R                | MS 1251 | KOI8-R | MS 1251 |
| 275                   | 210     | 355    | 314     |
| 276                   | 241     | 356    | 315     |
| 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          |

| Conversions Performed |              |        |              |
|-----------------------|--------------|--------|--------------|
| KO18-R                | Mac Cyrillic | KO18-R | Mac Cyrillic |
| 242                   | 257          | 323    | 361          |
| 243                   | 336          | 324    | 362          |
| 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          |

| Conversions Performed |              |        |              |
|-----------------------|--------------|--------|--------------|
| KO18-R                | Mac Cyrillic | KO18-R | Mac Cyrillic |
| 275                   | 40           | 356    | 215          |
| 276                   | 330          | 357    | 216          |
| 277                   | 332          | 360    | 217          |
| 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        |

| 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        |

| Conversions Performed |            |              |            |
|-----------------------|------------|--------------|------------|
| Mac Cyrillic          | ISO 8859-5 | Mac Cyrillic | ISO 8859-5 |
| 260-263               | 40         | 365          | 345        |
| 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     |

| Conversions Performed |        |              |        |
|-----------------------|--------|--------------|--------|
| Mac Cyrillic          | KO18-R | Mac Cyrillic | KO18-R |
| 213                   | 354    | 330          | 276    |
| 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    |

| Conversions Performed |        |              |        |
|-----------------------|--------|--------------|--------|
| Mac Cyrillic          | KOI8-R | Mac Cyrillic | KOI8-R |
| 256                   | 262    | 363          | 325    |
| 257                   | 242    | 364          | 306    |
| 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         |

| Conversions Performed |             |              |             |
|-----------------------|-------------|--------------|-------------|
| Mac Cyrillic          | PC Cyrillic | Mac Cyrillic | PC Cyrillic |
| 345                   | 245         | 367          | 347         |
| 346                   | 246         | 370          | 350         |
| 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     |

| Conversions Performed |         |              |         |
|-----------------------|---------|--------------|---------|
| Mac Cyrillic          | MS 1251 | Mac Cyrillic | MS 1251 |
| 215                   | 315     | 277          | 234     |
| 216                   | 316     | 300          | 274     |
| 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     |

| Conversions Performed |         |              |         |
|-----------------------|---------|--------------|---------|
| Mac Cyrillic          | MS 1251 | Mac Cyrillic | MS 1251 |
| 250                   | 256     | 335          | 250     |
| 252                   | 231     | 336          | 270     |
| 253                   | 200     | 337          | 377     |
| 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        |

| 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    |

| 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        |

| 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           | KO18-R | PC Cyrillic | KO18-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.

| 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     |

| Conversions Performed |         |             |         |
|-----------------------|---------|-------------|---------|
| PC Cyrillic           | MS 1251 | PC Cyrillic | MS 1251 |
| 232                   | 332     | 354         | 374     |
| 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          |

| Conversions Performed |              |             |              |
|-----------------------|--------------|-------------|--------------|
| PC Cyrillic           | Mac Cyrillic | PC Cyrillic | Mac Cyrillic |
| 256                   | 356          | 360         | 335          |
| 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<br>Code             | FROM<br>Filename<br>Element | TO Code Set<br>Target Code       | TO<br>Filename<br>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<br>(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<br>(KS C 5601-1987)  | ko_KR-johap                 | Korean UTF-8                     | ko_KR-UTF-8               |
| Korean Johap<br>(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<br>(KS C 5601-1987) | ko_KR-johap               |
| Korean UTF-8                      | ko_KR-UTF-8                 | Korean Johap<br>(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<br>Code | FROM<br>Filename | TO Code Set<br>Target Code | TO<br>Filename |
|-----------------------|------------------|----------------------------|----------------|
|-----------------------|------------------|----------------------------|----------------|

| Element |        | 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     |
| 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<br>(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<br>(GB 2312-1980) | zh_CN.euc |

## CODE SET CONVERSIONS SUPPORTED

| FROM Code Set<br>Code | FROM<br>Filename<br>Element | TO Code Set<br>Target Code             | TO<br>Filename<br>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<br>(CNS 11643-1992) | zh_TW-euc                 |

|                                        |                 |             |                 |
|----------------------------------------|-----------------|-------------|-----------------|
| UTF-8                                  | UTF-8           | ISO 2022-TW | zh_TW-iso2022-7 |
| Chinese/PRC EUC<br>(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<br>(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.

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.

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).                                  |          |         |
| 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. |          |         |

**Name** ieee802.11 – 802.11 kernel statistics

**Description** This page describes the kernel statistics that can be used to monitor attributes specific to the 802.11 physical layer. These statistics can be retrieved using [kstat\(1M\)](#). Not all 802.11 devices will support all statistics.

|                           |                                                                                                                                                     |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>tx_frags</code>     | Count of data and management fragments transmitted.                                                                                                 |
| <code>rx_frags</code>     | Count of data and management fragments received.                                                                                                    |
| <code>rx_dups</code>      | Count of duplicate frames received. Duplicates are determined by the sequence control field.                                                        |
| <code>mcast_tx</code>     | Count of broadcast and multicast frames transmitted.                                                                                                |
| <code>mcast_rx</code>     | Count of broadcast and multicast frames received.                                                                                                   |
| <code>tx_failed</code>    | Count of frames that could not be transmitted due to the retransmission limit being reached.                                                        |
| <code>tx_retrans</code>   | Count of frames successfully retransmitted after one or more retransmissions.                                                                       |
| <code>tx_reretrans</code> | Count of frames successfully retransmitted after more than one retransmission.                                                                      |
| <code>rts_success</code>  | Count of times a CTS was received in response to an RTS.                                                                                            |
| <code>rts_failure</code>  | Count of times a CTS was not received in response to an RTS.                                                                                        |
| <code>ack_failure</code>  | Count of times an ACK was expected but was not received.                                                                                            |
| <code>fcs_errors</code>   | Count of frames received with FCS errors.                                                                                                           |
| <code>wep_errors</code>   | Count of frames received with the WEP bit set but that either should not have been encrypted or that were discarded due to WEP not being supported. |

**See Also** [kstat\(1M\)](#)

**Name** ieee802.3, cap\_autoneg, cap\_1000fdx, cap\_1000hdx, cap\_100fdx, cap\_100hdx, cap\_10fdx, cap\_10hdx, cap\_rem\_fault, cap\_pause, cap\_asym\_pause, adv\_cap\_autoneg, adv\_cap\_1000fdx, adv\_cap\_1000hdx, adv\_cap\_100fdx, adv\_cap\_100hdx, adv\_cap\_10fdx, adv\_cap\_10hdx, adv\_cap\_pause, adv\_cap\_asym\_pause, adv\_rem\_fault, lp\_cap\_autoneg, lp\_cap\_1000fdx, lp\_cap\_1000hdx, lp\_cap\_100fdx, lp\_cap\_100hdx, lp\_cap\_10fdx, lp\_cap\_10hdx, lp\_cap\_pause, lp\_cap\_asym\_pause, lp\_rem\_fault, xcvr\_addr, xcvr\_id, xcvr\_inuse, link\_up, link\_duplex, link\_tx\_pause, link\_rx\_pause – Ethernet mii kstat and dladm parameters

**Description** This page describes the kernel statistics and the `dladm(1M)` configuration parameters used to monitor and configure the Ethernet physical layer.

The `cap_*` parameters exist in the kernel statistics for an Ethernet device. The parameters describe the maximum capability of a device. When the value of a statistic is 1, the device has the capability described. When the value is 0, the device does not have the capability.

The exceptions to this rule are the `cap_asym_pause` and `cap_pause` parameters which are explained later in this page.

|                            |                                                              |
|----------------------------|--------------------------------------------------------------|
| <code>cap_autoneg</code>   | Capable of auto-negotiation                                  |
| <code>cap_1000fdx</code>   | Capable of 1000 full duplex operation                        |
| <code>cap_1000hdx</code>   | Capable of 1000 half duplex operation                        |
| <code>cap_100fdx</code>    | Capable of 100 full duplex operation                         |
| <code>cap_100hdx</code>    | Capable of 100 half duplex operation                         |
| <code>cap_10fdx</code>     | Capable of 10 full duplex operation                          |
| <code>cap_10hdx</code>     | Capable of 10 half duplex operation                          |
| <code>cap_rem_fault</code> | Capable of reporting locally detected faults to link partner |

The `adv_cap_*` parameters exist in the kernel statistics and represent a mirror image of the `dladm adv_*_cap` parameter list for an Ethernet device. The `dladm adv_*_cap` tuning parameters allow fine grain control of the Ethernet device physical layer. The parameters are also a subset of the `cap_*` statistics. If the `cap_*` value is 0, the corresponding `adv_cap_*` must also be 0. The exceptions to this rule are the `adv_cap_asym_pause` and `adv_cap_pause` parameters.

When auto-negotiation is enabled, the `adv_*_cap` statistics show which capabilities are advertised to the link partner. When auto-negotiation is disabled in *forced mode*, the statistics precisely show how a link should function and that it must be matched on the link partner to achieve a valid link up.

Statistics with values other than 0 and 1 are also described in the following.

|                              |                                                      |
|------------------------------|------------------------------------------------------|
| <code>adv_cap_autoneg</code> | Advertise auto-negotiation capability                |
| <code>adv_cap_1000fdx</code> | Advertise 1000 full duplex capability                |
| <code>adv_cap_1000hdx</code> | Advertise 1000 half duplex capability                |
| <code>adv_cap_100fdx</code>  | Advertise 100 full duplex capability                 |
| <code>adv_cap_100hdx</code>  | Advertise 100 half duplex capability                 |
| <code>adv_cap_10fdx</code>   | Advertise 10 full duplex capability                  |
| <code>adv_cap_10hdx</code>   | Advertise 10 half duplex capability                  |
| <code>adv_rem_fault</code>   | Fault value reported by the local system to the peer |
|                              | 0 Link is good                                       |
|                              | 1 Fault                                              |

The `lp_cap_*` parameters exist as kernel statistics for an Ethernet device. The statistics are the advertised capabilities provided by the link partner on completion of auto-negotiation. If the capabilities match the capabilities provided in the local advertisement, the link can proceed to a link up state. If no match is found, the link remains down. In two other instances, `lp_cap_*` values might all be zero: (1) when a cable is not present and (2) when forced mode is enabled.

|                             |                                                     |
|-----------------------------|-----------------------------------------------------|
| <code>lp_cap_autoneg</code> | Link partner advertises auto-negotiation capability |
| <code>lp_cap_1000fdx</code> | Link partner advertises 1000 full duplex capability |
| <code>lp_cap_1000hdx</code> | Link partner advertises 1000 half duplex capability |
| <code>lp_cap_100fdx</code>  | Link partner advertises 100 full duplex capability  |
| <code>lp_cap_100hdx</code>  | Link partner advertises 100 half duplex capability  |
| <code>lp_cap_10fdx</code>   | Link partner advertises 10 full duplex capability   |
| <code>lp_cap_10hdx</code>   | Link partner advertises 10 half duplex capability   |
| <code>lp_rem_fault</code>   | Fault value the remote system reports               |
|                             | 0 Link is good                                      |
|                             | 1 Fault                                             |

The `xcvr_*` kernel statistics provide information about the physical layer device that is in use.

|                         |                                                                                                  |
|-------------------------|--------------------------------------------------------------------------------------------------|
| <code>xcvr_addr</code>  | MII address in the 0 to 31 range of the physical layer device in use for a given Ethernet device |
| <code>xcvr_id</code>    | MII transceiver manufacturer and device ID                                                       |
| <code>xcvr_inuse</code> | MII transceiver type, based on the following list:                                               |

|   |           |                                    |
|---|-----------|------------------------------------|
| 0 | other     | Undefined                          |
| 1 | none      | MII present, but nothing connected |
| 2 | 10Mb/s    | 10Mb/s Manchester encoding         |
| 3 | 100BaseT4 | 100 Mb/s 8B/6T                     |
| 4 | 100BaseX  | 100 Mb/s 4B/5B                     |
| 5 | 100BaseT2 | 100 Mb/s PAM5X5                    |
| 6 | 1000BaseX | 1000 Mb/s 8B/10B                   |
| 7 | 1000BaseT | 1000 Mb/s 4D-PAM5                  |

The above values define maximum capability. In many cases, lower speeds can occur. The `cap_*` statistics must be viewed to establish the range of capability.

The `link_*` kernel statistics show the link state at the local end of the connection.

|                          |   |                  |
|--------------------------|---|------------------|
| <code>link_up</code>     | 1 | Link is up       |
|                          | 0 | Link is down     |
| <code>link_duplex</code> | 2 | Full duplex link |
|                          | 1 | Half duplex link |
|                          | 0 | Unknown          |

The `cap_asym_pause`, `cap_pause`, `adv_cap_asym_pause`, and `adv_cap_pause` parameters do not follow the rules of other `cap_*` and `adv_cap_*` kstats or parameters. The `cap_*pause` kstats provide information about the capabilities supported by the device and constrain the values that may be set to the corresponding `adv_cap_*pause` parameters.

`cap_pause`            Symmetric pause capability.

`cap_asym_pause`      Asymmetric pause capability.

The `adv_cap_pause` and `adv_cap_asym_pause` statistics are limited by the available settings for `cap_pause` and `cap_asym_pause`. These statistics are read-only values whose settings may be administratively controlled by setting the `flowctrl` property supported by [dLadm\(1M\)](#). For a device that is fully capable of pausing both Rx (receive) and Tx (transmit) operations, the settings available are defined in the truth table that follows the `adv_cap_pause` and `adv_cap_asym_pause` parameter descriptions below.

`adv_cap_pause`            When `adv_cap_pause` is 1, the device can both assert and respond to flow control. This is the pre-Gigabit, symmetric mode of

operation, and implies a full (both send and receive) implementation of the PAUSE mechanism within the device. In addition, if `adv_cap_asym_pause` is 1, the device can operate either symmetrically or asymmetrically in either direction.

If `adv_cap_pause` is 0, advertised, flow-control behavior is determined by `adv_cap_asym_pause`. If the value of `adv_cap_asym_pause` is 1, the device can assert flow control, but cannot resend.

No flow control is available when both `adv_cap_pause` and `adv_cap_asym_pause` are 0.

`adv_cap_asym_pause` Asymmetric pause capability.

The `cap_asym_pause` and `cap_pause` statistics show the capability of a device and also limit the legal setting for `adv_cap_asym_pause` and `adv_cap_pause`. The following truth table describes the available `adv_cap_asym_pause` and `adv_cap_pause` settings limited by `cap_asym_pause` and `cap_pause` statistics. The abbreviations below are used in the table.

CA `cap_asym_pause`  
 CP `cap_pause`  
 AA `adv_cap_asym_pause`  
 AP `adv_cap_pause`

| CP | CA | AP | AA | Description                                                                                       |
|----|----|----|----|---------------------------------------------------------------------------------------------------|
| 0  | 0  | 0  | 0  | No pause in use.                                                                                  |
| 0  | 0  | x  | x  | Device not pause capable, cannot set.                                                             |
| 0  | 1  | 0  | 0  | Asymmetric Rx pause capable, but not advertised.                                                  |
| 0  | 1  | 0  | 1  | Asymmetric Rx pause capable and advertised.                                                       |
| 0  | 1  | 1  | 0  | Asymmetric Rx pause capable, but not advertised. Not capable of symmetric pause.                  |
| 0  | 1  | 1  | 1  | Asymmetric Rx pause capable and advertised. No symmetric pause capability or asymmetric Tx pause. |
| 1  | 0  | 0  | 0  | Symmetric pause capable, but not advertised.                                                      |
| 1  | 0  | 0  | 1  | Symmetric pause capable, advertising asymmetric Rx pause only.                                    |
| 1  | 0  | 1  | 0  | Symmetric pause capable, advertising symmetric Rx and Tx pause capability.                        |

|   |   |   |   |                                                                                 |
|---|---|---|---|---------------------------------------------------------------------------------|
| 1 | 0 | 1 | 1 | Symmetric pause capable and advertised.                                         |
| 1 | 1 | 0 | 0 | Symmetric and asymmetric pause capable, but not advertised.                     |
| 1 | 1 | 0 | 1 | Symmetric and asymmetric Tx pause capable. Only asymmetric Tx pause advertised. |
| 1 | 1 | 1 | 0 | Symmetric and symmetric Tx pause capable. Only symmetric pause advertised.      |
| 1 | 1 | 1 | 1 | Asymmetric Tx pause capable and advertised.                                     |

In the cases above, an error is posted when a device driver cannot advertise. A new setting is ignored and values revert to the previous setting.

The `lp_cap_pause` and the `lp_cap_asym_pause` provide the advertised capabilities of the link partners.

`lp_cap_pause`            When `lp_cap_pause` is 1, the link-partner can both assert and respond to flow control. This is the pre-Gigabit, symmetric mode of operation, and implies a full (both send and receive) implementation of the PAUSE mechanism within the device. In addition, if `lp_cap_asym_pause` is 1, the link-partner can operate either symmetrically or asymmetrically in either direction.

If `lp_cap_pause` is 0, the flow-control behavior supported by the link-partner is determined by `lp_cap_asym_pause`. If the value of `lp_cap_asym_pause` is 1, the link-partner can assert flow control, but cannot respond to any pause-frames sent to it.

No flow control is available when both `lp_cap_pause` and `lp_cap_asym_pause` are 0.

`lp_cap_asym_pause`    Asymmetric pause capability

When `adv_*pause_cap` and `lp_*pause_cap` are compared on completion of auto-negotiation, the chosen flow control mechanism for the link depends on what is most meaningful.

`link_tx_pause`        Link partner can assert flow control by sending pause frames when congestion is experienced.

`link_rx_pause`        Link partner can respond to pause frames received.

The following truth table illustrates the meaningful flow control combinations related to local and link partner configurations. The abbreviations below are used in the table.

AA        `adv_cap_asym_pause`

AP        `adv_cap_pause`

LAC    `lp_cap_asym_pause`  
 LPC    `lp_cap_pause`  
 LA     `link_asym_pause`  
 LP     `link_pause`

| AA | AP | LAC | LPC | LA | LP | Description                                                                    |
|----|----|-----|-----|----|----|--------------------------------------------------------------------------------|
| 1  | 0  | 1   | 1   | 1  | 0  | Local station will Tx a pause when Rx is congested.                            |
| 0  | 1  | 0   | 1   | 0  | 1  | Flow control in both Rx and Tx directions.                                     |
| x  | 1  | 1   | 0   | 1  | 1  | Local station honors received Pause frames by temporarily suspending Transmit. |
| x  | x  | x   | x   | 0  | 0  | All other combinations: Flow control not available on the link                 |

When forced mode is enabled, the current setting of `adv_cap_asym_pause` and `adv_cap_pause` are used for the link. The `link_asym_pause` and `link_pause` become equal to the current `adv_cap_asym_pause` and `adv_cap_pause` settings. The above table also applies in forced mode, but the link partner configuration must be checked to verify that flow control is operating on the link.

**See Also** [dladm\(1M\)](#), [driver.conf\(4\)](#), [bge\(7D\)](#), [d1pi\(7P\)](#), [eri\(7D\)](#), [gld\(7D\)](#), [hme\(7D\)](#), [qfe\(7d\)](#)

**Notes** When `adv_cap_autoneg` is set to 0, the highest priority speed and duplex is used for forced mode.

The highest priority is the highest speed at full duplex. The lowest priority is the lowest speed at half duplex.

MII transceivers can exist internally to a system or can be connected to an external MII connector. Typically, an internal transceiver has an `xcvr_addr` of 1, while an external connection has an `xcvr_addr` of 0.

**Name** ipfilter – IP packet filtering software

**Description** IP Filter is software that provides packet filtering capabilities on a Solaris system. On a properly setup system, it can be used to build a firewall.

Solaris IP Filter is installed with the Solaris operating system. However, packet filtering is not enabled by default. See [ipf\(1M\)](#) for a procedure to enable and activate the IP Filter feature.

**Interaction with Location Profiles** IP Filter configuration and activation is managed in Location profiles (refer to [netcfg\(1M\)](#) for more information about location profiles). These profiles are either fixed, meaning the network configuration is being managed in the traditional way, or reactive, meaning the network configuration is being managed automatically, reacting to changes in the network environment according to policy rules specified in the profiles.

When a fixed location (there can currently be only one, the DefaultFixed location) is active, changes made to the SMF repository will be applied to the location when it is disabled, and thus will be restored if that location is later re-enabled.

When a reactive location is active, changes should not be applied directly to the SMF repository; these changes will not be preserved in the location profile, and will thus be lost if the location is disabled, or if the system's network configuration, as managed by `svc:/network/physical:default` and `svc:/network/location:default`, is refreshed or restarted. Changes should instead be applied to the location itself, using the `netcfg` command; this will save the change to the location profile repository, and will also apply it to the SMF repository (if the change is made to the currently active location).

The `ipfilter` SMF service will be enabled if an IPv4 filter configuration file is specified in the `ipfilter-config-file` property. To enable additional types of filtering, the `ipfilter-v6-config-file`, `ipnat-config-file`, and `ippool-config-file` properties may also be specified.

**Service** The `ipfilter` SMF service supports the start, stop, restart, and refresh methods. The methods are invoked using [svcadm\(1M\)](#).

- |                      |                                                                                                                                                                                                                                                                 |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>start</code>   | Loads the <code>ipfilter</code> kernel module and activates any firewall or NAT rules as per the configuration.                                                                                                                                                 |
| <code>stop</code>    | Clears out all of the applied firewall and NAT rules and any active session information that has been created. Stopping the service with networking enabled should only be performed when there is no risk of any network traffic being able to enter the host. |
| <code>restart</code> | Performs a stop and then start of the <code>ipfilter</code> service. Using this method on an active firewall results in a window of exposure where traffic can enter and/or pass through the firewall without being filtered.                                   |
| <code>refresh</code> | Loads the current configuration and switches over from the old configuration to the new one without there being a moment in time when neither security policy is in active use.                                                                                 |

**Host-Based Firewall** To simplify IP Filter configuration management, a firewall framework is created to allow users to configure IP Filter by expressing firewall policy at system and service level. Given the user-defined firewall policy, the framework generates a set of IP Filter rules to enforce the desired system behavior. Users specify system and service firewall policies that allow or deny network traffic from certain hosts, subnets, and interface(s). The policies are translated into a set of active IPF rules to enforce the specified firewall policies.

Users can still specify their own `ipf` rule file if they choose not to take advantage of the framework. See [ipf\(1M\)](#) for how to enable customized rules and [ipf\(4\)](#) to find out `ipf` rule syntax.

IPF uses the [ipmon\(1M\)](#) service to log firewall events. The `ipmon` SMF service depends on the `ipfilter` SMF service. The `ipmon` service gets enabled temporarily by the `ipfilter` service start method as soon as `ipfilter` gets enabled automatically by “`svcadm enable ipfilter`”.

**Model** This section describes the host-based firewall framework. See [svc.ipfd\(1M\)](#) for details on how to configure firewall policies.

A three-layer approach with different precedence levels helps the user achieve the desired behaviors.

#### Global Default

Global Default - Default system-wide firewall policy. This policy is automatically inherited by all services unless services modify their firewall policy.

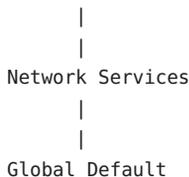
#### Network Services

Higher precedence than Global Default. A service's policy allows/disallows traffic to its specific ports, regardless of Global Default policy.

#### Global Override

Another system-wide policy that takes precedence over the needs of specific services in Network Services layer.

#### Global Override



A firewall policy includes a firewall mode and an optional set of network sources. Network sources are IP addresses, subnets, and local network interfaces, from all of which a system can receive incoming traffic. The basic set of firewall modes are:

#### None

No firewall, allow all incoming traffic.

**Deny**

Allow all incoming traffic but deny from specified source(s).

**Allow**

Deny all incoming traffic but allow from specified source(s).

**Layers in Detail** The first system-wide layer, Global Default, defines a firewall policy that applies to *any* incoming traffic, for example, allowing or blocking all traffic from an IP address. This makes it simple to have a policy that blocks all incoming traffic or all incoming traffic from unwanted source(s).

The Network Services layer contains firewall policies for local programs that provide service to remote clients, for example, `telnetd`, `sshd`, and `htptd`. Each of these programs, a network service, has its own firewall policy that controls access to its service. Initially, a service's policy is set to inherit Global Default policy, a "Use Global Default" mode. This makes it simple to set a single policy, at the Global Default layer, that can be inherited by all services.

When a service's policy is different from Global Default policy, the service's policy has higher precedence. If Global Default policy is set to block all traffic from a subnet, the SSH service could be configured to allow access from certain hosts in that subnet. The set of all policies for all network services comprises the Network Service layer.

The second system-wide layer, Global Override, has a firewall policy that also applies to any incoming network traffic. This policy has highest precedence and overrides policies in the other layers, specifically overriding the needs of network services. The example is when it is desirable to block known malicious source(s) regardless of services' policies.

**User Interaction** This framework leverages IP Filter functionality and is active only when `svc:/network/ipfilter` is enabled and inactive when `network/ipfilter` is disabled. Similarly, a network service's firewall policy is only active when that service is enabled and inactive when the service is disabled. A system with an active firewall has IP Filter rules for each running/enabled network service and system-wide policy(s) whose firewall mode is not None.

A user configures a firewall by setting the system-wide policies and policy for each network service. See `svc.ipfd(1M)` on how to configure a firewall policy.

The firewall framework composes of policy configuration and a mechanism to generate IP Filter rules from the policy and applying those rules to get the desired IP Filter configuration. A quick summary of the design and user interaction:

- system-wide policy(s) are stored in `network/ipfilter`
- network services' policies are stored in each SMF service
- a user activates a firewall by enabling `network/ipfilter` (see `ipf(1M)`)
- a user activates/deactivate a service's firewall by enabling/disabling that network service
- changes to system-wide or per-service firewall policy results in an update to the system's firewall rules

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Interface Stability | Committed       |

**See Also** [svcs\(1\)](#), [ipf\(1M\)](#), [ipmon\(1M\)](#), [ipnat\(1M\)](#), [netcfg\(1M\)](#), [svcadm\(1M\)](#), [svc.ipfd\(1M\)](#), [ipf\(4\)](#), [ipnat\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

*System Administration Guide: IP Services*

**Notes** The `ipfilter` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/ipfilter:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

IP Filter startup configuration files are stored in `/etc/ipf`.

**Name** isalist – the native instruction sets known to Solaris software

**Description** The possible instruction set names returned by `isalist(1)` and the `SI_ISALIST` command of `sysinfo(2)` are listed here.

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 `isalist(1)` and `sysinfo(2)`. In the following list of values, numbered entries generally represent increasing performance; lettered entries are either mutually exclusive or cannot be ordered.

This feature is obsolete and may be removed in a future version of Solaris. The lists below do not reflect all the extensions that have been made by modern processors. See `getisax(2)` for a better way to handle instruction set extensions.

SPARC Platforms Where appropriate, correspondence with a given value of the `-xarch` option of Sun's C 4.0 compiler is indicated. Other compilers might have similar options.

- |                                      |                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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, FSMULD, and all floating point operations on quad operands) can be emulated by the kernel on certain systems.                                                  |
| 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.                                                                                                                                                                                                                    |
| 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.                                                                                                                                                |
| 3. <code>sparcv8</code>              | Like <code>sparcv8-fsmuld</code> , except that FSMULD 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.                                                                                                                                                        |
| 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. |
| 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.                                                                          |
| 5b. <code>sparcv8plus+fmuladd</code> | Like <code>sparcv8plus</code> , with the addition of the Fujitsu SPARC64 floating multiply-add and multiply-subtract instructions.                                                                                                                                                                                                                |

---

|               |                                  |                                                                                                                                                                                                                                                                                         |
|---------------|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               | 6. <code>sparcv9</code>          | Indicates the SPARC V9 instruction set, as defined in <i>The SPARC Architecture Manual, Version 9</i> , Prentice-Hall, 1994.                                                                                                                                                            |
|               | 7a. <code>sparcv9+vis</code>     | Like <code>sparcv9</code> , with the addition of the UltraSPARC I Visualization Instructions.                                                                                                                                                                                           |
|               | 7b. <code>sparcv9+vis2</code>    | Like <code>sparcv9</code> , with the addition of the UltraSPARC III Visualization Instructions.                                                                                                                                                                                         |
|               | 7c. <code>sparcv9+fmuladd</code> | Like <code>sparcv9</code> , with the addition of the Fujitsu SPARC64 floating multiply-add and multiply-subtract instructions.                                                                                                                                                          |
| x86 Platforms | 1. <code>i386</code>             | The Intel 80386 instruction set, as described in <i>The i386 Microprocessor Programmer's Reference Manual</i> .                                                                                                                                                                         |
|               | 2. <code>i486</code>             | The Intel 80486 instruction set, as described in <i>The i486 Microprocessor Programmer's Reference Manual</i> . (This is effectively <code>i386</code> , plus the <code>CMPXCHG</code> , <code>BSWAP</code> , and <code>XADD</code> instructions.)                                      |
|               | 3. <code>pentium</code>          | The Intel Pentium instruction set, as described in <i>The Pentium Processor User's Manual</i> . (This is effectively <code>i486</code> , plus the <code>CPU_ID</code> instruction, and any features that the <code>CPU_ID</code> instruction indicates are present.)                    |
|               | 4. <code>pentium+mmx</code>      | Like <code>pentium</code> , with the MMX instructions guaranteed present.                                                                                                                                                                                                               |
|               | 5. <code>pentium_pro</code>      | The Intel PentiumPro instruction set, as described in <i>The PentiumPro Family Developer's Manual</i> . (This is effectively <code>pentium</code> , with the <code>CMOVcc</code> , <code>FCMOVcc</code> , <code>FCOMI</code> , and <code>RDPMC</code> instructions guaranteed present.) |
|               | 6. <code>pentium_pro+mmx</code>  | Like <code>pentium_pro</code> , with the MMX instructions guaranteed present.                                                                                                                                                                                                           |
|               | 7. <code>amd64</code>            | The AMD Opteron instruction set, as described in the <i>AMD64 Architecture Programmer's Manual</i> .                                                                                                                                                                                    |

**See Also** [isalist\(1\)](#), [getisax\(2\)](#), [sysinfo\(2\)](#)

**Name** kerberos – overview of Solaris Kerberos implementation

**Description** The Solaris Kerberos implementation, hereafter sometimes shortened to “Kerberos,” authenticates clients in a network environment, allowing for secure transactions. (A client may be a user or a network service.) Kerberos validates the identity of a client and the authenticity of transferred data. Kerberos is a *single-sign-on* system, meaning that a user needs to provide a password only at the beginning of a session. The Solaris Kerberos implementation is based on the Kerberos(TM) system developed at MIT, and is compatible with Kerberos V5 systems over heterogeneous networks.

Kerberos works by granting clients *tickets*, 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 Kerberos ticket may rlogin into another machine running Kerberos without having to identify itself. Because each client has a unique ticket, its identity is guaranteed.

To obtain tickets, a client must first initialize the Kerberos session, either by using the `kinit(1)` command or a PAM module. (See `pam_krb5(5)`.) `kinit` prompts for a password, and then communicates with a *Key Distribution Center* (KDC). The KDC returns a *Ticket-Granting Ticket* (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 `klist(1)` command.

Tickets are valid according to the system *policy* 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.

Tickets can be renewed using `kinit`. Tickets are also *forwardable*, allowing you to use a ticket granted on one machine on a different host. Tickets can be destroyed by using `kdestroy(1)`. It is a good idea to include a call to `kdestroy` in your `.logout` file.

Under Kerberos, a client is referred to as a *principal*. A principal takes the following form:

```
primary/instance@REALM
```

`primary`     A user, a host, or a service.

`instance`    A qualification of the primary. If the primary is a host — indicated by the keyword `host` — 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 `admin` or `root`, are privileged.

`realm`        The Kerberos equivalent of a domain; in fact, in most cases the realm is directly mapped to a DNS domain name. Kerberos 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), Kerberos 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\(4\)](#), [krb5envvar\(5\)](#)

*Oracle Solaris 11.1 Administration: Security Services*

**Notes** In previous releases of the Solaris operating system, the Solaris Kerberos implementation was referred to as the “Sun Enterprise Authentication Mechanism” (SEAM).

If you enter your username and kinit responds with this message:

```
Principal unknown (kerberos)
```

you have not been registered as a Kerberos user. See your system administrator or the *Oracle Solaris 11.1 Administration: Security Services*.

**Name** krb5\_auth\_rules – overview of Kerberos V5 authorization

**Description** When kerberized versions of the ftp, rcp, rlogin, rsh, ssh, telnet, or ssh clients are used to connect to a server, the identity of the originating user must be authenticated to the Kerberos V5 authentication system. Account access can then be authorized if appropriate entries exist in the `~/ .k5login` file, the `gsscred` table, or if the default GSS/Kerberos authentication rules successfully map the Kerberos principal name to Unix login name.

To avoid security problems, the `~/ .k5login` file must be owned by the remote user on the server the client is attempting to access. The file should contain a private authorization list comprised of Kerberos principal names of the form *principal/instance@realm*. The */instance* variable is optional in Kerberos principal names. For example, different principal names such as `jdb@ENG.ACME.COM` and `jdb/happy.eng.acme.com@ENG.ACME.COM` would each be legal, though not equivalent, Kerberos principals. The client is granted access if the `~/ .k5login` file is located in the login directory of the remote user account and if the originating user can be authenticated to one of the principals named in the file. See [gkadmin\(1M\)](#) and [kadm5.acl\(4\)](#) for more information on Kerberos principal names.

When no `~/ .k5login` file is found in the remote user's login account, the Kerberos V5 principal name associated with the originating user is checked against the `gsscred` table. If a `gsscred` table exists and the principal name is matched in the table, access is granted if the Unix user ID listed in the table corresponds to the user account the client is attempting to access. If the Unix user ID does not match, access is denied. See [gsscred\(1M\)](#).

For example, an originating user listed in the `gsscred` table with the principal name `jdb@ENG.ACME.COM` and the `uid 23154` is granted access to the `jdb-user` account if `23154` is also the `uid` of `jdb-user` listed in the user account database. See [passwd\(4\)](#).

Finally, if there is no `~/ .k5login` file and the Kerberos V5 identity of the originating user is not in the `gsscred` table, or if the `gsscred` table does not exist, the client is granted access to the account under the following conditions (default GSS/Kerberos auth rules):

- The user part of the authenticated principal name is the same as the Unix account name specified by the client.
- The realm part of the client and server are the same, unless the [krb5.conf\(4\)](#) `auth_to_local_realm` parameter is used to create equivalence.
- The Unix account name exists on the server.

For example, if the originating user has the principal name `jdb@ENG.ACME.COM` and if the server is in realm `SALES.ACME.COM`, the client would be denied access even if `jdb` is a valid account name on the server. This is because the realms `SALES.ACME.COM` and `ENG.ACME.COM` differ.

The `krb5.conf(4)` `auth_to_local_realm` parameter also affects authorization. Non-default realms can be equated with the default realm for authenticated name-to-local name mapping.

- Files**
- `~/ .k5login` Per user-account authorization file.
  - `/etc/passwd` System account file. This information may also be in a directory service. See [passwd\(4\)](#).

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Interface Stability | Committed       |

**See Also** [ftp\(1\)](#), [rcp\(1\)](#), [rsh\(1\)](#), [telnet\(1\)](#), [gkadmin\(1M\)](#), [gsscred\(1M\)](#), [kadm5.acl\(4\)](#), [krb5.conf\(4\)](#), [passwd\(4\)](#), [attributes\(5\)](#), [gss\\_auth\\_rules\(5\)](#)

**Name** krb5envvar – Kerberos environment variables

**Description** The Kerberos mechanism provides a number of environment variables to configure different behavior in order to meet applications' needs. Environment variables used within the Kerberos mechanism are:

#### KRB5\_KTNAME

Used by the mechanism to specify the location of the key table file. The variable can be set to the following value:

[ [*<kt type>*: ]*<file name>*]

where *<kt type>* can be FILE or WRFILE. FILE is for read operations; WRFILE is for write operations. *<file name>* is the location of the keytab file.

r

If KRB5\_KTNAME is not defined, the default value is:

FILE: /etc/krb5/krb5.keytab

The keytab file is used to store credentials persistently and is used commonly for service daemons.

Specifying the FILE type assumes that the subsequent operations on the associated file are readable by the invoking process. Care must be taken to ensure that the file is readable only by the set of principals that need to retrieve their unencrypted keys.

The WRFILE type is used by the [kadmin\(1M\)](#) command. Specifying this type allows the administrator to designate an alternate keytab file to write to without using extra command line arguments for file location.

#### KRB5CCNAME

Used by the mechanism to specify the location of the credential cache. The variable can be set to the following value:

[ [*<cc type>*: ]*<file name>*]

where *<cc type>* can be FILE or MEMORY. *<file name>* is the location of the principal's credential cache.

If KRB5CCNAME is not defined, the default value is:

FILE: /tmp/krb5cc\_*<uid>*

where *<uid>* is the user id of the process that created the cache file.

The credential cache file is used to store tickets that have been granted to the principal.

Specifying the FILE types assumes that subsequent operations on the associated file are readable and writable by the invoking process. Care must be taken to ensure that the file is

accessible only by the set of principals that need to access their credentials. If the credential file is in a directory to which other users have write access, you need to set that directory's sticky bit (see [chmod\(1\)](#)).

The MEMORY credential cache type is used only in special cases, such as when making a temporary cache for the life of the invoking process.

#### KRB5RCNAME

Used by the mechanism to specify the type and location of the replay cache. The variable can be set to the following value:

```
[ [<rc type>: ]<file name> ]
```

where *<rc type>* can be either FILE, MEMORY, or NONE. *<file name>* is relevant only when specifying the replay cache file type.

If not defined, the default value is:

```
FILE: /var/krb5/rcache/root/rc_<service>
```

...if the process is owned by root, or:

```
FILE: /var/krb5/rcache/rc_<service>
```

...if the process is owned by a user other than root. *<service>* is the service process name associated with the replay cache file.

The replay cache is used by Kerberos to detect the replay of authentication data. This prevents people who capture authentication messages on the network from authenticating to the server by resending these messages.

When specifying the FILE replay cache type, care must be taken to prevent the replay cache file from being deleted by another user. Make sure that every directory in the replay cache path is either writable only by the owner of the replay cache or that the sticky bit (“t”) is set on every directory in the replay cache path to which others have write permission.

When specifying the MEMORY replay cache type you need to weigh the trade-off of performance against the slight security risk created by using a non-persistent cache. The risk occurs during system reboots when the following condition obtains:

- The duration from the last write to the replay cache before reboot to the point when the Kerberized server applications are running is less than the Kerberos clockskew (see [krb5.conf\(4\)](#)).

When specifying the NONE replay cache time you need to understand that this disables the replay cache, and all security risks that this presents. This includes all the risks outlined in this section of the man page.

Under this condition, the server applications can accept a replay of Kerberos authentication data (up to the difference between the time of the last write and the clockskew). Typically, this is a small window of time. If the server applications take longer than the clockskew to start accepting connections there is no replay risk.

The risk described above is the same when using FILE replay cache types when the replay cache resides on swap file systems, such as /tmp and /var/run.

The performance improvement in MEMORY replay cache types over FILE types is derived from the absence of disk I/O. This is true even if the FILE replay cache is on a memory-backed file system, such as swap (/tmp and /var/run).

Note that MEMORY-type caches are per-process caches, therefore use of these types of caches must be carefully considered. One example of where MEMORY-type caches can be problematic is when an application uses more than one process for establishing security contexts. In such a case, memory replay caches are not shared across the processes, thus allowing potential for replay attacks.

#### KRB5\_CONFIG

Allows you to change the default location of the /etc/krb5/krb5.conf file to enable the Kerberos library code to read configuration parameters from another file specified by KRB5\_CONFIG. For example (using kinit from [ksh\(1\)](#)):

```
KRB5_CONFIG=/var/tmp/krb5.conf kinit
```

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE            |
|---------------------|----------------------------|
| Availability        | system/security/kerberos-5 |
| Interface Stability | Uncommitted                |

**See Also** [chmod\(1\)](#), [kinit\(1\)](#), [klist\(1\)](#), [ksh\(1\)](#), [kadmind\(1M\)](#), [kadmin\(1M\)](#), [krb5.conf\(4\)](#), [attributes\(5\)](#), [kerberos\(5\)](#)

**Name** kssl, KSSL – kernel SSL proxy

**Description** The KSSL is a transparent server side proxy for SSL/TLS protocol. It provides processing of SSL traffic in the kernel and thus improving performance by avoiding context switches and directly accessing kernel providers of Oracle Solaris Crypto Framework. With KSSL it is possible to provide SSL protection even for applications which are only able to communicate in clear text over TCP.

KSSL is configured in the kernel and passes/accepts clear text data from an application. Together they are visible to the clients as single SSL server.

The server side application for which KSSL is configured is unaware that it is receiving data previously protected by SSL. KSSL receives SSL traffic on one port, for example, 443, performs processing and passes clear text data to the application listening on another port, for example, 8080. Similarly, for the outgress direction, application sends clear text data and KSSL produces SSL records and sends them to the client. Therefore, the application does not have to be setup for SSL.

Multiple KSSL instances can be configured on the system, each with separate set of properties such as port, certificate, key or cipher suites. See [ksslcfg\(1M\)](#). Each KSSL instance in the kernel is tracked as SMF service. See [smf\(5\)](#).

KSSL provides SSL processing for records passed with TCP over both IPv4 and IPv6.

KSSL supports the following protocols: SSLv3, TLSv1.0

**See Also** [ksslcfg\(1M\)](#), [smf\(5\)](#)

T. Dierks, C. Allen, *RFC 2246, The TLS Protocol Version 1.0*, The Internet Society, 1999.

**Name** labels – Solaris Trusted Extensions label attributes

**Description** Labels are attributes that are used in mandatory policy decisions. Labels are associated, either explicitly or implicitly, with all subjects (generally processes) and objects (generally things with data such as files) that are accessible to subjects. The default Trusted Extensions mandatory policy labels are defined by a site's security administrator in [label\\_encodings\(4\)](#).

**Mandatory Policy** Various mandatory policies might be delivered in the lifetime of Solaris Trusted Extensions.

The default mandatory policy of Trusted Extensions is a Mandatory Access Control (MAC) policy that is equivalent to that of the Bell-LaPadula Model of the Lattice, the Simple Security Property, and the \*-Property (Star Property), with restricted write up. The default mandatory policy is also equivalent to the Goguen and Meseguer model of Non-Inteference.

For this MAC policy, two labels are always defined: `admin_low` and `admin_high`. The site's security administrator defines all other labels in [label\\_encodings\(4\)](#). `admin_low` is associated with all normal user readable (viewable) Trusted Extensions objects. `admin_high` is associated with all other Trusted Extensions objects. Only administrative users have MAC read (view) access to `admin_high` objects and only administrative users have MAC write (modify) access to `admin_low` objects or `admin_high` objects.

**Human Readable Labels** Users interact with labels as strings. Graphical user interfaces and command line interfaces present the strings as defined in [label\\_encodings\(4\)](#). Human readable labels are classified at the label that they represent. Thus the string for a label A is only readable (viewable, translatable to or from human readable to opaque `m_label_t`) by a subject whose label allows read (view) access to that label.

**Internal Text Labels** In order to store labels in publicly accessible (`admin_low`) name service databases, an unclassified internal text form is used. This textual form is not intended to be used in any interfaces other than those that are provided with the Trusted Extensions software release that created this textual form of the label.

**Labels and Applications** Applications interact with labels as opaque (`m_label_t`) structures. The semantics of these opaque structures are defined by a string to `m_label_t` translation. This translation is defined in [label\\_encodings\(4\)](#). Various Application Programming Interfaces (API) translate between strings and `m_label_t` structures. Various APIs test access of subject-related labels to object-related labels.

**Attributes** See [attributes\(5\)](#) for description of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Interface Stability | See below.      |

The labels implementation is Committed for systems that implement the Defense Intelligence Agency (DIA) MAC policy of [label\\_encodings\(4\)](#). Other policies might exist in a future release of Trusted Extensions that might make obsolete or supplement [label\\_encodings](#).

---

Internal text labels are Not-an-Interface and might change with any release of Trusted Extensions. They are intended only for input and generation on the same release of Trusted Extensions software.

As a potential porting aid for Trusted Solaris 8 applications, the opaque structure names `bs_label_t`, `blevel_t`, and `bclear_t` are defined to be equivalent to `m_label_t`. Like `m_label_t`, these types must be ported as opaque pointers. The same must be done with the various Trusted Solaris 8 label interfaces. These Trusted Solaris 8 structures and interfaces are Obsolete and might be removed from a future release of Trusted Extensions.

**See Also** `chk_encodings(1M)`, `blcompare(3TSOL)`, `label_to_str(3TSOL)`, `m_label_alloc(3TSOL)`, `m_label_dup(3TSOL)`, `m_label_free(3TSOL)`, `str_to_label(3TSOL)`, `label_encodings(4)`, `attributes(5)`

Bell, D. E., and LaPadula, L. J. *Secure Computer Systems: Unified Exposition and Multics Interpretation*, MTR-2997 Rev. 2, MITRE Corp., Bedford Mass., March 1976. NTIS AD-A023 588/7.

Goguen, J. A., and Meseguer, J.: *Security Policies and Security Models*, Proceedings 1982 Symposium on Security and Privacy, IEEE Computer Society Press, 1982, p 11-20.

Goguen, J. A., and Meseguer, J.: *Unwinding and Interference Control*, Proceedings 1984 Symposium on Security and Privacy, IEEE Computer Society Press, 1984, p 75-86.

*Compartmented Mode Workstation Labeling: Encodings Format*

**Notes** The functionality described on this manual page is available only if the system is configured with Trusted Extensions.

**Name** largefile – large file status of utilities

**Description** A *large file* is a regular file whose size is greater than or equal to 2 Gbyte ( $2^{31}$  bytes). A *small file* is a regular file whose size is less than 2 Gbyte.

Large file aware utilities A utility is called *large file aware* 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 `file` utility supports the `-m` option for an alternative “magic” file and the `-f` 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.

The following `/usr/bin` utilities are large file aware:

|                        |                       |                       |                         |                         |
|------------------------|-----------------------|-----------------------|-------------------------|-------------------------|
| <code>adb</code>       | <code>aliasadm</code> | <code>awk</code>      | <code>bdiff</code>      | <code>cat</code>        |
| <code>chgrp</code>     | <code>chmod</code>    | <code>chown</code>    | <code>cksum</code>      | <code>cmp</code>        |
| <code>compress</code>  | <code>cp</code>       | <code>csd</code>      | <code>csplit</code>     | <code>cut</code>        |
| <code>dd</code>        | <code>dircmp</code>   | <code>du</code>       | <code>egrep</code>      | <code>fgrep</code>      |
| <code>file</code>      | <code>find</code>     | <code>ftp</code>      | <code>getconf</code>    | <code>grep</code>       |
| <code>gzip</code>      | <code>head</code>     | <code>join</code>     | <code>jsh</code>        | <code>ksh88</code>      |
| <code>ksh</code>       | <code>ln</code>       | <code>ls</code>       | <code>mailcompat</code> | <code>mailstats</code>  |
| <code>mdb</code>       | <code>mkdir</code>    | <code>mkfifo</code>   | <code>more</code>       | <code>mv</code>         |
| <code>nawk</code>      | <code>page</code>     | <code>paste</code>    | <code>pathchk</code>    | <code>pg</code>         |
| <code>praliases</code> | <code>rcp</code>      | <code>remsh</code>    | <code>rksh88</code>     | <code>rksh</code>       |
| <code>rm</code>        | <code>rmdir</code>    | <code>rsh</code>      | <code>sed</code>        | <code>sh</code>         |
| <code>sort</code>      | <code>split</code>    | <code>sum</code>      | <code>tail</code>       | <code>tar</code>        |
| <code>tee</code>       | <code>test</code>     | <code>touch</code>    | <code>tr</code>         | <code>uncompress</code> |
| <code>uudcode</code>   | <code>uuencode</code> | <code>vacation</code> | <code>wc</code>         | <code>zcat</code>       |

The following `/usr/xpg4/bin` utilities are large file aware:

|                    |                    |                    |                    |                 |
|--------------------|--------------------|--------------------|--------------------|-----------------|
| <code>awk</code>   | <code>cp</code>    | <code>chgrp</code> | <code>chown</code> | <code>du</code> |
| <code>egrep</code> | <code>fgrep</code> | <code>file</code>  | <code>grep</code>  | <code>ln</code> |

---

```

ls             more             mv             rm             sed
sh            sort            tail          tr

```

The following `/usr/xpg6/bin` utilities are large file aware:

```

getconf      ls             tr

```

The following `/usr/sbin` utilities are large file aware:

```

editmap      install      makemap      mkfile      mknod
mmdir       swap

```

The following `/usr/lib` utilities are large file aware:

```

mail.local   sendmail     smrsh

```

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.

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`

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/xpg6/bin` utility is large file safe:

ed

The following `/usr/sbin` utilities are large file safe:

lpfilter            lpforms

The following `/usr/ucb` utilities are large file safe:

Mail                lpr

**See Also** [lf64\(5\)](#), [lfcompile\(5\)](#), [lfcompile64\(5\)](#)

**Name** ldap – LDAP as a naming repository

**Description** LDAP refers to Lightweight Directory Access Protocol, which is an industry standard for accessing directory servers. By initializing the client using `ldapclient(1M)` and using the keyword `ldap` in the name service switch file, `/etc/nsswitch.conf`, Oracle Solaris clients can obtain naming information from an LDAP server. Information such as usernames, hostnames, and passwords are stored on the LDAP server in a Directory Information Tree or DIT. The DIT consists of entries which in turn are composed of attributes. Each attribute has a type and one or more values.

Oracle Solaris LDAP clients use the LDAP v3 protocol to access naming information from LDAP servers. The LDAP server must support the object classes and attributes defined in *RFC2307bis (draft)*, which maps the naming service model on to LDAP. As an alternate to using the schema defined in *RFC2307bis (draft)*, the system can be configured to use other schema sets and the schema mapping feature is configured to map between the two. Refer to the *Oracle Solaris Administration: Naming and Directory Services* for more details.

The `ldapclient(1M)` utility can make an Oracle Solaris machine an LDAP client by setting up the appropriate directories, files, and configuration information. The LDAP client caches this configuration information in local cache files. This configuration information is accessed through the `ldap_cachemgr(1M)` daemon. This daemon also refreshes the information in the configuration files from the LDAP server, providing better performance and security. The `ldap_cachemgr` must run at all times for the proper operation of the naming services.

There are two types of configuration information, the information available through a profile, and the information configured per client. The profile contains all the information as to how the client accesses the directory. The credential information for proxy user is configured on a per client basis and is not downloaded through the profile.

The profile contains server-specific parameters that are required by all clients to locate the servers for the desired LDAP domain. This information could be the server's IP address and the search base Distinguished Name (DN), for instance. It is configured on the client from the default profile during client initialization and is periodically updated by the `ldap_cachemgr` daemon when the expiration time has elapsed.

Client profiles can be stored on the LDAP server and can be used by the `ldapclient` utility to initialize an LDAP client. Using the client profile is the easiest way to configure a client machine. See `ldapclient(1M)`.

Credential information includes client-specific parameters that are used by a client. This information could be the Bind DN (LDAP “login” name) of the client and the password. If these parameters are required, they are manually defined during the initialization through `ldapclient(1M)`.

The naming information is stored in containers on the LDAP server. A container is a non-leaf entry in the DIT that contains naming service information. Containers are similar to maps in

NIS. A default mapping between the NIS databases and the containers in LDAP is presented below. The location of these containers as well as their names can be overridden through the use of `serviceSearchDescriptors`. For more information, see `ldapClient(1M)`.

| Database   | Object Class    | Container                 |
|------------|-----------------|---------------------------|
| passwd     | posixAccount    | ou=people,dc=...          |
|            | shadowAccount   |                           |
| group      | posixGroup      | ou=Group,dc=...           |
| services   | ipService       | ou=Services,dc=...        |
| protocols  | ipProtocol      | ou=Protocols,dc=...       |
| rpc        | oncRpc          | ou=Rpc,dc=...             |
| hosts      | ipHost          | ou=Hosts,dc=...           |
| ipnodes    | ipHost          | ou=Hosts,dc=...           |
| ethers     | ieee802Device   | ou=Ethers,dc=...          |
| bootparams | bootableDevice  | ou=Ethers,dc=...          |
| networks   | ipNetwork       | ou=Networks,dc=...        |
| netmasks   | ipNetwork       | ou=Networks,dc=...        |
| netgroup   | nisNetgroup     | ou=Netgroup,dc=...        |
| aliases    | mailGroup       | ou=Aliases,dc=...         |
| publickey  | nisKeyObject    |                           |
| generic    | nisObject       | nisMapName=...,dc=...     |
| printers   | printerService  | ou=Printers,dc=...        |
| auth_attr  | SolarisAuthAttr | ou=SolarisAuthAttr,dc=... |
| prof_attr  | SolarisProfAttr | ou=SolarisProfAttr,dc=... |
| exec_attr  | SolarisExecAttr | ou=SolarisProfAttr,dc=... |
| user_attr  | SolarisUserAttr | ou=people,dc=...          |

The security model for clients is defined by a combination of the credential level to be used, the authentication method, and the PAM modules to be used. The credential level defines what credentials the client should use to authenticate to the directory server, and the authentication method defines the method of choice. Both these can be set with multiple values. The Oracle Solaris LDAP supports the following values for credential level :

anonymous  
proxy  
self

The Oracle Solaris LDAP supports the following values for authentication method:

none  
simple  
sasl/CRAM-MD5  
sasl/DIGEST-MD5  
sasl/GSSAPI  
tls:simple  
tls:sasl/CRAM-MD5  
tls:sasl/DIGEST-MD5

When the credential level is configured as `self`, DNS must be configured and the authentication method must be `sasl/GSSAPI`. The hosts and ipnodes in `/etc/nsswitch.conf` must be configured to use DNS, for example `hosts: dns files` and `ipnodes: dns files`.

`sasl/GSSAPI` automatically uses GSSAPI confidentiality and integrity options, if they are configured on the directory server.

The credential level of `self` enables per-user naming service lookups, or lookups that use the GSSAPI credentials of the user when connecting to the directory server. Currently the only GSSAPI mechanism supported in this model is Kerberos V5. Kerberos must be configured before you can use this credential level. See [kerberos\(5\)](#) for details.

More protection can be provided by means of access control, allowing the server to grant access for certain containers or entries. Access control is specified by Access Control Lists (ACLs) that are defined and stored in the LDAP server. The Access Control Lists on the LDAP server are called Access Control Instructions (ACIs) by the the SunOne Directory Server. Each ACL or ACI specifies one or more directory objects, for example, the `cn` attribute in a specific container, one or more clients to whom you grant or deny access, and one or more access rights that determine what the clients can do to or with the objects. Clients can be users or applications. Access rights can be specified as read and write, for example. Refer to the *Oracle Solaris Administration: Naming and Directory Services* regarding the restrictions on ACLs and ACIs when using LDAP as a naming repository.

A sample `nsswitch.conf(4)` file called `nsswitch.ldap` is provided in the `/etc` directory. This is copied to `/etc/nsswitch.conf` by the `ldapclient(1M)` utility. This file uses LDAP as a repository for the different databases in the `nsswitch.conf` file.

The following is a list of the user commands related to LDAP:

`idsconfig(1M)`      Prepares a SunOne Directory Server to be ready to support Solaris LDAP clients.

- `ldapaddent(1M)` Creates LDAP entries from corresponding /etc files.
- `ldapclient(1M)` Initializes LDAP clients, or generates a configuration profile to be stored in the directory.
- `ldaplist(1)` Lists the contents of the LDAP naming space.

|              |                                                                                    |                                                                                                                                                                                                    |
|--------------|------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Files</b> | <code>/var/ldap/ldap_client_cred</code><br><code>/var/ldap/ldap_client_file</code> | Files that contain the LDAP configuration of the client. Do not manually modify these files. Their content is not guaranteed to be human readable. Use <code>ldapclient(1M)</code> to update them. |
|              | <code>/etc/nsswitch.conf</code>                                                    | Configuration file for the name-service switch.                                                                                                                                                    |
|              | <code>/etc/nsswitch.ldap</code>                                                    | Sample configuration file for the name-service switch configured with LDAP and files.                                                                                                              |
|              | <code>/etc/pam.conf</code>                                                         | PAM framework configuration file.                                                                                                                                                                  |
|              | <code>/etc/pam.d/service</code>                                                    | Alternate PAM framework configuration files.                                                                                                                                                       |

**See Also** `ldaplist(1)`, `idsconfig(1M)`, `ldap_cachemgr(1M)`, `ldapaddent(1M)`, `ldapclient(1M)`, `nsswitch.conf(4)`, `pam.conf(4)`, `kerberos(5)`, `pam_authtok_check(5)`, `pam_authtok_get(5)`, `pam_authtok_store(5)`, `pam_dhkeys(5)`, `pam_ldap(5)`, `pam_passwd_auth(5)`, `pam_unix_account(5)`, `pam_unix_auth(5)`, `pam_unix_session(5)`

*Oracle Solaris Administration: Naming and Directory Services*

**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 [lfcompile64\(5\)](#) 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** 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.

**< aio.h >**

|                   |                     |
|-------------------|---------------------|
| struct aiocb      | struct aiocb64      |
| off_t aio_offset; | off64_t aio_offset; |

**< sys/dirent.h >**

|               |                 |
|---------------|-----------------|
| struct dirent | struct dirent64 |
| ino_t d_ino;  | ino64_t d_ino;  |
| off_t d_off;  | off64_t d_off;  |

**< sys/fcntl.h >**

|                |                  |
|----------------|------------------|
| struct flock   | struct flock64   |
| off_t l_start; | off64_t l_start; |
| off_t l_len;   | off64_t l_len;   |
| F_SETLK        | F_SETLK64        |
| F_SETLKW       | F_SETLKW64       |
| F_GETLK        | F_GETLK64        |
| F_FREESP       | F_FREESP64       |
| F_ALLOCSF      | F_ALLOCSF64      |
|                | O_LARGEFILE      |

**<sys/stdio.h>**

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>**

|                          |                            |
|--------------------------|----------------------------|
| <code>off_t;</code>      | <code>off64_t;</code>      |
| <code>ino_t;</code>      | <code>ino64_t;</code>      |
| <code>blkcnt_t;</code>   | <code>blkcnt64_t;</code>   |
| <code>fsblkcnt_t;</code> | <code>fsblkcnt64_t;</code> |
| <code>fsfilcnt_t;</code> | <code>fsfilcnt64_t;</code> |

**<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>**

|                                              |                                                  |
|----------------------------------------------|--------------------------------------------------|
| <code>int aio_cancel(...,</code>             | <code>int aio_cancel64(...,</code>               |
| <code>struct aiocb *);</code>                | <code>struct aiocb64 *);</code>                  |
| <code>int aio_error(</code>                  | <code>int aio_error64(</code>                    |
| <code>const struct aiocb *);</code>          | <code>const struct aiocb64 *);</code>            |
| <code>int aio_fsync(...,</code>              | <code>int aio_fsync64(...,</code>                |
| <code>struct aiocb *);</code>                | <code>struct aiocb64 *);</code>                  |
| <code>int aio_read(struct aiocb *);</code>   | <code>int aio_read64(struct aiocb64 *);</code>   |
| <code>int aio_return(struct aiocb *);</code> | <code>int aio_return64(struct aiocb64 *);</code> |
| <code>int aio_suspend(</code>                | <code>int aio_suspend64(</code>                  |

```

const struct aiocb *, ...);
int aio_waitn(aiocb_t*[],
...);
int aio_write(struct aiocb*);
int lio_listio(...,
const struct aiocb *, ...);

```

```

const struct aiocb64 *, ...);
int aio_waitn64(aiocb64_t*[],
...);
int aio_write64(struct aiocb64*);
int lio_listio64(...,
const struct aiocb64 *, ...);

```

### <dirent.h>

```

int alphasort(
const struct dirent **,
const struct dirent **)
struct dirent *readdir();
struct dirent *readdir_r();
int scandir(...,
struct dirent *(*[]),
int (*)(const struct dirent *),
int (*)(const struct dirent **,
const struct dirent **))

```

```

int alphasort64(
const struct dirent64 **,
const struct dirent64 **)
struct dirent64 *readdir64();
struct dirent64 *readdir64_r();
int scandir64(...,
struct dirent64 *(*[]),
int (*)(const struct dirent64 *),
int (*)(const struct dirent64 **,
const struct dirent64 **))

```

### <fcntl.h>

```

int attropen();
int creat();
int open();
int openat();
int posix_fadvise()
int posix_fallocate()

```

```

int attropen64();
int creat64();
int open64();
int openat64();
int posix_fadvise64()
int posix_fallocate64()

```

### <ftw.h>

```
int ftw(...,  
const struct stat *, ...);
```

```
int ftw64(...,  
const struct stat64 *, ...);
```

```
int nftw(..  
const struct stat *, ...);
```

```
int nftw64(...,  
const struct stat64 *, ...);
```

### <libgen.h>

```
char *copylist(..., off_t);
```

```
char *copylist64(..., off64_t);
```

### <stdio.h>

```
int fgetpos();  
FILE *fopen();  
FILE *freopen();  
int fseeko(..., off_t, ...);  
int fsetpos(...,  
const fpos_t *);  
off_t ftello();  
FILE *tmpfile();
```

```
int fgetpos64();  
FILE *fopen64();  
FILE *freopen64();  
int fseeko64(..., off64_t, ...);  
int fsetpos64(...,  
const fpos64_t *);  
off64_t ftello64();  
FILE *tmpfile64();
```

### <stdlib.h>

```
int mkstemp();
```

```
int mkstemp64();
```

### <sys/async.h>

```
int aioread(..., off_t, ...);  
int aiowrite(..., off_t, ...);
```

```
int aioread64(..., off64_t, ...);  
int aiowrite64(..., off64_t, ...);
```

### <sys/dirent.h>

```
int getdents(..., dirent);
```

```
int getdents64(..., dirent64);
```

### <sys/mman.h>

```
void mmap(..., off_t);
```

```
void mmap64(..., off64_t);
```

### <sys/resource.h>

```
int getrlimit(...,
struct rlimit *);
int setrlimit(...,
const struct rlimit *);
```

```
int getrlimit64(...,
struct rlimit64 *);
int setrlimit64(...,
const struct rlimit64 *);
```

### <sys/sendfile.h>

```
ssize_t sendfile(...,
off_t *, ...);
ssize_t sendfilev(..., const
struct sendfilevec *, ...);
```

```
ssize_t sendfile64(...,
off64_t *, ...);
ssize_t sendfilev64(..., const
struct sendfilevec64 *, ...);
```

### <sys/stat.h>

```
int fstat(..., struct stat *);
int fstatat(...,
struct stat *, int);
int lstat(..., struct stat *);
int stat(..., struct stat *);
```

```
int fstat64(..., struct stat64 *);
int fstatat64(...,
struct stat64 *, int);
int lstat64(..., struct stat64 *);
int stat64(..., struct stat64 *);
```

### <sys/statvfs.h>

```
int statvfs(...,
struct statvfs *);
int fstatvfs(...,
struct statvfs *);
```

```
int statvfs64(...,
struct statvfs64 *);
int fstatvfs64(...,
struct statvfs64 *);
```

#### <ucbinclude/stdio.h>

```
FILE *fopen()
FILE *freopen()
```

```
FILE *fopen64()
FILE *freopen64()
```

#### <ucbinclude/sys/dir.h>

```
int alphasort(
struct direct **,
struct direct **);
struct direct *readdir();
int scandir(...,
struct direct *(*[]), ...);
```

```
int alphasort64(
struct direct64 **,
struct direct64 **);
struct direct64 *readdir64();
int scandir64(...,
struct direct64 *(*[]), ...);
```

#### <unistd.h>

```
int lockf(..., off_t);
off_t lseek(..., off_t, ...);
int ftruncate(..., off_t);
ssize_t pread(..., off_t);
ssize_t pwrite(..., off_t);
int truncate(..., off_t);
```

```
int lockf64(..., off64_t);
off64_t lseek64(..., off64_t, ...);
int ftruncate64(..., off64_t);
ssize_t pread64(..., off64_t);
ssize_t pwrite64(..., off64_t);
int truncate64(..., off64_t);
```

**See Also** [lfcompile\(5\)](#), [lfcompile64\(5\)](#)

**Name** lfcompile – large file compilation environment for 32-bit applications

**Description** All 64-bit applications can manipulate large files by default. The methods described on this page allow 32-bit applications to manipulate large files.

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).

Each interface named `xxx()` that needs to access 64-bit entities to access large files maps to a `xxx64()` call in the resulting binary. All relevant data types are defined to be of correct size (for example, `off_t` has a typedef definition for a 64-bit entity).

An application compiled in this environment is able to use the `xxx()` source interfaces to access both large and small files, rather than having to explicitly utilize the transitional `xxx64()` interface calls to access large files. See the [lfcompile64\(5\)](#) manual page for information regarding the transitional compilation environment.

Applications can be compiled in the large file compilation environment by using the following methods:

- Use the [getconf\(1\)](#) utility with one or more of the arguments listed in the table below. This method is recommended for portable applications.

| 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                                                                 |

- Set the compile-time flag `_FILE_OFFSET_BITS` 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.

**Access to Additional Large File Interfaces**

The `fseek()` and `ftell()` functions *do not* map to functions named `fseek64()` and `ftell64()`; rather, the large file additions `fseeko()` and `ftello()`, have functionality identical to `fseek()` and `ftell()` and *do* map to the 64-bit functions `fseeko64()` and `ftello64()`. Applications wishing to access large files should use `fseeko()` and `ftello()` in place of `fseek()` and `ftell()`. See the [fseek\(3C\)](#) and [ftell\(3C\)](#) manual pages for information about `fseeko()` and `ftello()`.

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** Compile a program with a “large” `off_t` that uses `fseeko()`, `ftello()`, and `yacc`.

The following example compiles a program with a “large” `off_t` and 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** Compile 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_CFLAG`S a.c      \
    `getconf LFS_LDFLAG`S         \
    `getconf LFS_LIB`S            \
```

**EXAMPLE 3** Compile a program with a “default” `off_t` that uses `fseeko()` and `ftello()`.

```
$ c89 -D_LARGEFILE_SOURCE a.c
```

**See Also** [csh\(1\)](#), [getconf\(1\)](#), [ksh\(1\)](#), [yacc\(1\)](#), [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** 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.

**Name** lfcompile64 – transitional compilation environment

**Description** 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.

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).

The transitional compilation environment exports all the explicit 64-bit functions (`xxx64()`) and types in addition to all the regular functions (`xxx()`) and types. Both `xxx()` and `xxx64()` functions are available to the program source. A 32-bit application must use the `xxx64()` functions in order to access large files. See the [lf64\(5\)](#) manual page for a complete listing of the 64-bit transitional interfaces.

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 `xxx()` source interfaces to access both large and small files, rather than having to explicitly utilize the transitional `xxx64()` interface calls to access large files. See the [lfcompile\(5\)](#) manual page for more information regarding the large file compilation environment.

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.

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.

**Access to Additional Large File Interfaces** 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 [standards\(5\)](#)).

- Set the compile-time flag `_LARGEFILE64_SOURCE` to 1 before including any headers.
- Use the [getconf\(1\)](#) command with one or more of the following arguments:

| argument                     | purpose                                                                               |
|------------------------------|---------------------------------------------------------------------------------------|
| <code>LFS64_CFLAGS</code>    | obtain compilation flags necessary to enable the transitional compilation environment |
| <code>LFS64_LDFLAGS</code>   | obtain link editor options                                                            |
| <code>LFS64_LIBS</code>      | obtain link library names                                                             |
| <code>LFS64_LINTFLAGS</code> | 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 [csh\(1\)](#)), 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_LINTFLAG`S ... \
    `getconf LFS64_LIB`S
```

**See Also** [getconf\(1\)](#), [lseek\(2\)](#), [fopen\(3C\)](#), [lf64\(5\)](#), [standards\(5\)](#)

**Name** locale – subset of a user's environment that depends on language and cultural conventions

**Description** A locale 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:

|             |                                                                           |
|-------------|---------------------------------------------------------------------------|
| LC_CTYPE    | Character classification and case conversion.                             |
| LC_COLLATE  | Collation order.                                                          |
| LC_TIME     | Date and time formats.                                                    |
| LC_NUMERIC  | Numeric formatting.                                                       |
| LC_MONETARY | Monetary formatting.                                                      |
| LC_MESSAGES | Formats of informative and diagnostic messages and interactive responses. |

The standard utilities base their behavior on the current locale, as defined in the ENVIRONMENT VARIABLES 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 `setlocale(3C)`.

Locales other than those supplied by the implementation can be created by the application via the `localedef(1)` utility. The value that is used to specify a locale when using environment variables will be the string specified as the *name* operand to `localedef` when the locale was created. The strings “C” and “POSIX” are reserved as identifiers for the POSIX locale.

Applications can select the desired locale by invoking the `setlocale()` function with the appropriate value. If the function is invoked with an empty string, such as:

```
setlocale(LC_ALL, "");
```

the value of the corresponding environment variable is used. If the environment variable is unset or is set to the empty string, the `setlocale()` function sets the appropriate environment.

**Locale Definition** Locales can be described with the file format accepted by the `localedef` utility.

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.

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 LC\_. The category trailer consists of the string END, followed by one or more blank characters and the string used in the corresponding category header.

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>*, as 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.

Example:

```
c "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.

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.

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.

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.

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.

**upper** Define characters to be classified as upper-case letters.

In the POSIX locale, the 26 upper-case letters are included:

```
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, no character specified for the keywords `cntrl`, `digit`, `punct`, or `space` can be specified. The upper-case letters A to Z are automatically included in this class.

**lower** Define characters to be classified as lower-case letters. In the POSIX locale, the 26 lower-case letters are included:

```
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, no character specified for the keywords `cntrl`, `digit`, `punct`, or `space` can be specified. The lower-case letters a to z of the portable character set are automatically included in this class.

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <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><code>0 1 2 3 4 5 6 7 8 9</code></p> <p>are included.</p> <p>In a locale definition file, only the digits <code>0</code>, <code>1</code>, <code>2</code>, <code>3</code>, <code>4</code>, <code>5</code>, <code>6</code>, <code>7</code>, <code>8</code>, and <code>9</code> can be specified, and in contiguous ascending sequence by numerical value. The digits <code>0</code> to <code>9</code> 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>alnum</code> | <p>Define characters to be classified as letters and numeric digits. Only the characters specified for the <code>alpha</code> and <code>digit</code> keywords are specified. Characters specified for the keywords <code>alpha</code> and <code>digit</code> are automatically included in this class.</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> <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>                                                                                                                                                                                                                                                                                                                                                                          |

---

|        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| punct  | Define characters to be classified as punctuation characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|        | 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.                                                                                                                                                                                                                                                                                                                                                          |
|        | 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.                                                                                                                                                                                                                                                                         |
| graph  | Define characters to be classified as printable characters, not including the space character.                                                                                                                                                                                                                                                                                                                                                                                                                        |
|        | 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.                                                                                                                                                                                                                                                                                                                                 |
|        | 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.                                                                                                                                                                                             |
| print  | Define characters to be classified as printable characters, including the space character.                                                                                                                                                                                                                                                                                                                                                                                                                            |
|        | In the POSIX locale, all characters in class <code>graph</code> are included; no characters in class <code>cntrl</code> are included.                                                                                                                                                                                                                                                                                                                                                                                 |
|        | 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.                                                                                                                                                                       |
| xdigit | Define the characters to be classified as hexadecimal digits.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|        | In the POSIX locale, only:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|        | <pre>0 1 2 3 4 5 6 7 8 9 A B C D E F a b c d e f</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|        | are included.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|        | 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. |
|        | The definition of character class <code>xdigit</code> requires that the characters included in character class <code>digit</code> be included here also.                                                                                                                                                                                                                                                                                                                                                              |

|                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>blank</code>          | <p>Define characters to be classified as blank characters.</p> <p>In the POSIX locale, only the space and tab characters are included.</p> <p>In a locale definition file, the characters space and tab are automatically included in this class.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <code>charclass</code>      | <p>Define one or more locale-specific character class names as strings separated by semicolons. 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>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <code>charclass-name</code> | <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 <a href="#">wctype(3C)</a> function, in regular expression and shell pattern-matching bracket expressions, and by the <code>tr(1)</code> command.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <code>toupper</code>        | <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> <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>are mapped to the corresponding 26 upper-case characters:</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, 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 <code>lower</code> and <code>upper</code> can be specified. The lower-case letters <code>a</code> to <code>z</code>, and their corresponding upper-case letters <code>A</code> to <code>Z</code>, of the portable character set are automatically included in this mapping, but only when the <code>toupper</code> keyword is omitted from the locale definition.</p> |
| <code>tolower</code>        | <p>Define the mapping of upper-case letters to lower-case letters.</p> <p>In the POSIX locale, at a minimum, the 26 upper-case characters:</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>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

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 `wcsxfrm(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). 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, the pairs of collating elements are re-compared 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.

|                                |                                                                                                                                                                                 |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>copy</code>              | Specify the name of an existing locale which is used as the definition of this category. If this keyword is specified, no other keyword is specified.                           |
| <code>collating-element</code> | Define a collating-element symbol representing a multi-character collating element. This keyword is optional.                                                                   |
| <code>collating-symbol</code>  | Define a collating symbol for use in collation order statements. This keyword is optional.                                                                                      |
| <code>order_start</code>       | Define collation rules. This statement is followed by one or more collation order statements, assigning character collation values and collation weights to collating elements. |
| <code>order_end</code>         | 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.

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.

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 (;). 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:

- |                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>forward</code>  | Specifies that comparison operations for the weight level proceed from start of string towards the end of string.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <code>backward</code> | Specifies that comparison operations for the weight level proceed from end of string towards the beginning of string.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <code>position</code> | Specifies that comparison operations for the weight level will consider the relative position of elements in the strings not subject to <code>IGNORE</code> . The string containing an element not subject to <code>IGNORE</code> after the fewest collating elements subject to <code>IGNORE</code> from the start of the compare will collate first. If both strings contain a character not subject to <code>IGNORE</code> 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 <code>IGNORE</code> are considered in the same manner. |

The directives `forward` and `backward` are mutually exclusive.

Example:

```
order_start    forward;backward
```

If no operands are specified, a single `forward` operand is assumed.

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 use of the ellipsis symbol ties the definition to a specific coded character set and may preclude the definition from being portable between implementations.

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.

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.

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.

`int_curr_symbol` 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 standard. The fourth character is the character used to separate the international currency symbol from the monetary quantity.

`currency_symbol` The string used as the local currency symbol.

`mon_decimal_point` The operand is a string containing the symbol that is used as the decimal delimiter (radix character) in monetary formatted quantities.

`mon_thousands_sep` 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.

`mon_grouping` 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 `mon_grouping` keyword. Assuming that the value to be formatted is `123456789` and the `mon_thousands_sep` 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 `localeconv` function to accommodate this grouping.

<code>mon_grouping</code>	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	1234567898	"\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>	In an application conforming to the SUSv3 standard, an integer set to 1 if the <code>currency_symbol</code> precedes the value for a monetary quantity with a non-negative value, and set to 0 if the symbol succeeds the value.  In an application <i>not</i> conforming to the SUSv3 standard, an integer set to 1 if the <code>currency_symbol</code> or <code>int_currency_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>	In an application conforming to the SUSv3 standard, an integer set to 0 if no space separates the <code>currency_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.  In an application <i>not</i> conforming to the SUSv3 standard, 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.

n_cs_precedes	<p>In an application conforming to the SUSv3 standard, an integer set to 1 if the <code>currency_symbol</code> precedes the value for a monetary quantity with a negative value, and set to 0 if the symbol succeeds the value.</p> <p>In an application <i>not</i> conforming to the SUSv3 standard, an integer set to 1 if the <code>currency_symbol</code> or <code>int_currency_symbol</code> precedes the value for a monetary quantity with a negative value, and set to 0 if the symbol succeeds the value.</p>
n_sep_by_space	<p>In an application conforming to the SUSv3 standard, an integer set to 0 if no space separates the <code>currency_symbol</code> from the value for a monetary quantity with a 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.</p> <p>In an application <i>not</i> conforming to the SUSv3 standard, 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, 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>
p_sign_posn	<p>An integer set to a value indicating the positioning of the <code>positive_sign</code> for a monetary quantity with a non-negative value. The following integer values are recognized for both <code>p_sign_posn</code> and <code>n_sign_posn</code>:</p> <p>In an application conforming to the SUSv3 standard:</p> <ul style="list-style-type: none"> <li>0 Parentheses enclose the quantity and the <code>currency_symbol</code>.</li> <li>1 The sign string precedes the quantity and the <code>currency_symbol</code>.</li> <li>2 The sign string succeeds the quantity and the <code>currency_symbol</code>.</li> <li>3 The sign string precedes the <code>currency_symbol</code>.</li> <li>4 The sign string succeeds the <code>currency_symbol</code>.</li> </ul> <p>In an application <i>not</i> conforming to the SUSv3 standard:</p> <ul style="list-style-type: none"> <li>0 Parentheses enclose the quantity and the <code>currency_symbol</code> or <code>int_curr_symbol</code>.</li> <li>1 The sign string precedes the quantity and the <code>currency_symbol</code> or <code>int_curr_symbol</code>.</li> </ul>

	<ol style="list-style-type: none"><li>2 The sign string succeeds the quantity and the <code>currency_symbol</code> or <code>int_curr_symbol</code>.</li><li>3 The sign string precedes the <code>currency_symbol</code> or <code>int_curr_symbol</code>.</li><li>4 The sign string succeeds the <code>currency_symbol</code> or <code>int_curr_symbol</code>.</li></ol>
<code>n_sign_posn</code>	An integer set to a value indicating the positioning of the <code>negative_sign</code> for a negative formatted monetary quantity.
<code>int_p_cs_precedes</code>	An integer set to 1 if the <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>int_n_cs_precedes</code>	An integer set to 1 if the <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>int_p_sep_by_space</code>	An integer set to 0 if no space separates the <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>int_n_sep_by_space</code>	An integer set to 0 if no space separates the <code>int_curr_symbol</code> from the value for a monetary quantity with a 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>int_p_sign_posn</code>	<p>An integer set to a value indicating the positioning of the <code>positive_sign</code> for a positive monetary quantity formatted with the international format. The following integer values are recognized for <code>int_p_sign_posn</code> and <code>int_n_sign_posn</code>:</p> <ol style="list-style-type: none"><li>0 Parentheses enclose the quantity and the <code>int_curr_symbol</code>.</li><li>1 The sign string precedes the quantity and the <code>int_curr_symbol</code>.</li><li>2 The sign string precedes the quantity and the <code>int_curr_symbol</code>.</li><li>3 The sign string precedes the <code>int_curr_symbol</code>.</li><li>4 The sign string succeeds the <code>int_curr_symbol</code>.</li></ol>
<code>int_n_sign_posn</code>	An integer set to a value indicating the positioning of the <code>negative_sign</code> for a negative monetary quantity formatted with the international format.

The following table shows the result of various combinations:

		p_sep_by_space		
		2	1	0
p_cs_precedes=1	p_sign_posn=0	(\$1.25)	(\$1.25)	(\$1.25)
	p_sign_posn=1	+\$1.25	+\$1.25	+\$1.25
	p_sign_posn=2	\$1.25+	\$1.25+	\$1.25+
	p_sign_posn=3	+\$1.25	+\$1.25	+\$1.25
	p_sign_posn=4	+\$1.25	+\$1.25	+\$1.25
p_cs_precedes=0	p_sign_posn=0	(1.25 \$)	(1.25 \$)	(1.25\$)
	p_sign_posn=1	+1.25 \$	+1.25 \$	+1.25\$
	p_sign_posn=2	1.25\$ +	1.25 \$+	1.25\$+
	p_sign_posn=3	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 depicts 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
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
int_p_cs_precedes    -1
```

```
int_p_sep_by_space    -1
int_n_cs_precedes    -1
int_n_sep_by_space    -1
int_p_sign_posn      -1
int_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.

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 <code>-1</code> , then the size of the previous group (if any) will be repeatedly used for the remainder of the digits. If the last integer is <code>-1</code> , then no further grouping will be performed. The non-monetary numeric formatting definitions for the POSIX locale follow. The code listing depicts 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 **strptime()**, **wcsftime()**, and **strptime()** 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 \\, \a, \b, \f, \n, \r, \t, \v.
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 descriptors. In addition, the string can contain the escape sequences \\, \a, \b, \f, \n, \r, \t, \v.
d_fmt	Define the appropriate date representation, corresponding to the %x 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 \\, \a, \b, \f, \n, \r, \t, \v.
t_fmt	Define the appropriate time representation, corresponding to the %X 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 \\, \a, \b, \f, \n, \r, \t, \v.
am_pm	Define the appropriate representation of the <i>ante meridiem</i> and <i>post meridiem</i> strings, corresponding to the %p 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.
t_fmt_ampm	Define the appropriate time representation in the 12-hour clock format with am_pm, corresponding to the %r 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.
era	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:  <i>direction:offset:start_date:end_date:era_name:era_format</i>  according to the definitions below. There can be as many era description segments as are necessary to describe the different eras.  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.

	<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> .
	<i>offset</i>	The number of the year closest to the <i>start_date</i> in the era, corresponding to the %Eg and %Ey field descriptors.
	<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 –* or +*. The value –* indicates that the ending date is the beginning of time. The value +* 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.
era_d_fmt		Define the format of the date in alternative era notation, corresponding to the %Ex field descriptor.
era_t_fmt		Define the locale's appropriate alternative time format, corresponding to the %EX field descriptor.
era_d_t_fmt		Define the locale's appropriate alternative date and time format, corresponding to the %Ec field descriptor.
alt_digits		Define alternative symbols for digits, corresponding to the %0 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 %0 modifier indicates that the string corresponding to the value specified via 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 <code>&lt;langinfo.h&gt;</code> and used as arguments to the <code>nl_langinfo(3C)</code> function.
ABDAY_x		The abbreviated weekday names (for example Sun), where <i>x</i> is a number from 1 to 7.

DAY_ <i>x</i>	The full weekday names (for example Sunday), where <i>x</i> is a number from 1 to 7.
ABMON_ <i>x</i>	The abbreviated month names (for example Jan), where <i>x</i> is a number from 1 to 12.
MON_ <i>x</i>	The full month names (for example January), where <i>x</i> is a number from 1 to 12.
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 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:

*direction*:*offset*:*start\_date*:*end\_date*:*era\_name*:*era\_format*

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.

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.

*direction* 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*.

*offset* The number of the year closest to the *start\_date* in the era.

*start\_date* A date in the form *yyyy/mm/dd*, where *yyyy*, *mm*, and *dd* are the year, month and day numbers respectively of the start of the era. Years prior to AD 1 are represented as negative numbers.

*end\_date* The ending date of the era, in the same format as the *start\_date*, or one of the two special values, –\* or +\*. The

value `-*` indicates that the ending date is the beginning of 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 `%0` 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 `strptime(3C)`, `wcsftime(3C)`, and `strptime(3C)` functions.

localedef Keyword	langinfo Constant	Conversion Specifier
abday	ABDAY_x	%a
day	DAY_x	%A
abmon	ABMON_x	%b
mon	MON	%B
d_t_fmt	D_T_FMT	%c
date_fmt	DATE_FMT	%C
d_fmt	D_FMT	%x
t_fmt	T_FMT	%X
am_pm	AM_STR	%p
am_pm	PM_STR	%p
t_fmt_ampm	T_FMT_AMPM	%r
era	ERA	%EC, %Eg,

Localedef Keyword	Langinfo Constant	Conversion Specifier
		%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	%0

**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 %0d 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” The above example is for illustrative purposes only. The `%0` 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
<code>yesexpr</code>	YESEXPR	"^[yY]"
<code>noexpr</code>	NOEXPR	"^[nN]"
<code>yesstr</code>	YESSTR	"yes"
<code>nostr</code>	NOSTR	"no"

In an application conforming to the SUSv3 standard, the information on `yesstr` and `nostr` is not available.

**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\(3C\)](#), [wcsoll\(3C\)](#), [wcsftime\(3C\)](#), [wcsxfrm\(3C\)](#), [wctype\(3C\)](#), [attributes\(5\)](#), [charmap\(5\)](#), [extensions\(5\)](#), [regex\(5\)](#)

**Name** locale\_alias – locale name aliases and their corresponding canonical locale names

**Description** There are two sets of locale name aliases that are accepted and supported in the system:

1. Locale name aliases that are accepted and mapped to corresponding canonical locale names, if any, during locale selection process as specified in [setlocale\(3C\)](#) and message object or message catalog processing as specified in [gettext\(1\)](#), [catopen\(3C\)](#), and [gettext\(3C\)](#).

During the mapping process, the codeset name portion of the locale name aliases are normalized by extracting only alphanumeric characters with to-lower case conversions to have a better success ratio of possible mappings. As an example, with this normalization, a locale name alias such as AR\_AA.UTF-8 is normalized into AR\_AA.utf8 as a search domain value prior to actual comparisons to a set of pre-normalized locale alias names in internal mapping table to find the canonical locale name.

The supported locale name aliases in machine order are shown below. It is possible that the current system might not have the corresponding canonical locales due to either the locales are not currently installed in the system or not yet supported by Solaris.

Locale Name Alias	Canonical Locale Name
-----	-----
AR_AE	ar_AE.UTF-8
AR_AE.UTF-8	ar_AE.UTF-8
AR_BH	ar_BH.UTF-8
AR_BH.UTF-8	ar_BH.UTF-8
AR_DZ	ar_DZ.UTF-8
AR_DZ.UTF-8	ar_DZ.UTF-8
AR_EG	ar_EG.UTF-8
AR_EG.UTF-8	ar_EG.UTF-8
AR_JO	ar_JO.UTF-8
AR_JO.UTF-8	ar_JO.UTF-8
AR_KW	ar_KW.UTF-8
AR_KW.UTF-8	ar_KW.UTF-8
AR_MA	ar_MA.UTF-8
AR_MA.UTF-8	ar_MA.UTF-8
AR_OM	ar_OM.UTF-8
AR_OM.UTF-8	ar_OM.UTF-8
AR_QA	ar_QA.UTF-8
AR_QA.UTF-8	ar_QA.UTF-8
AR_SA	ar_SA.UTF-8
AR_SA.UTF-8	ar_SA.UTF-8
AR_TN	ar_TN.UTF-8
AR_TN.UTF-8	ar_TN.UTF-8
AR_YE	ar_YE.UTF-8
AR_YE.UTF-8	ar_YE.UTF-8
AS_IN	as_IN.UTF-8
AS_IN.UTF-8	as_IN.UTF-8
AZ_AZ	az_AZ.UTF-8

AZ_AZ.UTF-8	az_AZ.UTF-8
BE_BY	be_BY.UTF-8
BE_BY.UTF-8	be_BY.UTF-8
BG_BG	bg_BG.UTF-8
BG_BG.UTF-8	bg_BG.UTF-8
BN_IN	bn_IN.UTF-8
BN_IN.UTF-8	bn_IN.UTF-8
CA_ES	ca_ES.UTF-8
CA_ES.UTF-8	ca_ES.UTF-8
CA_ES.UTF-8@euro	ca_ES.UTF-8
CA_ES@euro	ca_ES.UTF-8
CS_CZ	cs_CZ.UTF-8
CS_CZ.UTF-8	cs_CZ.UTF-8
DA_DK	da_DK.UTF-8
DA_DK.UTF-8	da_DK.UTF-8
DE_AT	de_AT.UTF-8
DE_AT.UTF-8	de_AT.UTF-8
DE_AT.UTF-8@euro	de_AT.UTF-8
DE_AT@euro	de_AT.UTF-8
DE_CH	de_CH.UTF-8
DE_CH.UTF-8	de_CH.UTF-8
DE_DE	de_DE.UTF-8
DE_DE.UTF-8	de_DE.UTF-8
DE_DE.UTF-8@euro	de_DE.UTF-8
DE_DE@euro	de_DE.UTF-8
DE_LU	de_LU.UTF-8
DE_LU.UTF-8	de_LU.UTF-8
DE_LU.UTF-8@euro	de_LU.UTF-8
DE_LU@euro	de_LU.UTF-8
EL_GR	eġ_GR.UTF-8
EL_GR.UTF-8	eġ_GR.UTF-8
EN_AU	en_AU.UTF-8
EN_AU.UTF-8	en_AU.UTF-8
EN_CA	en_CA.UTF-8
EN_CA.UTF-8	en_CA.UTF-8
EN_GB	en_GB.UTF-8
EN_GB.UTF-8	en_GB.UTF-8
EN_GB.UTF-8@euro	en_GB.UTF-8
EN_GB@euro	en_GB.UTF-8
EN_HK	en_HK.UTF-8
EN_HK.UTF-8	en_HK.UTF-8
EN_IE	en_IE.UTF-8
EN_IE.UTF-8	en_IE.UTF-8
EN_IE.UTF-8@euro	en_IE.UTF-8
EN_IE@euro	en_IE.UTF-8
EN_IN	en_IN.UTF-8
EN_IN.UTF-8	en_IN.UTF-8
EN_NZ	en_NZ.UTF-8

EN_NZ.UTF-8	en_NZ.UTF-8
EN_PH	en_PH.UTF-8
EN_PH.UTF-8	en_PH.UTF-8
EN_SG	en_SG.UTF-8
EN_SG.UTF-8	en_SG.UTF-8
EN_US	en_US.UTF-8
EN_US.UTF-8	en_US.UTF-8
ES_AR	es_AR.UTF-8
ES_AR.UTF-8	es_AR.UTF-8
ES_BO	es_BO.UTF-8
ES_BO.UTF-8	es_BO.UTF-8
ES_CL	es_CL.UTF-8
ES_CL.UTF-8	es_CL.UTF-8
ES_CO	es_CO.UTF-8
ES_CO.UTF-8	es_CO.UTF-8
ES_CR	es_CR.UTF-8
ES_CR.UTF-8	es_CR.UTF-8
ES_DO	es_DO.UTF-8
ES_DO.UTF-8	es_DO.UTF-8
ES_EC	es_EC.UTF-8
ES_EC.UTF-8	es_EC.UTF-8
ES_ES	es_ES.UTF-8
ES_ES.UTF-8	es_ES.UTF-8
ES_ES.UTF-8@euro	es_ES.UTF-8
ES_ES@euro	es_ES.UTF-8
ES_GT	es_GT.UTF-8
ES_GT.UTF-8	es_GT.UTF-8
ES_HN	es_HN.UTF-8
ES_HN.UTF-8	es_HN.UTF-8
ES_MX	es_MX.UTF-8
ES_MX.UTF-8	es_MX.UTF-8
ES_NI	es_NI.UTF-8
ES_NI.UTF-8	es_NI.UTF-8
ES_PA	es_PA.UTF-8
ES_PA.UTF-8	es_PA.UTF-8
ES_PE	es_PE.UTF-8
ES_PE.UTF-8	es_PE.UTF-8
ES_PR	es_PR.UTF-8
ES_PR.UTF-8	es_PR.UTF-8
ES_PY	es_PY.UTF-8
ES_PY.UTF-8	es_PY.UTF-8
ES_SV	es_SV.UTF-8
ES_SV.UTF-8	es_SV.UTF-8
ES_US	es_US.UTF-8
ES_US.UTF-8	es_US.UTF-8
ES_UY	es_UY.UTF-8
ES_UY.UTF-8	es_UY.UTF-8
ES_VE	es_VE.UTF-8

ES_VE.UTF-8	es_VE.UTF-8
ET_EE	et_EE.UTF-8
ET_EE.UTF-8	et_EE.UTF-8
FI_FI	fi_FI.UTF-8
FI_FI.UTF-8	fi_FI.UTF-8
FI_FI.UTF-8@euro	fi_FI.UTF-8
FI_FI@euro	fi_FI.UTF-8
FR_BE	fr_BE.UTF-8
FR_BE.UTF-8	fr_BE.UTF-8
FR_BE.UTF-8@euro	fr_BE.UTF-8
FR_BE@euro	fr_BE.UTF-8
FR_CA	fr_CA.UTF-8
FR_CA.UTF-8	fr_CA.UTF-8
FR_CH	fr_CH.UTF-8
FR_CH.UTF-8	fr_CH.UTF-8
FR_FR	fr_FR.UTF-8
FR_FR.UTF-8	fr_FR.UTF-8
FR_FR.UTF-8@euro	fr_FR.UTF-8
FR_FR@euro	fr_FR.UTF-8
FR_LU	fr_LU.UTF-8
FR_LU.UTF-8	fr_LU.UTF-8
FR_LU.UTF-8@euro	fr_LU.UTF-8
FR_LU@euro	fr_LU.UTF-8
GU_IN	gu_IN.UTF-8
GU_IN.UTF-8	gu_IN.UTF-8
HE_IL	he_IL.UTF-8
HE_IL.UTF-8	he_IL.UTF-8
HI_IN	hi_IN.UTF-8
HI_IN.UTF-8	hi_IN.UTF-8
HR_HR	hr_HR.UTF-8
HR_HR.UTF-8	hr_HR.UTF-8
HU_HU	hu_HU.UTF-8
HU_HU.UTF-8	hu_HU.UTF-8
ID_ID	id_ID.UTF-8
ID_ID.UTF-8	id_ID.UTF-8
IS_IS	is_IS.UTF-8
IS_IS.UTF-8	is_IS.UTF-8
IT_CH	it_CH.UTF-8
IT_CH.UTF-8	it_CH.UTF-8
IT_IT	it_IT.UTF-8
IT_IT.UTF-8	it_IT.UTF-8
IT_IT.UTF-8@euro	it_IT.UTF-8
IT_IT@euro	it_IT.UTF-8
JA_JP	ja_JP.UTF-8
JA_JP.UTF-8	ja_JP.UTF-8
KK_KZ	kk_KZ.UTF-8
KK_KZ.UTF-8	kk_KZ.UTF-8
KN_IN	kn_IN.UTF-8

KN_IN.UTF-8	kn_IN.UTF-8
KO_KR	ko_KR.UTF-8
KO_KR.UTF-8	ko_KR.UTF-8
LT_LT	lt_LT.UTF-8
LT_LT.UTF-8	lt_LT.UTF-8
LV_LV	lv_LV.UTF-8
LV_LV.UTF-8	lv_LV.UTF-8
MK_MK	mk_MK.UTF-8
MK_MK.UTF-8	mk_MK.UTF-8
ML_IN	ml_IN.UTF-8
ML_IN.UTF-8	ml_IN.UTF-8
MR_IN	mr_IN.UTF-8
MR_IN.UTF-8	mr_IN.UTF-8
MS_MY	ms_MY.UTF-8
MS_MY.UTF-8	ms_MY.UTF-8
MT_MT	mt_MT.UTF-8
MT_MT.UTF-8	mt_MT.UTF-8
NL_BE	nl_BE.UTF-8
NL_BE.UTF-8	nl_BE.UTF-8
NL_BE.UTF-8@euro	nl_BE.UTF-8
NL_BE@euro	nl_BE.UTF-8
NL_NL	nl_NL.UTF-8
NL_NL.UTF-8	nl_NL.UTF-8
NL_NL.UTF-8@euro	nl_NL.UTF-8
NL_NL@euro	nl_NL.UTF-8
OR_IN	or_IN.UTF-8
OR_IN.UTF-8	or_IN.UTF-8
PA_IN	pa_IN.UTF-8
PA_IN.UTF-8	pa_IN.UTF-8
PL_PL	pl_PL.UTF-8
PL_PL.UTF-8	pl_PL.UTF-8
PT_BR	pt_BR.UTF-8
PT_BR.UTF-8	pt_BR.UTF-8
PT_PT	pt_PT.UTF-8
PT_PT.UTF-8	pt_PT.UTF-8
PT_PT.UTF-8@euro	pt_PT.UTF-8
PT_PT@euro	pt_PT.UTF-8
RO_RO	ro_RO.UTF-8
RO_RO.UTF-8	ro_RO.UTF-8
RU_RU	ru_RU.UTF-8
RU_RU.UTF-8	ru_RU.UTF-8
SK_SK	sk_SK.UTF-8
SK_SK.UTF-8	sk_SK.UTF-8
SL_SI	sl_SI.UTF-8
SL_SI.UTF-8	sl_SI.UTF-8
SQ_AL	sq_AL.UTF-8
SQ_AL.UTF-8	sq_AL.UTF-8
SV_SE	sv_SE.UTF-8

SV_SE.UTF-8	sv_SE.UTF-8
TA_IN	ta_IN.UTF-8
TA_IN.UTF-8	ta_IN.UTF-8
TE_IN	te_IN.UTF-8
TE_IN.UTF-8	te_IN.UTF-8
TH_TH	th_TH.UTF-8
TH_TH.UTF-8	th_TH.UTF-8
TR_TR	tr_TR.UTF-8
TR_TR.UTF-8	tr_TR.UTF-8
UK_UA	uk_UA.UTF-8
UK_UA.UTF-8	uk_UA.UTF-8
VI_VN	vi_VN.UTF-8
VI_VN.UTF-8	vi_VN.UTF-8
ZH_CN	zh_CN.UTF-8
ZH_CN.UTF-8	zh_CN.UTF-8
ZH_HK	zh_HK.UTF-8
ZH_HK.UTF-8	zh_HK.UTF-8
ZH_SG	zh_SG.UTF-8
ZH_SG.UTF-8	zh_SG.UTF-8
ZH_TW	zh_TW.UTF-8
ZH_TW.UTF-8	zh_TW.UTF-8
Zh_CN	zh_CN.GB18030
Zh_CN.GB18030	zh_CN.GB18030
Zh_TW	zh_TW.BIG5
Zh_TW.big5	zh_TW.BIG5
af_ZA.utf8	af_ZA.UTF-8
ar	ar_EG.ISO8859-6
ar_AE.utf8	ar_AE.UTF-8
ar_BH.utf8	ar_BH.UTF-8
ar_DZ.utf8	ar_DZ.UTF-8
ar_EG	ar_EG.ISO8859-6
ar_EG.iso88596	ar_EG.ISO8859-6
ar_EG.utf8	ar_EG.UTF-8
ar_IQ.utf8	ar_IQ.UTF-8
ar_JO.utf8	ar_JO.UTF-8
ar_KW.utf8	ar_KW.UTF-8
ar_LY.utf8	ar_LY.UTF-8
ar_MA.utf8	ar_MA.UTF-8
ar_OM.utf8	ar_OM.UTF-8
ar_QA.utf8	ar_QA.UTF-8
ar_SA.utf8	ar_SA.UTF-8
ar_TN.utf8	ar_TN.UTF-8
ar_YE.utf8	ar_YE.UTF-8
as_IN.utf8	as_IN.UTF-8
az_AZ.utf8	az_AZ.UTF-8
be_BY.utf8	be_BY.UTF-8
bg_BG	bg_BG.ISO8859-5
bg_BG.utf8	bg_BG.UTF-8

bn_IN	bn_IN.UTF-8
bn_IN.utf8	bn_IN.UTF-8
bs_BA	bs_BA.ISO8859-2
bs_BA.iso88592	bs_BA.ISO8859-2
bs_BA.utf8	bs_BA.UTF-8
ca	ca_ES.ISO8859-1
ca_ES	ca_ES.ISO8859-1
ca_ES.8859-15	ca_ES.ISO8859-15
ca_ES.8859-15@euro	ca_ES.ISO8859-15
ca_ES.ISO8859-15@euro	ca_ES.ISO8859-15
ca_ES.iso88591	ca_ES.ISO8859-1
ca_ES.utf8	ca_ES.UTF-8
ca_ES@euro	ca_ES.ISO8859-15
cs	cs_CZ.ISO8859-2
cs_CZ	cs_CZ.ISO8859-2
cs_CZ.iso88592	cs_CZ.ISO8859-2
cs_CZ.utf8	cs_CZ.UTF-8
da	da_DK.ISO8859-1
da.ISO8859-15	da_DK.ISO8859-15
da_DK	da_DK.ISO8859-1
da_DK.8859-15	da_DK.ISO8859-15
da_DK.iso88591	da_DK.ISO8859-1
da_DK.iso885915	da_DK.ISO8859-15
da_DK.iso885915@euro	da_DK.ISO8859-15
da_DK.utf8	da_DK.UTF-8
de	de_DE.ISO8859-1
de.ISO8859-15	de_DE.ISO8859-15
de.UTF-8	de_DE.UTF-8
de_AT	de_AT.ISO8859-1
de_AT.8859-15	de_AT.ISO8859-15
de_AT.8859-15@euro	de_AT.ISO8859-15
de_AT.ISO8859-15@euro	de_AT.ISO8859-15
de_AT.iso88591	de_AT.ISO8859-1
de_AT.utf8	de_AT.UTF-8
de_AT@euro	de_AT.ISO8859-15
de_BE.utf8	de_BE.UTF-8
de_CH	de_CH.ISO8859-1
de_CH.iso88591	de_CH.ISO8859-1
de_CH.utf8	de_CH.UTF-8
de_DE	de_DE.ISO8859-1
de_DE.8859-15	de_DE.ISO8859-15
de_DE.8859-15@euro	de_DE.ISO8859-15
de_DE.ISO8859-15@euro	de_DE.ISO8859-15
de_DE.UTF-8@euro	de_DE.UTF-8
de_DE.iso88591	de_DE.ISO8859-1
de_DE.utf8	de_DE.UTF-8
de_DE@euro	de_DE.ISO8859-15
de_LU.utf8	de_LU.UTF-8

eł	eł_GR.ISO8859-7
eł.UTF-8	eł_CY.UTF-8
eł.sun_eu_greek	eł_GR.ISO8859-7
eł_CY.utf8	eł_CY.UTF-8
eł_GR	eł_GR.ISO8859-7
eł_GR.ISO8859-7@euro	eł_GR.ISO8859-7
eł_GR.iso88597	eł_GR.ISO8859-7
eł_GR.utf8	eł_GR.UTF-8
en_AU	en_AU.ISO8859-1
en_AU.iso88591	en_AU.ISO8859-1
en_AU.utf8	en_AU.UTF-8
en_BW.utf8	en_BW.UTF-8
en_CA	en_CA.ISO8859-1
en_CA.iso88591	en_CA.ISO8859-1
en_CA.utf8	en_CA.UTF-8
en_GB	en_GB.ISO8859-1
en_GB.8859-15@euro	en_GB.ISO8859-15
en_GB.IBM-1252	en_GB.ANSI1252
en_GB.iso88591	en_GB.ISO8859-1
en_GB.iso885915	en_GB.ISO8859-15
en_GB.iso885915@euro	en_GB.ISO8859-15
en_GB.utf8	en_GB.UTF-8
en_HK.utf8	en_HK.UTF-8
en_IE	en_IE.ISO8859-1
en_IE.8859-15	en_IE.ISO8859-15
en_IE.8859-15@euro	en_IE.ISO8859-15
en_IE.ISO8859-15@euro	en_IE.ISO8859-15
en_IE.iso88591	en_IE.ISO8859-1
en_IE.utf8	en_IE.UTF-8
en_IE@euro	en_IE.ISO8859-15
en_IN.utf8	en_IN.UTF-8
en_NZ	en_NZ.ISO8859-1
en_NZ.iso88591	en_NZ.ISO8859-1
en_NZ.utf8	en_NZ.UTF-8
en_PH.utf8	en_PH.UTF-8
en_SG.utf8	en_SG.UTF-8
en_US	en_US.ISO8859-1
en_US.8859-15	en_US.ISO8859-15
en_US.iso88591	en_US.ISO8859-1
en_US.iso885915	en_US.ISO8859-15
en_US.utf8	en_US.UTF-8
en_ZW.utf8	en_ZW.UTF-8
es	es_ES.ISO8859-1
es.ISO8859-15	es_ES.ISO8859-15
es.UTF-8	es_ES.UTF-8
es_AR	es_AR.ISO8859-1
es_AR.iso88591	es_AR.ISO8859-1
es_AR.utf8	es_AR.UTF-8

es_BO	es_BO.ISO8859-1
es_BO.iso88591	es_BO.ISO8859-1
es_BO.utf8	es_BO.UTF-8
es_CL	es_CL.ISO8859-1
es_CL.iso88591	es_CL.ISO8859-1
es_CL.utf8	es_CL.UTF-8
es_CO	es_CO.ISO8859-1
es_CO.iso88591	es_CO.ISO8859-1
es_CO.utf8	es_CO.UTF-8
es_CR	es_CR.ISO8859-1
es_CR.iso88591	es_CR.ISO8859-1
es_CR.utf8	es_CR.UTF-8
es_DO.utf8	es_DO.UTF-8
es_EC	es_EC.ISO8859-1
es_EC.iso88591	es_EC.ISO8859-1
es_EC.utf8	es_EC.UTF-8
es_ES	es_ES.ISO8859-1
es_ES.8859-15	es_ES.ISO8859-15
es_ES.8859-15@euro	es_ES.ISO8859-15
es_ES.ISO8859-15@euro	es_ES.ISO8859-15
es_ES.UTF-8@euro	es_ES.UTF-8
es_ES.iso88591	es_ES.ISO8859-1
es_ES.utf8	es_ES.UTF-8
es_ES@euro	es_ES.ISO8859-15
es_GT	es_GT.ISO8859-1
es_GT.iso88591	es_GT.ISO8859-1
es_GT.utf8	es_GT.UTF-8
es_HN.utf8	es_HN.UTF-8
es_MX	es_MX.ISO8859-1
es_MX.iso88591	es_MX.ISO8859-1
es_NI.utf8	es_NI.UTF-8
es_PA	es_PA.ISO8859-1
es_PA.iso88591	es_PA.ISO8859-1
es_PA.utf8	es_PA.UTF-8
es_PE	es_PE.ISO8859-1
es_PE.iso88591	es_PE.ISO8859-1
es_PE.utf8	es_PE.UTF-8
es_PR.utf8	es_PR.UTF-8
es_PY	es_PY.ISO8859-1
es_PY.iso88591	es_PY.ISO8859-1
es_PY.utf8	es_PY.UTF-8
es_SV	es_SV.ISO8859-1
es_SV.iso88591	es_SV.ISO8859-1
es_SV.utf8	es_SV.UTF-8
es_US.utf8	es_US.UTF-8
es_UY	es_UY.ISO8859-1
es_UY.8859-15	es_UY.ISO8859-15
es_UY.iso88591	es_UY.ISO8859-1

es_UY.utf8	es_UY.UTF-8
es_VE	es_VE.ISO8859-1
es_VE.iso88591	es_VE.ISO8859-1
es_VE.utf8	es_VE.UTF-8
et	et_EE.ISO8859-15
et_EE	et_EE.ISO8859-15
et_EE.iso885915	et_EE.ISO8859-15
et_EE.utf8	et_EE.UTF-8
fi	fi_FI.ISO8859-1
fi.ISO8859-15	fi_FI.ISO8859-15
fi_FI	fi_FI.ISO8859-1
fi_FI.8859-15	fi_FI.ISO8859-15
fi_FI.8859-15@euro	fi_FI.ISO8859-15
fi_FI.ISO8859-15@euro	fi_FI.ISO8859-15
fi_FI.iso88591	fi_FI.ISO8859-1
fi_FI.utf8	fi_FI.UTF-8
fi_FI@euro	fi_FI.ISO8859-15
fr	fr_FR.ISO8859-1
fr.ISO8859-15	fr_FR.ISO8859-15
fr.UTF-8	fr_FR.UTF-8
fr_BE	fr_BE.ISO8859-1
fr_BE.8859-15	fr_BE.ISO8859-15
fr_BE.8859-15@euro	fr_BE.ISO8859-15
fr_BE.ISO8859-15@euro	fr_BE.ISO8859-15
fr_BE.UTF-8@euro	fr_BE.UTF-8
fr_BE.iso88591	fr_BE.ISO8859-1
fr_BE.utf8	fr_BE.UTF-8
fr_BE@euro	fr_BE.ISO8859-15
fr_CA	fr_CA.ISO8859-1
fr_CA.iso88591	fr_CA.ISO8859-1
fr_CA.utf8	fr_CA.UTF-8
fr_CH	fr_CH.ISO8859-1
fr_CH.iso88591	fr_CH.ISO8859-1
fr_CH.utf8	fr_CH.UTF-8
fr_FR	fr_FR.ISO8859-1
fr_FR.8859-15	fr_FR.ISO8859-15
fr_FR.8859-15@euro	fr_FR.ISO8859-15
fr_FR.ISO8859-15@euro	fr_FR.ISO8859-15
fr_FR.UTF-8@euro	fr_FR.UTF-8
fr_FR.iso88591	fr_FR.ISO8859-1
fr_FR.utf8	fr_FR.UTF-8
fr_FR@euro	fr_FR.ISO8859-15
fr_LU.utf8	fr_LU.UTF-8
fr_LU@euro	fr_LU.ISO8859-15
gu_IN	gu_IN.UTF-8
gu_IN.utf8	gu_IN.UTF-8
he	he_IL.ISO8859-8
he_IL	he_IL.ISO8859-8

he_IL.iso88598	he_IL.ISO8859-8
he_IL.utf8	he_IL.UTF-8
hi_IN	hi_IN.UTF-8
hi_IN.utf8	hi_IN.UTF-8
hr_HR	hr_HR.ISO8859-2
hr_HR.iso88592	hr_HR.ISO8859-2
hr_HR.utf8	hr_HR.UTF-8
hu	hu_HU.ISO8859-2
hu_HU	hu_HU.ISO8859-2
hu_HU.iso88592	hu_HU.ISO8859-2
hu_HU.utf8	hu_HU.UTF-8
hy_AM	hy_AM.UTF-8
hy_AM.utf8	hy_AM.UTF-8
id_ID.utf8	id_ID.UTF-8
is_IS	is_IS.ISO8859-1
is_IS.iso88591	is_IS.ISO8859-1
is_IS.utf8	is_IS.UTF-8
it	it_IT.ISO8859-1
it.ISO8859-15	it_IT.ISO8859-15
it.UTF-8	it_IT.UTF-8
it_CH.utf8	it_CH.UTF-8
it_IT	it_IT.ISO8859-1
it_IT.8859-15	it_IT.ISO8859-15
it_IT.8859-15@euro	it_IT.ISO8859-15
it_IT.ISO8859-15@euro	it_IT.ISO8859-15
it_IT.UTF-8@euro	it_IT.UTF-8
it_IT.iso88591	it_IT.ISO8859-1
it_IT.utf8	it_IT.UTF-8
it_IT@euro	it_IT.ISO8859-15
ja	ja_JP.eucJP
ja_JP	ja_JP.eucJP
ja_JP.IBM-eucJP	ja_JP.eucJP
ja_JP.eucjp	ja_JP.eucJP
ja_JP.utf8	ja_JP.UTF-8
ka_GE.utf8	ka_GE.UTF-8
kk_KZ.utf8	kk_KZ.UTF-8
kn_IN	kn_IN.UTF-8
kn_IN.utf8	kn_IN.UTF-8
ko	ko_KR.EUC
ko.UTF-8	ko_KR.UTF-8
ko_KR	ko_KR.EUC
ko_KR.IBM-eucKR	ko_KR.EUC
ko_KR.euckr	ko_KR.EUC
ko_KR.utf8	ko_KR.UTF-8
ku_TR.utf8	ku_TR.UTF-8
ky_KG	ky_KG.UTF-8
ky_KG.utf8	ky_KG.UTF-8
lt	lt_LT.ISO8859-13

lt_LT	lt_LT.ISO8859-13
lt_LT.iso885913	lt_LT.ISO8859-13
lt_LT.utf8	lt_LT.UTF-8
lv	lv_LV.ISO8859-13
lv_LV	lv_LV.ISO8859-13
lv_LV.iso885913	lv_LV.ISO8859-13
lv_LV.utf8	lv_LV.UTF-8
mk_MK	mk_MK.ISO8859-5
mk_MK.iso88595	mk_MK.ISO8859-5
mk_MK.utf8	mk_MK.UTF-8
ml_IN	ml_IN.UTF-8
ml_IN.utf8	ml_IN.UTF-8
mr_IN	mr_IN.UTF-8
mr_IN.utf8	mr_IN.UTF-8
ms_MY.utf8	ms_MY.UTF-8
mt_MT.utf8	mt_MT.UTF-8
nb_NO	nb_NO.ISO8859-1
nb_NO.iso88591	nb_NO.ISO8859-1
nb_NO.utf8	nb_NO.UTF-8
nĳ	nĳ_NL.ISO8859-1
nĳ.ISO8859-15	nĳ_NL.ISO8859-15
nĳ_BE	nĳ_BE.ISO8859-1
nĳ_BE.8859-15	nĳ_BE.ISO8859-15
nĳ_BE.8859-15@euro	nĳ_BE.ISO8859-15
nĳ_BE.ISO8859-15@euro	nĳ_BE.ISO8859-15
nĳ_BE.iso88591	nĳ_BE.ISO8859-1
nĳ_BE.utf8	nĳ_BE.UTF-8
nĳ_BE@euro	nĳ_BE.ISO8859-15
nĳ_NL	nĳ_NL.ISO8859-1
nĳ_NL.8859-15	nĳ_NL.ISO8859-15
nĳ_NL.8859-15@euro	nĳ_NL.ISO8859-15
nĳ_NL.ISO8859-15@euro	nĳ_NL.ISO8859-15
nĳ_NL.iso88591	nĳ_NL.ISO8859-1
nĳ_NL.utf8	nĳ_NL.UTF-8
nĳ_NL@euro	nĳ_NL.ISO8859-15
nn_NO	nn_NO.ISO8859-1
nn_NO.iso88591	nn_NO.ISO8859-1
nn_NO.utf8	nn_NO.UTF-8
no	nb_NO.ISO8859-1
no_NO	nb_NO.ISO8859-1
no_NO.ISO8859-1@bokmal	nb_NO.ISO8859-1
no_NO.ISO8859-1@nynorsk	nn_NO.ISO8859-1
no_NY	nn_NO.ISO8859-1
or_IN	or_IN.UTF-8
or_IN.utf8	or_IN.UTF-8
pa_IN	pa_IN.UTF-8
pa_IN.utf8	pa_IN.UTF-8
pĳ	pĳ_PL.ISO8859-2

pl.UTF-8	pl_PL.UTF-8
pl_PL	pl_PL.ISO8859-2
pl_PL.iso88592	pl_PL.ISO8859-2
pl_PL.utf8	pl_PL.UTF-8
pt	pt_PT.ISO8859-1
pt.ISO8859-15	pt_PT.ISO8859-15
pt_BR	pt_BR.ISO8859-1
pt_BR.iso88591	pt_BR.ISO8859-1
pt_BR.utf8	pt_BR.UTF-8
pt_PT	pt_PT.ISO8859-1
pt_PT.8859-15	pt_PT.ISO8859-15
pt_PT.8859-15@euro	pt_PT.ISO8859-15
pt_PT.ISO8859-15@euro	pt_PT.ISO8859-15
pt_PT.iso88591	pt_PT.ISO8859-1
pt_PT.utf8	pt_PT.UTF-8
pt_PT@euro	pt_PT.ISO8859-15
ro_RO	ro_RO.ISO8859-2
ro_RO.iso88592	ro_RO.ISO8859-2
ro_RO.utf8	ro_RO.UTF-8
ru	ru_RU.ISO8859-5
ru.UTF-8	ru_RU.UTF-8
ru.koi8-r	ru_RU.KOI8-R
ru_RU	ru_RU.ISO8859-5
ru_RU.iso88595	ru_RU.ISO8859-5
ru_RU.koi8r	ru_RU.KOI8-R
ru_RU.utf8	ru_RU.UTF-8
ru_UA.utf8	ru_UA.UTF-8
sh	bs_BA.ISO8859-2
sh_BA	bs_BA.ISO8859-2
sh_BA.ISO8859-2@bosnia	bs_BA.ISO8859-2
sh_BA.UTF-8	bs_BA.UTF-8
sk_SK	sk_SK.ISO8859-2
sk_SK.iso88592	sk_SK.ISO8859-2
sk_SK.utf8	sk_SK.UTF-8
sl_SI	sl_SI.ISO8859-2
sl_SI.iso88592	sl_SI.ISO8859-2
sl_SI.utf8	sl_SI.UTF-8
sq_AL	sq_AL.ISO8859-2
sq_AL.utf8	sq_AL.UTF-8
sr_CS.UTF-8	sr_RS.UTF-8
sr_CS.iso88595	sr_CS.ISO8859-5
sr_ME	sr_ME.UTF-8
sr_ME.utf8	sr_ME.UTF-8
sr_RS	sr_RS.UTF-8
sr_RS.utf8	sr_RS.UTF-8
sv	sv_SE.ISO8859-1
sv.ISO8859-15	sv_SE.ISO8859-15
sv.UTF-8	sv_SE.UTF-8

sv_SE	sv_SE.ISO8859-1
sv_SE.8859-15	sv_SE.ISO8859-15
sv_SE.iso88591	sv_SE.ISO8859-1
sv_SE.iso885915	sv_SE.ISO8859-15
sv_SE.iso885915@euro	sv_SE.ISO8859-15
sv_SE.utf8	sv_SE.UTF-8
ta_IN	ta_IN.UTF-8
ta_IN.utf8	ta_IN.UTF-8
te_IN	te_IN.UTF-8
te_IN.utf8	te_IN.UTF-8
th	th_TH.TIS620
th_TH	th_TH.TIS620
th_TH.ISO8859-11	th_TH.TIS620
th_TH.TIS-620	th_TH.TIS620
th_TH.utf8	th_TH.UTF-8
tr	tr_TR.ISO8859-9
tr_TR	tr_TR.ISO8859-9
tr_TR.iso88599	tr_TR.ISO8859-9
tr_TR.utf8	tr_TR.UTF-8
uk_UA.utf8	uk_UA.UTF-8
vi_VN	vi_VN.UTF-8
vi_VN.utf8	vi_VN.UTF-8
zh	zh_CN.EUC
zh.GBK	zh_CN.GBK
zh.UTF-8	zh_CN.UTF-8
zh_CN	zh_CN.EUC
zh_CN.IBM-eucCN	zh_CN.EUC
zh_CN.gb18030	zh_CN.GB18030
zh_CN.gb2312	zh_CN.EUC
zh_CN.gbk	zh_CN.GBK
zh_CN.utf8	zh_CN.UTF-8
zh_HK	zh_HK.BIG5HK
zh_HK.big5hkscs	zh_HK.BIG5HK
zh_HK.utf8	zh_HK.UTF-8
zh_SG.utf8	zh_SG.UTF-8
zh_TW	zh_TW.EUC
zh_TW.IBM-eucTW	zh_TW.EUC
zh_TW.big5	zh_TW.BIG5
zh_TW.euctw	zh_TW.EUC
zh_TW.utf8	zh_TW.UTF-8

2. Obsoleted Solaris locale names that are additionally checked against to find message object or message catalog files as specified in [gettext\(1\)](#), [catopen\(3C\)](#), and [gettext\(3C\)](#)

When the current locale is one of the listed canonical locales and there is no matching message object or message catalog file for the current running program to open with using the current locale name, the messaging functions additionally check on the existence of the message object or the message catalog file to open by utilizing the additional locale names shown at below as aliases:

Canonical Locale Name	Additional Locale Names Checked
-----	-----
ar_EG.ISO8859-6	ar
bg_BG.ISO8859-5	bg_BG
bs_BA.ISO8859-2	sh, sh_BA, sh_BA.ISO8859-2@bosnia
bs_BA.UTF-8	sh_BA.UTF-8
ca_ES.ISO8859-1	ca, ca_ES
ca_ES.ISO8859-15	ca_ES.ISO8859-15@euro
cs_CZ.ISO8859-2	cs, cs_CZ
da_DK.ISO8859-1	da, da_DK
da_DK.ISO8859-15	da.ISO8859-15
de_AT.ISO8859-1	de_AT
de_AT.ISO8859-15	de_AT.ISO8859-15@euro
de_CH.ISO8859-1	de_CH
de_DE.ISO8859-1	de, de_DE
de_DE.ISO8859-15	de.ISO8859-15, de_DE.ISO8859-15@euro
de_DE.UTF-8	de.UTF-8, de_DE.UTF-8@euro
el_CY.UTF-8	el.UTF-8
el_GR.ISO8859-7	el, el.sun_eu_greek, el_GR, el_GR.ISO8859-7@euro
en_AU.ISO8859-1	en_AU
en_CA.ISO8859-1	en_CA
en_GB.ISO8859-1	en_GB
en_IE.ISO8859-1	en_IE
en_IE.ISO8859-15	en_IE.ISO8859-15@euro
en_NZ.ISO8859-1	en_NZ
en_US.ISO8859-1	en_US
es_AR.ISO8859-1	es_AR
es_BO.ISO8859-1	es_BO
es_CL.ISO8859-1	es_CL
es_CO.ISO8859-1	es_CO
es_CR.ISO8859-1	es_CR
es_EC.ISO8859-1	es_EC
es_ES.ISO8859-1	es, es_ES
es_ES.ISO8859-15	es.ISO8859-15, es_ES.ISO8859-15@euro
es_ES.UTF-8	es.UTF-8, es_ES.UTF-8@euro
es_GT.ISO8859-1	es_GT
es_MX.ISO8859-1	es_MX
es_NI.ISO8859-1	es_NI
es_PA.ISO8859-1	es_PA
es_PE.ISO8859-1	es_PE
es_PY.ISO8859-1	es_PY
es_SV.ISO8859-1	es_SV
es_UY.ISO8859-1	es_UY
es_VE.ISO8859-1	es_VE
et_EE.ISO8859-15	et, et_EE
fi_FI.ISO8859-1	fi, fi_FI

fi_FI.ISO8859-15	fi.ISO8859-15, fi_FI.ISO8859-15@euro
fr_BE.ISO8859-1	fr_BE
fr_BE.ISO8859-15	fr_BE.ISO8859-15@euro
fr_BE.UTF-8	fr_BE.UTF-8@euro
fr_CA.ISO8859-1	fr_CA
fr_CH.ISO8859-1	fr_CH
fr_FR.ISO8859-1	fr, fr_FR
fr_FR.ISO8859-15	fr.ISO8859-15, fr_FR.ISO8859-15@euro
fr_FR.UTF-8	fr.UTF-8, fr_FR.UTF-8@euro
he_IL.ISO8859-8	he, he_IL
hr_HR.ISO8859-2	hr_HR
hu_HU.ISO8859-2	hu, hu_HU
is_IS.ISO8859-1	is_IS
it_IT.ISO8859-1	it, it_IT
it_IT.ISO8859-15	it.ISO8859-15, it_IT.ISO8859-15@euro
it_IT.UTF-8	it.UTF-8, it_IT.UTF-8@euro
ja_JP.eucJP	ja
ko_KR.EUC	ko
ko_KR.UTF-8	ko.UTF-8
lt_LT.ISO8859-13	lt, lt_LT
lv_LV.ISO8859-13	lv, lv_LV
mk_MK.ISO8859-5	mk_MK
nb_NO.ISO8859-1	no, no_NO, no_NO.ISO8859-1@bokmal
nL_BE.ISO8859-1	nL_BE
nL_BE.ISO8859-15	nL_BE.ISO8859-15@euro
nL_NL.ISO8859-1	nL, nL_NL
nL_NL.ISO8859-15	nL.ISO8859-15, nL_NL.ISO8859-15@euro
nn_NO.ISO8859-1	no_NO.ISO8859-1@nynorsk, no_NY
pl_PL.ISO8859-2	pl, pl_PL
pl_PL.UTF-8	pl.UTF-8
pt_BR.ISO8859-1	pt_BR
pt_PT.ISO8859-1	pt, pt_PT
pt_PT.ISO8859-15	pt.ISO8859-15, pt_PT.ISO8859-15@euro
ro_RO.ISO8859-2	ro_RO
ru_RU.ISO8859-5	ru, ru_RU
ru_RU.KOI8-R	ru.koi8-r
ru_RU.UTF-8	ru.UTF-8
sk_SK.ISO8859-2	sk_SK
sl_SI.ISO8859-2	sl_SI
sq_AL.ISO8859-2	sq_AL
sr_ME.ISO8859-5	sr_SP, sr_YU, sr_YU.ISO8859-5
sr_ME.UTF-8	sr_CS, sr_CS.UTF-8
sr_RS.ISO8859-5	sr_SP, sr_YU, sr_YU.ISO8859-5
sr_RS.UTF-8	sr_CS, sr_CS.UTF-8
sv_SE.ISO8859-1	sv, sv_SE
sv_SE.ISO8859-15	sv.ISO8859-15
sv_SE.UTF-8	sv.UTF-8
th_TH.TIS620	th, th_TH, th_TH.ISO8859-11

tr_TR.ISO8859-9	tr, tr_TR
zh_CN.EUC	zh
zh_CN.GBK	zh.GBK
zh_CN.UTF-8	zh.UTF-8
zh_TW.EUC	zh_TW

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

**See Also** [gettext\(1\)](#), [catopen\(3C\)](#), [gettext\(3C\)](#), [setlocale\(3C\)](#), [attributes\(5\)](#), [environ\(5\)](#)

**Name** man – macros to format Reference Manual pages

**Synopsis** nroff -man *filename*...

troff -man *filename*...

**Description** These macros are used to lay out the reference pages in this manual. If *filename* 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.

Any text argument *t* may be zero to six words. Quotes may be used to include SPACE characters in a “word”. If *text* 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.

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* are `ens`.

Type font and size are reset to default values before each paragraph, and after processing font and size setting macros.

These strings are predefined by `-man`:

`\*R` ‘@’, ‘(Reg)’ in nroff.

`\*S` Change to default type size.

**Requests** \* n.t.l. = next text line; p.i. = prevailing indent

<i>Request</i>	<i>Cause</i>	<i>If no</i>	<i>Explanation</i>
	<i>Break</i>	<i>Argument</i>	
<code>.Bt</code>	no	<code>t=n.t.l.*</code>	Text is in bold font.
<code>.BI<i>t</i></code>	no	<code>t=n.t.l.</code>	Join words, alternating bold and italic.
<code>.BR<i>t</i></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>i</i></code>	yes	<code>i=p.i.*</code>	Begin paragraph with hanging indent. Set prevailing indent to <i>i</i> .
<code>.It</code>	no	<code>t=n.t.l.</code>	Text is italic.
<code>.IB<i>t</i></code>	no	<code>t=n.t.l.</code>	Join words, alternating italic and bold.
<code>.IP<i>x i</i></code>	yes	<code>x=""</code>	Same as <code>.TP</code> with tag <i>x</i> .
<code>.IR<i>t</i></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>	
.IXt	no	-	Index macro, for SunSoft internal use.
.LP	yes	-	Begin left-aligned paragraph. Set prevailing indent to .5i.
.P	yes	-	Same as .LP.
.PDd	no	d=.4v	Set vertical distance between paragraphs.
.PP	yes	-	Same as .LP.
.RE	yes	-	End of relative indent. Restores prevailing indent.
.RBt	no	t=n.t.l.	Join words, alternating roman and bold.
.RI t	no	t=n.t.l.	Join words, alternating roman and italic.
.RSi	yes	i=p.i.	Start relative indent, increase indent by <i>i</i> . Sets prevailing indent to .5i for nested indents.
.SBt	no	-	Reduce size of text by 1 point, make text bold.
.SHt	yes	-	Section Heading.
.SMt	no	t=n.t.l.	Reduce size of text by 1 point.
.SSt	yes	t=n.t.l.	Section Subheading.
.THn s d f m	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.
.TPi	yes	i=p.i.	Begin indented paragraph, with the tag given on the next text line. Set prevailing indent to <i>i</i> .
.TXt p	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 *title*[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 `troff(1)` commands or escapes, and no macro requests.

**.SH SYNOPSIS**

Commands:

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.

Syntactic symbols appear in roman face:

[ ]

An argument, when surrounded by brackets is optional.

|

Arguments separated by a vertical bar are exclusive. You can supply only one item from such a list.

...

Arguments followed by an ellipsis can be repeated. When an ellipsis follows a bracketed set, the expression within the brackets can be repeated.

Functions:

If required, the data declaration, or `#include` directive, is shown first, followed by the function declaration. Otherwise, the function declaration is shown.

**.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 returns to the calling program and the conditions that cause these values to be returned.

### .SH EXIT STATUS

A list of the values the utility returns 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.

SMF service `svc:/application/man-index` is used to generate the index files for `-f`, `-k`, and `-K` options of `man(1)`. The SMF service is online by default. The index files for the options are generated during package/system installation if the FMRI is specified as an `restart_fmri` actuator. Any man page delivering packages that use the FMRI as `restart_fmri` actuator for its file actions must deliver a unique symbolic link to their man page directory at the following directory:

```
/usr/share/man/index.d/
```

The unique symbolic link names should be of package FMRI without scheme, slash character (`/`) replaced with `%2F`, and optionally have some or all sequences of numbers of version. For instance, for the following package FMRI:

```
pkg://opensolaris.org/library/libc@5.11,5.11-0.75:20071001T163427
```

could supply one of the symbolic link names such as:

```
opensolaris.org%2Flibrary%2Flibc  
opensolaris.org%2Flibrary%2Flibc@5.11
```

The method of the SMF service, upon requested to generate index files using `restart_fmri` during package install, checks the symbolic links. Based on the newness of the installed files, index files are generated only as necessary.

---

Unless used in an IPS package as `restart_fmri`, by default, it generates index files in `/usr/share/man/` and `/usr/gnu/share/man/`. When used in an IPS package as `restart_fmri`, it generates and places index files in the same directory as the man page source directory of the package.

**Files** `/usr/share/man/index.d/` Directory to save symbolic links to man page directories needing index files for `-f`, `-k`, and `-K` queries.

`/usr/share/man/man_index/*` Table of Contents and keyword database.

Generated files include:

- `/usr/share/man/man_index/man.idx`
- `/usr/share/man/man_index/man.dic`
- `/usr/share/man/man_index/man.frq`
- `/usr/share/man/man_index/man.pos`

`/usr/share/lib/tmac/an`

**See Also** [man\(1\)](#), [nroff\(1\)](#), [troff\(1\)](#), [whatis\(1\)](#)

Dale Dougherty and Tim O'Reilly, *Unix Text Processing*

**Name** mansun – macros to format Reference Manual pages

**Synopsis** nroff -mansun *filename*...

troff -mansun *filename*...

**Description** These macros are used to lay out the reference pages in this manual. Note: if *filename* contains format input for a preprocessor, the commands shown above must be piped through the appropriate preprocessor. This is handled automatically by [man\(1\)](#). See the “Conventions” section.

Any text argument *t* may be zero to six words. Quotes may be used to include SPACE characters in a “word”. If *text* 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.

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* are ens.

Type font and size are reset to default values before each paragraph, and after processing font and size setting macros.

These strings are predefined by -mansun:

`\*R` ‘*o*’, ‘(Reg)’ in nroff.

`\*S` Change to default type size.

Requests \* n.t.l. = next text line; p.i. = prevailing indent

<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 li...</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> .

<i>Request</i>	<i>Cause</i>	<i>If no</i>	<i>Explanation</i>
	<i>Break</i>	<i>Argument</i>	
.IR <i>t</i>	no	<i>t</i> =n.t.l.	Join words, alternating italic and Roman.
.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</i> =.4v	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</i> =n.t.l.	Join words, alternating Roman and bold.
.RI <i>t</i>	no	<i>t</i> =n.t.l.	Join words, alternating Roman and italic.
.RS <i>i</i>	yes	<i>i</i> =p.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</i> =n.t.l.	Reduce size of text by 1 point.
.SS <i>t</i>	yes	<i>t</i> =n.t.l.	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</i> =p.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** *title* [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 `troff(1)` commands or escapes, and no macro requests. It is used to generate the `windex` database, which is used by the `what-is(1)` command.
- .SH** SYNOPSIS
- Commands:    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.
- Syntactic symbols appear in Roman face:
- [ ]    An argument, when surrounded by brackets is optional.
  - |    Arguments separated by a vertical bar are exclusive. You can supply only one item from such a list.
  - . . .    Arguments followed by an ellipsis can be repeated. When an ellipsis follows a bracketed set, the expression within the brackets can be repeated.
- Functions:    If required, the data declaration, or `#include` directive, is shown first, followed by the function declaration. Otherwise, the function declaration is shown.
- .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*

**Name** me – macros for formatting papers

**Synopsis** nroff -me [*options*] *filename*...

troff -me [*options*] *filename*...

**Description** 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\(1\)](#).

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:

.bp        begin new page  
.br        break output line here  
.sp *n*     insert *n* spacing lines  
.ls *n*     (line spacing) *n*=1 single, *n*=2 double space  
.na        no alignment of right margin  
.ce *n*     center next *n* lines  
.ul *n*     underline next *n* lines  
.sz +*n*    add *n* to point size

Output of the [eqn\(1\)](#), [neqn\(1\)](#), [refer\(1\)](#), and [tbl\(1\)](#) preprocessors for equations and tables is acceptable as input.

**Requests** In the following list, “initialization” refers to the first .pp, .lp, .ip, .np, .sh, or .uh macro. This list is incomplete.

<i>Request</i>	<i>Initial</i>	<i>Cause</i>	<i>Explanation</i>
	<i>Value</i>	<i>Break</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.

<i>Request</i>	<i>Initial Value</i>	<i>Cause Break</i>	<i>Explanation</i>
.)c	-	yes	End centered block.
.)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>Request</i>	<i>Initial Value</i>	<i>Cause Break</i>	<i>Explanation</i>
			C to center the equation.
.GE	-	yes	End <i>gremlin</i> picture.
.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 'x'y'z	""	no	Set even footer to <i>x y z</i> .
.eh 'x'y'z	""	no	Set even header to <i>x y z</i> .
.fo 'x'y'z	""	no	Set footer to <i>x y z</i> .
.hx	-	no	Suppress headers and footers on

<i>Request</i>	<i>Initial Value</i>	<i>Cause Break</i>	<i>Explanation</i>
			next page.
.he 'x'y'z	""	no	Set header to $x y z$ .
.hl	-	yes	Draw a horizontal line.
.i x	no	no	Italicize $x$ ; if $x$ missing, italic text follows.
.ip $x y$	no	yes	Start indented paragraph, with hanging tag $x$ . Indentation is $y$ ens (default 5).
.lp	yes	yes	Start left-blocked paragraph.
.lo	-	no	Read in a file of local macros of the form $. *x$ . Must be given before initialization.
.np	1	yes	Start numbered paragraph.
.of 'x'y'z	""	no	Set odd footer to $x y z$ .
.oh 'x'y'z	""	no	Set odd header to $x y z$ .
.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 $n x$	-	yes	Section head follows, font automatically bold. $n$ is level of section, $x$ is title of section.
.sk	no	no	Leave the next page blank. Only one page is remembered ahead.

<i>Request</i>	<i>Initial Value</i>	<i>Cause Break</i>	<i>Explanation</i>
<code>.sm x</code>	-	no	Set <i>x</i> in a smaller pointsize.
<code>.sz +n</code>	10p	no	Augment the point size by <i>n</i> points.
<code>.th</code>	no	no	Produce the paper in thesis format. Must be given before initialization.
<code>.tp</code>	no	yes	Begin title page.
<code>.u x</code>	-	no	Underline argument (even in <code>troff</code> ). (Nofill only).
<code>.uh</code>	-	yes	Like <code>.sh</code> but unnumbered.
<code>.xp x</code>	-	no	Print index <i>x</i> .

**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** mech\_spnego – Simple and Protected GSS-API Negotiation Mechanism

**Synopsis** /usr/lib/gss/mech\_spnego.so.1

**Description** The SPNEGO security mechanism for GSS-API allows GSS-API applications to negotiate the actual security mechanism to be used in the GSS-API session. mech\_spnego.so.1 is a shared object module that is dynamically opened by applications that specify the SPNEGO Object Identifier (OID) in calls to the GSS-API functions (see [libgss\(3LIB\)](#)).

SPNEGO is described by IETF RFC 2478 and is intended to be used in environments where multiple GSS-API mechanisms are available to the client or server and neither side knows what mechanisms are supported by the other.

When SPNEGO is used, it selects the list of mechanisms to advertise by reading the GSS mechanism configuration file, /etc/gss/mech (see [mech\(4\)](#)), and by listing all active mechanisms except for itself.

**Options** SPNEGO may be configured to function in two ways. The first way is to interoperate with Microsoft SSPI clients and servers that use the Microsoft "Negotiate" method, which is also based on SPNEGO. The Microsoft "Negotiate" mechanism does not strictly follow the IETF RFC. Therefore, use special handling in order to enable full interoperability. In order to interoperate, place option "[ msinterop ]" at the end of the SPNEGO line in /etc/gss/mech.

This is an example (from /etc/gss/mech):

```
spnego    1.3.6.1.5.5.2    mech_spnego.so    [ msinterop ]
```

Without the "[ msinterop ]" option, mech\_spnego will follow the strict IETF RFC 2478 specification and will not be able to negotiate with Microsoft applications that try to use the SSPI "Negotiate" mechanism.

**Interfaces** mech\_spnego.so.1 has no public interfaces. It is only activated and used through the GSS-API interface provided by libgss.so.1 (see [libgss\(3LIB\)](#)).

**Files**

/usr/lib/gss/mech_spnego.so.1	shared object file
/usr/lib/sparcv9/gss/mech_spnego.so.1	SPARC 64-bit shared object file
/usr/lib/amd64/gss/mech_spnego.so.1	x86 64-bit shared object file

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUWNspnego
MT Level	Safe

**See Also** [Intro\(3\)](#), [libgss\(3LIB\)](#), [mech\(4\)](#), [attributes\(5\)](#)

*Developer's Guide to Oracle Solaris 11 Security*

**Name** mm – text formatting (memorandum) macros

**Synopsis** nroff -mm [*options*] *filename*...

troff -mm [*options*] *filename*...

**Description** This package of `nroff(1)` and `troff(1)` 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 `col(1)`. All external -mm macros are defined below.

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.

Many `nroff` and `troff` 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:

.bp      begin new page  
 .br      break output line  
 .spn     insert n spacing lines  
 .cen     center next n lines  
 .lsn     line spacing:  $n=1$  single,  $n=2$  double space  
 .na      no alignment of right margin

Font and point size changes with `\f` and `\s` are also allowed; for example, `\fIword\fR` will italicize *word*. Output of the `tbl(1)`, `eqn(1)` and `refer(1)` preprocessors for equations, tables, and references is acceptable as input.

**Requests** Here is a table of macros.

Macro Name	Initial Value	Break? Reset?	Explanation
.1C	on	y,y	one column format on a new page
.2C [ <i>l</i> ]	–	y,y	two column format <i>l</i> =line length
.AE	–	y	end abstract
.AL [ <i>t</i> ] [ <i>i</i> ] [ <i>s</i> ]	$t=1; i= .L; s=0$	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 indention <i>i</i> ; separation <i>s</i>

Macro Name	Initial Value	Break? Reset?	Explanation
.AS $m [ n ]$	$n=0$	y	begin abstract
.AU	–	y	author's name
.AV $x$	–	y	signature and date line of verifier $x$
.B $x$	–	n	embolden $x$ ; if no $x$ , switch to boldface
.BE	–	y	end block text
.BI $x y$	–	n	embolden $x$ and underline $y$
.BL	–	y	bullet list
.BR $x y$	–	n	embolden $x$ and use Roman font for $y$
.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 [ $p$ ][ $f$ ][ $rp$ ]	$p=L;f=N$	y	start floating display; position $p=[L,C,CB]$ L=left; I=indent; C=center; CB=center block fill $f=[N,Y]$ ; right position $rp$ (fill only)
.DL [ $i$ ][ $s$ ]	–	y	start dash list
.DS [ $p$ ][ $f$ ][ $rp$ ]	$p=L;f=N$	y	begin static display (see .DF for argument descriptions)
.EC $x [ n ]$	$n=1$	y	equation title; equation $x$ ; number $n$
.EF $x$	–	n	even footer appears at the bottom of even-numbered pages; $x="l' c' r"$ l=left; c=center; r=right
.EH $x$	–	n	even header appears at the top of even-numbered pages; $x="l' c' r"$ l=left; c=center; r=right
.EN	–	y	end displayed equation produced by eqn
.EQ	–	y	break out equation produced by eqn
.EX $x [ n ]$	$n=1$	y	exhibit title; exhibit $x$
			number $n$

Macro Name	Initial Value	Break? Reset?	Explanation
.FD [ <i>f</i> ] [ <i>r</i> ]	$f=10;r=1$	n	set footnote style format $f=[0-11]$ ; renumber $r=[0,1]$
.FE	–	y	end footnote
.FG <i>x</i> [ <i>n</i> ]	$n=1$	y	figure title; figure <i>x</i> ; number <i>n</i>
.FS	–	n	start footnote
.Hl [ <i>t</i> ]	–	y	produce numbered heading level $l=[1-7]$ ; 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> ]	$s=0$	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> ]	$s=0$	y	start marked list; mark <i>m</i> indentation <i>i</i> ; separation $s=[0,1]$
.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; $x="l'c'r"$ $l$ =left; $c$ =center; $r$ =right
.OF <i>x</i>	–	n	odd header appears at the top of odd-numbered pages; $x="l'c'r"$ $l$ =left; $c$ =center; $r$ =right
.OP	–	y	skip to the top of an odd-number page
.P [ <i>t</i> ]	$t=0$	y,y	begin paragraph; $t=[0,1]$ $0$ =justified; $1$ =indented
.PF <i>x</i>	–	n	page footer appears at the bottom of every page; $x="l'c'r"$ $l$ =left; $c$ =center; $r$ =right

Macro Name	Initial Value	Break? Reset?	Explanation
.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
.VL <i>i</i> [ <i>m</i> ] [ <i>s</i> ]	<i>m</i> =0; <i>s</i> =0	y	start variable-item list; indentation <i>i</i> mark-indentation <i>m</i> ; separation <i>s</i>

**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:

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)

Name	String's Function
\*(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	text/doctools

**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.

**Name** ms – text formatting macros

**Synopsis** nroff -ms [options] filename...

troff -ms [options] filename...

**Description** This package of `nroff(1)` and `troff(1)` 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 `col(1)`. All external -ms macros are defined below.

Note: this -ms macro package is an extended version written at Berkeley and is a superset of the standard -ms 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.

Many `nroff` and `troff` 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:

.bp      begin new page  
 .br      break output line  
 .sp *n*    insert *n* spacing lines  
 .ce *n*    center next *n* lines  
 .ls *n*    line spacing: *n*=1 single, *n*=2 double space  
 .na      no alignment of right margin

Font and point size changes with `\f` and `\s` are also allowed; for example, `\fIword\fR` will italicize *word*. Output of the `tbl(1)`, `eqn(1)` and `refer(1)` preprocessors for equations, tables, and references is acceptable as input.

## Requests

Macro Name	Initial Value	Break? Reset?	Explanation
.AB <i>x</i>	–	y	begin abstract; if <i>x</i> =no do not label abstract
.AE	–	y	end abstract
.AI	–	y	author's institution
.AM	–	n	better accent mark definitions
.AU	–	y	author's name
.B <i>x</i>	–	n	embolden <i>x</i> ; if no <i>x</i> , switch to boldface
.B1	–	y	begin text to be enclosed in a box

Macro Name	Initial Value	Break? Reset?	Explanation
.B2	–	y	end boxed text and print it
.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	–	y,y	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>	I	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>	–	y,y	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	–	y,y	left (block) paragraph.

Macro Name	Initial Value	Break? Reset?	Explanation
.MC <i>x</i>	–	y,y	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
.NH <i>x y</i>	–	y,y	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	–	y,y	paragraph with first line indented
.PT	- % -	n	page title, printed at head of page
.PX <i>x</i>	–	y	print index (table of contents); <i>x</i> =no suppresses title
.QP	–	y,y	quote paragraph (indented and shorter)
.R	on	n	return to Roman font
.RE	5n	y,y	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	y,y	right shift: start level of relative indentation
.SH	–	y,y	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>	–	y	print table of contents at end; <i>x</i> =no suppresses title
.TE	–	y	end of table processed by tbl
.TH	–	y	end multi-page header of table
.TL	–	y	title in boldface and two points larger
.TM	off	n	UC Berkeley thesis mode
.TS <i>x</i>	–	y,y	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>	–	y	another index entry; <i>x</i> =page or no for none; <i>y</i> =indent

Macro Name	Initial Value	Break? Reset?	Explanation
.XE	–	y	end index entry (or series of .IX entries)
.XP	–	y,y	paragraph with first line indented, others indented
.XS x y	–	y	begin index entry; x=page or no for none; y=indent
.1C	on	y,y	one column format, on a new page
.2C	–	y,y	begin two column format
.] –	–	n	beginning of refer 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.

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>\*<sup>ˆ</sup></code>	grave accent (before letter)
<code>\*<sup>ˆ</sup></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** mutex – concepts relating to mutual exclusion locks

**Description** Mutual exclusion locks (mutexes) prevent multiple threads from simultaneously executing critical sections of code which access shared data (that is, mutexes are used to serialize the execution of threads). All mutexes must be global. A successful call to acquire a mutex will cause another thread that is also trying to lock the same mutex to block until the owner thread unlocks the mutex.

Mutexes can synchronize threads within the same process or in other processes. Mutexes can be used to synchronize threads between processes if the mutexes are allocated in writable memory and shared among the cooperating processes (see [mmap\(2\)](#)), and have been initialized for this task.

The following table lists mutex functions and the actions they perform.

FUNCTION	ACTION
<code>mutex_init</code>	Initialize a mutex.
<code>mutex_destroy</code>	Destroy a mutex.
<code>mutex_lock</code>	Lock a mutex.
<code>mutex_trylock</code>	Attempt to lock a mutex.
<code>mutex_unlock</code>	Unlock a mutex.
<code>pthread_mutex_init</code>	Initialize a mutex.
<code>pthread_mutex_destroy</code>	Destroy a mutex.
<code>pthread_mutex_lock</code>	Lock a mutex.
<code>pthread_mutex_trylock</code>	Attempt to lock a mutex.
<code>pthread_mutex_unlock</code>	Unlock a mutex.

**Initialization** Mutexes are either intra-process or inter-process, depending upon the argument passed implicitly or explicitly to the initialization of that mutex. A statically allocated mutex does not need to be explicitly initialized; by default, a statically allocated mutex is initialized with all zeros and its scope is set to be within the calling process.

For inter-process synchronization, a mutex needs to be allocated in memory shared between these processes. Since the memory for such a mutex must be allocated dynamically, the mutex needs to be explicitly initialized with the appropriate attribute that indicates inter-process use.

**Locking and Unlocking** A critical section of code is enclosed by a call to lock the mutex and the call to unlock the mutex to protect it from simultaneous access by multiple threads. Only one thread at a time may possess mutually exclusive access to the critical section of code that is enclosed by the mutex-locking call and the mutex-unlocking call, whether the mutex's scope is intra-process

or inter-process. A thread calling to lock the mutex either gets exclusive access to the code starting from the successful locking until its call to unlock the mutex, or it waits until the mutex is unlocked by the thread that locked it.

Mutexes have ownership, unlike semaphores. Only the thread that locked a mutex, (that is, the owner of the mutex), should unlock it.

If a thread waiting for a mutex receives a signal, upon return from the signal handler, the thread resumes waiting for the mutex as if there was no interrupt.

**Caveats** Mutexes are almost like data – they can be embedded in data structures, files, dynamic or static memory, and so forth. Hence, they are easy to introduce into a program. However, too many mutexes can degrade performance and scalability of the application. Because too few mutexes can hinder the concurrency of the application, they should be introduced with care. Also, incorrect usage (such as recursive calls, or violation of locking order, and so forth) can lead to deadlocks, or worse, data inconsistencies.

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

**See Also** [mmap\(2\)](#), [shmop\(2\)](#), [mutex\\_destroy\(3C\)](#), [mutex\\_init\(3C\)](#), [mutex\\_lock\(3C\)](#), [mutex\\_trylock\(3C\)](#), [mutex\\_unlock\(3C\)](#), [pthread\\_create\(3C\)](#), [pthread\\_mutex\\_destroy\(3C\)](#), [pthread\\_mutex\\_init\(3C\)](#), [pthread\\_mutex\\_lock\(3C\)](#), [pthread\\_mutex\\_trylock\(3C\)](#), [pthread\\_mutex\\_unlock\(3C\)](#), [pthread\\_mutexattr\\_init\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

**Notes** In the current implementation of threads, [pthread\\_mutex\\_lock\(\)](#), [pthread\\_mutex\\_unlock\(\)](#), [mutex\\_lock\(\)](#), [mutex\\_unlock\(\)](#), [pthread\\_mutex\\_trylock\(\)](#), and [mutex\\_trylock\(\)](#) do not validate the mutex type. Therefore, an uninitialized mutex or a mutex with an invalid type does not return EINVAL. Interfaces for mutexes with an invalid type have unspecified behavior.

By default, if multiple threads are waiting for a mutex, the order of acquisition is undefined.

The system does not support multiple mappings to the same logical synch object if it is initialized as process-private (USYNC\_THREAD for Solaris, PTHREAD\_PROCESS\_PRIVATE for POSIX). If you need to [mmap\(2\)](#) a synch object to different locations within the same address space, then the synch object should be initialized as a shared object (USYNC\_PROCESS for Solaris, PTHREAD\_PROCESS\_SHARED for POSIX).

**Name** mwac, MWAC – Mandatory Write Access Control

**Description** Mandatory Write Access Control (MWAC) implements a new policy in the Oracle Solaris operating environment, that allows for fine-grained control over the writability of objects on otherwise read-only file systems.

In the current instance of the Oracle Solaris operating environment, MWAC is available only to non-global zones. The global zone implements the MWAC policy for non-global zones, preventing any overruling of the policy from within the non-global zone.

Zones marked as read-only have their root file system write-protected by MWAC. Only the file system objects that are write-listed by the read-only-profile are writable. See [zonecfg\(1M\)](#). Other file system objects are read-only.

Creating links to objects that are read-only by virtue of the MWAC-policy is not allowed.

**See Also** [ln\(1\)](#), [zoneadm\(1M\)](#), [zonecfg\(1M\)](#), [link\(2\)](#), [pathconf\(2\)](#)

*Oracle Solaris 11.1 Administration: Security Services*

**Name** nfssec – overview of NFS security modes

**Description** The `mount_nfs(1M)` and `share_nfs(1M)` commands each provide a way to specify the security mode to be used on an NFS file system through the `sec=mode` option. *mode* can be `sys`, `dh`, `krb5`, `krb5i`, `krb5p`, or `none`. These security modes can also be added to the automount maps. `mount_nfs(1M)` allows you to specify a single security mode; `share_nfs(1M)` allows you to specify multiple modes (or `none`). With multiple modes, an NFS client can choose any of the modes in the list.

The `sec=mode` option on the `share_nfs(1M)` 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 *mode* to use. If the NFS connection uses the NFS Version 2 protocol, then the NFS client uses the default security mode, which is currently `sys`. NFS clients may force the use of a specific security mode by specifying the `sec=mode` 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.

If the NFS client wants to authenticate the NFS server using a particular (stronger) security mode, the client wants 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.

The NFS security modes are described below. Of these, the `krb5`, `krb5i`, `krb5p` 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 [kerberos\(5\)](#).

<code>sys</code>	Use AUTH_SYS 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.
<code>dh</code>	Use a Diffie-Hellman public key system (AUTH_DES, which is referred to as AUTH_DH in <i>RFC 2695: Authentication Mechanisms for ONC RPC</i> ).
<code>krb5</code>	Use Kerberos V5 protocol to authenticate users before granting access to the shared filesystem.
<code>krb5i</code>	Use Kerberos V5 authentication with integrity checking (checksums) to verify that the data has not been tampered with.
<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.

none	Use null authentication (AUTH_NONE). NFS clients using AUTH_NONE have no identity and are mapped to the anonymous user nobody by NFS servers. A client using a security mode other than the one with which a Solaris NFS server shares the file system has its security mode mapped to AUTH_NONE. In this case, if the file system is shared with <i>sec=none</i> , users from the client are mapped to the anonymous user. The NFS security mode none is supported by <a href="#">share_nfs(1M)</a> .
<i>sec=mode[:mode]. . .</i>	<p>Sharing uses one or more of the specified security modes. The <i>mode</i> in the <i>sec=mode</i> option must be a node name supported on the client. If the <i>sec=</i> option is not specified, the default security mode used is AUTH_SYS. Multiple <i>sec=</i> options can be specified on the command line, although each mode can appear only once.</p> <p>Each <i>sec=</i> option specifies modes that apply to any subsequent <i>window=</i>, <i>rw</i>, <i>ro</i>, <i>rw=</i>, <i>ro=</i> and <i>root=</i> options that are provided before another <i>sec=</i> option. Each additional <i>sec=</i> resets the security mode context, so that more <i>window=</i>, <i>rw</i>, <i>ro</i>, <i>rw=</i>, <i>ro=</i> and <i>root=</i> options can be supplied for additional modes.</p>

The NFSv4 server constructs a shared file system name space which is identical to the real file system name space on the server, including directories which are not actually shared, if they lead to shared directories. The constructed parts of the name space are known as the pseudo-*fs*. The pseudo-*fs* is always read-only.

As with NFSv3, the security mode of the shared directory is controlled using the *sec=mode* option of [share\\_nfs\(1M\)](#). However, the security mode of pseudo-*fs* objects is the union of the various security modes of the shared directories below.

When an NFSv4 client performs a mount, the client traverses the server's name space, from the root, down to the directory being mounted. Using the features of the NFSv4 protocol, the client may negotiate the security flavor of the directories as it proceeds down. If no *sec=mode* option is given to *mount\_nfs* or an automounter map entry, then the client will do full negotiation for each directory down to the mount point, changing security flavors as needed. If *sec=mode* option is given, the client is constrained to use the requested security mode for all operations.

#### Examples EXAMPLE 1 Sharing /var with Kerberos Authentication and Integrity Protection

The following example shares /var with Kerberos authentication and integrity protection:

```
share -F nfs -o sec=krb5i /var
```

**EXAMPLE 2** Sharing /var with Kerberos Authentication and Privacy Protection

The following example shares /var with Kerberos authentication and privacy protection:

```
share -F nfs -o sec=krb5p /var
```

**EXAMPLE 3** Sharing /var with Kerberos Authentication and Optionally Falling Back to AUTH\_SYS Authentication

The following example shares /var with Kerberos authentication and optionally falls back to AUTH\_SYS authentication:

```
share -F nfs -o sec=krb5:sys /var
```

**EXAMPLE 4** Sharing /var with Kerberos Authentication Allowing read/write Operations for Kerberos Authenticated Users and Optionally Falling Back to AUTH\_SYS Authentication Allowing only Read Operations

The following example shares /var with Kerberos authentication allowing read/write operations for Kerberos authenticated users and optionally falls back to AUTH\_SYS authentication allowing only read operations:

```
share -F nfs -o sec=krb5,rw,sec=sys,ro /var
```

**Files** /etc/nfssec.conf NFS security service configuration file

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/nfs

**See Also** [automount\(1M\)](#), [kclient\(1M\)](#), [mount\\_nfs\(1M\)](#), [share\\_nfs\(1M\)](#), [rpc\\_clnt\\_auth\(3NSL\)](#), [secure\\_rpc\(3NSL\)](#), [nfssec.conf\(4\)](#), [attributes\(5\)](#), [kerberos\(5\)](#)

*RFC 2695: Authentication Mechanisms for ONC RPC*

**Notes** /etc/nfssec.conf lists the NFS security services. Do not edit this file. It is not intended to be user-configurable. See [kclient\(1M\)](#).

**Name** nis, NIS, NIS+ – network information name service

**Description** NIS, formerly known as the Yellow Pages or YP, is the name of the network information name service in common use in networks on which Sun and other vendors' network nodes reside. The most recent version of NIS is version 2.

All commands and functions that use NIS version 2 are prefixed by the letters `yp` as in `ypmatch(1)`, `ypcat(1)`, `yp_match(3NSL)`, and `yp_first(3NSL)`.

The NIS+ name service is no longer shipped with the Solaris operating system. Tools to aid the migration from NIS+ to LDAP are available in the current Solaris release.

**Name** nwam – network auto-magic configuration management

**Description** Network configuration can be managed automatically, according to default policy rules defined by the system, or to user-defined policy rules. This management is performed by the `nwamd` daemon.

The NWAM facility is provided by the `network/physical:default` SMF service. NWAM configuration management is enabled if the `netcfg/active_ncp` property is set to the name of a reactive NCP. Alternatively, traditional network configuration is performed if the `netcfg/active_ncp` property is set to `DefaultFixed`.

**Operation** The default configuration policy, implemented in the Automatic NCP, is to have all physically connected Ethernet links active; if no Ethernet links are available, a single wireless link is made active. DHCP is used to obtain IP addresses for all active links. This default policy can be changed by creating alternate Network Configuration Profiles (NCPs). The `nwam-manager` GUI tool or the `nwamd(1M)` command can be used to create and modify NCPs.

NWAM management is also available for higher-layer network configuration, such as name services and network security protocols. This configuration is specified in a Location profile.

Profiles can be managed using `netcfg(1M)`. This is the primary tool for specifying the activation policy for all reactive profiles. It can also be used to specify locations as well as simple NCP components.

Configuration changes made using `dladm(1M)` or `ipadm(1M)` are applied to the currently active NCP, whether that NCP is reactive (that is, actively managed by `nwamd`) or fixed.

Locations manage a more diverse set of configuration objects. The active system configuration for these objects resides in SMF; the NWAM facility manages that configuration by storing settings in a profile repository, and applying those settings to the SMF repository when the location is activated.

For reactive locations, any changes made to the SMF repository while the location is active will be lost when the location is disabled; rather, changes should be made to the NWAM profile using `netcfg`. These changes will be saved to the NWAM repository, and will also be applied to the active system configuration if made to the currently active location.

For the `DefaultFixed` location, changes made to the SMF repository will be saved to the NWAM repository when the location is disabled. They will thus be restored the next time the `DefaultFixed` location is enabled.

**Service Properties** The following list takes the form:

```
property_group/property_name property_type default_value  
Description
```

```
nwamd/debug boolean false  
    Enables debug logging using daemon.debug.
```

`nwamd/autoconf` boolean `false`

Indicates whether open WLANs should be connected automatically, in the absence of a better (more preferred) choice.

`nwamd/ncu_wait_time` count `60`

The number of seconds to wait for an NCU (or link/interface NCU pair, as appropriate) to come up before trying the next available NCU. The bringup activity is not cancelled, and might eventually succeed, at which time the more preferred NCU is activated and the alternate might be disabled, depending on the specified configuration conditions.

`nwamd/condition_check_interval` count `120`

The number of seconds between periodic condition checks for conditionally activated objects. Minimum value is 30 seconds.

`nwamd/scan_interval` count `120`

The number of seconds between periodic wireless scans.

`nwamd/scan_level` astring `weak`

A signal strength threshold; if the currently connected AP drops below this signal level, and equivalent APs (of the same ESSID) are available at higher signal strength, the existing connection is dropped in favor of a connection to an AP with stronger signal.

`nwamd/strict_bssid` boolean `false`

If true, both ESSID and BSSID must be matched in order to connect to a previously connected WLAN. If false, only an ESSID match is required.

`netcfg/active_ncp` astring `Automatic`

The currently active NCP. This property should not be set by the user; it is used internally by the NWAM service for persistence across restarts. The appropriate user interface to change the currently active NCP is by means of the `nwam-manager` GUI or the `netadm` command's `enable` subcommand.

Enabling the reserved NCP name `DefaultFixed` disables the NWAM automatic configuration management and enables traditional, fixed network configuration.

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	<code>system/core-os</code>
Interface Stability	<code>Volatile</code>

**See Also** [svcs\(1\)](#), [dladm\(1M\)](#), [ipadm\(1M\)](#), [netcfg\(1M\)](#), [netcfgd\(1M\)](#), [nwamd\(1M\)](#), [svadm\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

See also `nwam-manager(1M)`, available in the JDS/GNOME man page collection.

**Notes** The networking service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/physical:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

<b>Name</b>	openssl – OpenSSL cryptographic and Secure Sockets Layer toolkit
<b>Description</b>	<p>OpenSSL is a cryptography toolkit that implements the Secure Sockets Layer (SSLv2/v3) and Transport Layer Security (TLS v1) network protocols.</p> <p>The following features are omitted from the binaries for issues including but not limited to patents, trademark, and US export restrictions: ECC, IDEA, MDC2, RC3, RC5, 4758_CCA Engine, AEP Engine, Atalla Engine, CHIL Engine, CSWIFT Engine, GMP Engine, NURON Engine, PadLock Engine, Sureware Engine, and UBSEC Engine.</p>
The PKCS#11 Engine	<p>A new PKCS#11 engine has been included with ENGINE name pkcs11. The engine was developed in Sun and is not integrated in the OpenSSL project.</p> <p>The PKCS#11 engine is configured to use the Oracle Solaris Cryptographic Framework. See <a href="#">cryptoadm(1M)</a> for configuration information.</p> <p>The PKCS#11 engine can support the following set of mechanisms: CKM_AES_CBC, CKM_AES_ECB, CKM_BLOWFISH_CBC, CKM_DES_CBC, CKM_DES_ECB, CKM_DES3_CBC, CKM_DES3_ECB, CKM_DSA, CKM_MD5, CKM_RC4, CKM_RSA_PKCS, CKM_RSA_X_509, CKM_SHA_1, CKM_SHA224, CKM_SHA256, CKM_SHA384, CKM_SHA512, CKM_SHA224_HMAC, CKM_SHA224_HMAC_GENERAL, and CKM_SHA224_KEY_DERIVATION.</p> <p>The set of mechanisms available depends on installed Crypto Framework providers. To see what mechanisms can be offloaded to the Cryptographic Framework through the PKCS#11 engine on a given machine, run the following command:</p> <pre>/usr/sfw/bin/openssl engine pkcs11 -vvv -t -c</pre> <p>Due to requirements of the PKCS#11 standard regarding <a href="#">fork(2)</a> behavior, some applications that use the OpenSSL EVP interfaces and <code>fork()</code> with active crypto contexts might experience unexpected behavior.</p>
The Dynamic Engine Support	<p>The dynamic engine support has been enabled, which allows an external engine, in the form of a shared library, to be dynamically bound and used by an OpenSSL-based application.</p> <p>Run the following command to see if the dynamic engine is supported:</p> <pre>\$openssl engine dynamic (dynamic) Dynamic engine loading support</pre>
The devcrypto Engine	<p>This engine is implemented as a separate shared library, and it can be used by an OpenSSL application through the OpenSSL dynamic engine support. This engine was developed in Sun and is not integrated in the OpenSSL project.</p> <p>Files for the devcrypto engine library are:</p> <pre>/lib/openssl/engines/libdevcrypto.so /lib/openssl/engines/64/libdevcrypto.so</pre>

The devcrypto engine accesses only the kernel hardware providers from the Oracle Solaris Cryptographic Framework. To see the hardware provider information on a system, use the [cryptoadm\(1M\)](#) command.

The devcrypto engine supports the following set of mechanisms: CKM\_AES\_CBC, CKM\_AES\_CTR, CKM\_AES\_ECB, CKM\_BLOWFISH\_CBC, CKM\_DES\_CBC, CKM\_DES\_ECB, CKM\_DES3\_CBC, CKM\_DES3\_ECB, and CKM\_RC4.

The set of mechanisms available depends on hardware providers installed and enabled in the Cryptographic Framework. To see what mechanisms are supported by the devcrypto engine on a particular machine, run the following command:

```
$openssl engine dynamic -pre SO_PATH:/lib/openssl/engines/libdevcrypto.so\
-pre LOAD -t -c
```

Using FIPS Mode Currently a FIPS-140 certified mode is not available in Oracle Solaris.

Building an OpenSSL Application To build an OpenSSL application, use the following cc command line options:

```
cc [ flag... ] file... -lcrypto -lssl [ library... ]
```

Accessing RSA Keys in PKCS#11 Keystores OpenSSL can access RSA keys in PKCS#11 keystores using the following functions of the ENGINE API:

```
EVP_PKEY *ENGINE_load_private_key(ENGINE *e,
    const char *key_id, UI_METHOD *ui_method,
    void *callback_data)
```

```
EVP_PKEY *ENGINE_load_public_key(ENGINE *e,
    const char *key_id, UI_METHOD *ui_method,
    void *callback_data)
```

`key_id`, formerly for filenames only, can be now also set to a PKCS#11 URI. The `EVP_PKEY` structure is newly allocated and caller is responsible to free the structure later. To avoid clashes with existing filenames, `file://` prefix for filenames is now also accepted but only when the PKCS#11 engine is in use. The PKCS#11 URI specification follows:

```
pkcs11:[token=<label>][:manuf=<label>][;serial=<label>]
    [;model=<label>][;object=<label>]
    [;objecttype=(public|private|cert)]
    [;passphrasedialog=(builtin|exec:<file>)]
```

The ordering of keywords is not significant. The PKCS#11 engine uses the keystore for the slot chosen for public key operations, which is `metaslot` on a standard configured machine. Currently, the PKCS#11 engine ignores the `objecttype` keyword. The only mandatory keyword is `object` which is the key object label. For information on how to use a different, possibly hardware, keystore with `metaslot`, see [libpkcs11\(3LIB\)](#).

The token PIN is provided by way of the `passphrasedialog` keyword and is either read from the terminal (`builtin`) or from the output of an external command (`exec:<file>`). The PIN

is used to log into the token and by default is deleted from the memory then. The keyword `pin` is intentionally not provided due to inherent security problems of possible use of a password in the process arguments.

Due to fork safety issues the application must re-login if the child continues to use the PKCS#11 engine. It is done inside of the engine automatically if fork is detected and in that case, `exec:<file>` option of the `passphrasedialog` keyword can be used. Alternatively, an environment variable `OPENSSL_PKCS11_PIN_CACHING_POLICY` can be used to allow the PIN to be cached in memory and reused in the child. It can be set to `none` which is the default, `memory` to store the PIN in memory, and `mlocked-memory` to keep the PIN in a locked page using `mlock(3C)`. `PRIV_PROC_LOCK_MEMORY` privilege is required in that case.

Sensitive parts of private keys are never read from the token to the process memory no matter whether the key is tagged with sensitive flag or not. The PKCS#11 engine uses the public components as a search key to get a PKCS#11 object handle to the private key.

To use the RSA keys by reference, high level API functions such as `RSA_public_decrypt()`, `EVP_PKEY_set1_RSA()`, or `EVP_SignInit()` must be used. Low level functions might go around the engine and fail to make use of the feature.

**Additional Documentation**

Extensive additional documentation for OpenSSL modules is available in the `/usr/share/man/man1openssl`, `/usr/share/man/man3openssl`, `/usr/share/man/man5openssl`, and `/usr/share/man/man7openssl` directories.

To view the license terms, attribution, and copyright for OpenSSL, run `pkg info -l license library/security/openssl`.

**Examples** **EXAMPLE 1** Generating and Printing a Public Key

The following example generates and prints a public key stored in an already initialized PKCS#11 keystore. Notice the use of `-engine pkcs11` and `-inform e`.

```
$ pktool gencert keystore=pkcs11 label=mykey \
  subject="CN=test" keytype=rsa keylen=1024 serial=01
$ openssl rsa -in "pkcs11:object=mykey;passphrasedialog=builtin" \
  -pubout -text -engine pkcs11 -inform e
```

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	library/security/openssl, library/security/openssl
Interface Stability	Volatile

**See Also** [crle\(1\)](#), [cryptoadm\(1M\)](#), [libpkcs11\(3LIB\)](#), [attributes\(5\)](#), [privileges\(5\)](#)

[/usr/share/man/man1openssl/openssl.1openssl](#),  
[/usr/share/man/man1openssl/CRYPTO\\_num\\_locks.3openssl](#),  
[/usr/share/man/man3openssl/engine.3](#), [/usr/share/man/man3openssl/evp.3](#)

**Name** pam\_allow – PAM authentication, account, session and password management PAM module to allow operations

**Synopsis** pam\_allow.so.1

**Description** The pam\_allow module implements all the PAM service module functions and returns PAM\_SUCCESS for all calls. Opposite functionality is available in the [pam\\_deny\(5\)](#) module.

Proper Solaris authentication operation requires [pam\\_unix\\_cred\(5\)](#) be stacked above pam\_allow.

The following options are interpreted:

debug Provides [syslog\(3C\)](#) debugging information at the LOG\_AUTH | LOG\_DEBUG level.

**Errors** PAM\_SUCCESS is always returned.

**Examples** EXAMPLE 1 Allowing ssh none

The following example is a pam.conf fragment that illustrates how to allow the SSHv2 userauth of “none”:

```
sshd-none auth required pam_unix_cred.so.1
sshd-none auth sufficient pam_allow.so.1
sshd-none account sufficient pam_allow.so.1
sshd-none session sufficient pam_allow.so.1
sshd-none password sufficient pam_allow.so.1
```

The equivalent configuration using /etc/pam.d/ would be the following entries in /etc/pam.d/sshd-none:

```
auth required pam_unix_cred.so.1
auth sufficient pam_allow.so.1
account sufficient pam_allow.so.1
session sufficient pam_allow.so.1
password sufficient pam_allow.so.1
```

EXAMPLE 2 Allowing Kiosk Automatic Login Service

The following example is a pam.conf fragment that illustrates how to allow gdm kiosk automatic login:

```
gdm-autologin auth required pam_unix_cred.so.1
gdm-autologin auth sufficient pam_allow.so.1
```

The equivalent configuration using /etc/pam.d/ would be the following entries in /etc/pam.d/gdm-autologin:

```
auth required pam_unix_cred.so.1
auth sufficient pam_allow.so.1
```

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT Level	MT-Safe with exceptions

**See Also** [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam\\_sm\(3PAM\)](#), [syslog\(3C\)](#), [pam.conf\(4\)](#), [attributes\(5\)](#), [pam\\_deny\(5\)](#), [pam\\_unix\\_cred\(5\)](#)

**Notes** The interfaces in [libpam\(3LIB\)](#) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

This module is intended to be used to either allow access to specific services names, or to all service names not specified (by specifying it as the default service stack).

**Name** pam\_authtok\_check – authentication and password management module

**Synopsis** pam\_authtok\_check.so.1

**Description** 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:

length	The password length should not be less than the minimum specified in /etc/default/passwd.
circular shift	The password should not be a circular shift of the login name. This check may be disabled in /etc/default/passwd.
complexity	The password should contain at least the minimum number of characters described by the parameters MINALPHA, MINNONALPHA, MINDIGIT, and MINSPECIAL. Note that MINNONALPHA describes the same character classes as MINDIGIT and MINSPECIAL combined; therefore the user cannot specify both MINNONALPHA and MINSPECIAL (or MINDIGIT). The user must choose which of the two options to use. Furthermore, the WHITESPACE parameter determines whether whitespace characters are allowed. If unspecified MINALPHA is 2, MINNONALPHA is 1 and WHITESPACE is yes
variation	The old and new passwords must differ by at least the MINDIFF value specified in /etc/default/passwd. If unspecified, the default is 3. For accounts in name services which support password history checking, if prior history is defined, the new password must not match the prior passwords.
dictionary check	The password must not be based on a dictionary word. The list of words to be used for the site's dictionary can be specified with DICTIONLIST. It should contain a comma-separated list of filenames, one word per line. The database that is created from these files is stored in the directory named by DICTIONDBDIR (defaults to /var/passwd). See <a href="#">mkpwdict(1M)</a> for information on pre-generating the database. If neither DICTIONLIST nor DICTIONDBDIR is specified, no dictionary check is made.
upper/lower case	The password must contain at least the minimum of upper- and lower-case letters specified by the MINUPPER and MINLOWER values in /etc/default/passwd. If unspecified, the defaults are 0.

**maximum repeats** The password must not contain more consecutively repeating characters than specified by the `MAXREPEATS` value in `/etc/default/passwd`. If unspecified, no repeat character check is made.

The following option may be passed to the module:

**force\_check** If the `PAM_NO_AUTH Tok_CHECK` flag set, `force_check` ignores this flag. The `PAM_NO_AUTH Tok_CHECK` flag can be set to bypass password checks (see [pam\\_chauthtok\(3PAM\)](#)).

**server\_policy** If the account authority for the user, as specified by `PAM_USER`, is not files or NIS, and if `server_policy` is specified, this module does not perform any password-strength checks. Instead, it leaves it to the account authority to validate the new password against its own set of rules.

**debug** [syslog\(3C\)](#) debugging information at the `LOG_DEBUG` level

**Return Values** 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` See [passwd\(1\)](#) for a description of the contents.

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT Level	MT-Safe with exceptions

**See Also** [passwd\(1\)](#), [pam\(3PAM\)](#), [mkpwdict\(1M\)](#), [pam\\_chauthtok\(3PAM\)](#), [syslog\(3C\)](#), [libpam\(3LIB\)](#), [pam.conf\(4\)](#), [passwd\(4\)](#), [shadow\(4\)](#), [attributes\(5\)](#), [pam\\_authtok\\_get\(5\)](#), [pam\\_authtok\\_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.

**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** 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.

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.

If a previous module has set PAM\_OLDAUTHTOK prior to the invocation of pam\_authtok\_get, this module turns into a NO-OP and immediately returns PAM\_SUCCESS.

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.

The following option can be passed to the module:

debug     [syslog\(3C\)](#) debugging information at the LOG\_DEBUG level

**Errors** The authentication service returns the following error codes:

PAM\_SUCCESS       Successfully obtains authentication token

PAM\_SYSTEM\_ERR    Fails to retrieve username, username is NULL or empty

The password management service returns the following error codes:

PAM\_SUCCESS       Successfully obtains authentication token

PAM\_AUTHTOK\_ERR   Authentication token manipulation error

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
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\\_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.

**Name** pam\_authtok\_store – password management module

**Synopsis** pam\_authtok\_store.so.1

**Description** pam\_authtok\_store provides functionality to the PAM password management stack. It provides one function: pam\_sm\_chauthtok().

When invoked with flags set to PAM\_UPDATE\_AUTHTOK, this module updates the authentication token for the user specified by PAM\_USER.

The authentication token PAM\_OLDAUTHTOK can be used to authenticate the user against repositories that need updating (NIS, LDAP). After successful updates, the new authentication token stored in PAM\_AUTHTOK is the user's valid password.

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\(4\)](#).

The following option can be passed to the module:

`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** 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:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
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\\_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.

If the PAM\_REPOSITORY *item\_type* is set and a service module does not recognize the type, the service module does not process any information, and returns PAM\_IGNORE. If the PAM\_REPOSITORY *item\_type* is not set, a service module performs its default action.

**Name** pam\_deny – PAM authentication, account, session and password management PAM module to deny operations

**Synopsis** pam\_deny.so.1

**Description** The pam\_deny module implements all the PAM service module functions and returns the module type default failure return code for all calls.

The following options are interpreted:

debug `syslog(3C)` debugging information at the LOG\_AUTH|LOG\_DEBUG levels

**Errors** The following error codes are returned:

PAM\_ACCT\_EXPIRED If pam\_sm\_acct\_mgmt is called.

PAM\_AUTH\_ERR If pam\_sm\_authenticate is called.

PAM\_AUTHOK\_ERR If pam\_sm\_chauthtok is called.

PAM\_CRED\_ERR If pam\_sm\_setcred is called.

PAM\_SESSION\_ERR If pam\_sm\_open\_session or pam\_sm\_close\_session is called.

**Examples** EXAMPLE 1 Disallowing ssh none authentication

The following example is a pam.conf fragment that illustrates how to deny the SSHv2 userauth of “none”:

```
sshd-none    auth      requisite pam_deny.so.1
sshd-none    account  requisite pam_deny.so.1
sshd-none    session  requisite pam_deny.so.1
sshd-none    password requisite pam_deny.so.1
```

The equivalent configuration in /etc/pam.d/ would be the following entries in /etc/pam.d/sshd-none:

```
auth      requisite pam_deny.so.1
account  requisite pam_deny.so.1
session  requisite pam_deny.so.1
password requisite pam_deny.so.1
```

EXAMPLE 2 Disallowing any service not explicitly defined

The following example is a pam.conf fragment that illustrates how to deny any PAM service which is not explicitly defined in the PAM configuration:

```
other      auth      requisite pam_deny.so.1
other      account  requisite pam_deny.so.1
other      session  requisite pam_deny.so.1
other      password requisite pam_deny.so.1
```

**EXAMPLE 2** Disallowing any service not explicitly defined *(Continued)*

The equivalent configuration in `/etc/pam.d/` would be the following entries in `/etc/pam.d/other`:

```
auth      requisite  pam_deny.so.1
account   requisite  pam_deny.so.1
session   requisite  pam_deny.so.1
password  requisite  pam_deny.so.1
```

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT-Level	MT-Safe with exceptions

**See Also** [su\(1M\)](#), [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam\\_sm\\_authenticate\(3PAM\)](#), [syslog\(3C\)](#), [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\\_account\(5\)](#), [pam\\_unix\\_auth\(5\)](#), [pam\\_unix\\_session\(5\)](#), [privileges\(5\)](#)

**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_deny` module is intended to deny access to a specified service. The other service name may be used to deny access to services not explicitly specified.

**Name** pam\_dhkeys – authentication Diffie-Hellman keys management module

**Synopsis** pam\_dhkeys . so . 1

**Description** The pam\_dhkeys . so . 1 service module provides functionality to two PAM services: Secure RPC authentication and Secure RPC authentication token management.

Secure RPC authentication differs from regular Unix authentication because ONC RPCs use Secure RPC as the underlying security mechanism.

The following options may be passed to the module:

debug      `syslog(3C)` debugging information at LOG\_DEBUG level  
nowarn      Turn off warning messages

**Authentication Services** 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\_AUTHTOK item. 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)`.

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)`.

PAM\_REINITIALIZE\_CRED and PAM\_REFRESH\_CRED are not supported and return PAM\_IGNORE.

**Authentication Token Management** 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.

**Errors** The authentication service returns the following error codes:

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_AUTHTOK is not set, no credentials are present or there is a wrong password.
PAM_BUF_ERR	Module ran out of memory.

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	Committed
MT Level	MT-Safe with exceptions

**See Also** [keylogin\(1\)](#), [keylogout\(1\)](#), [pam\(3PAM\)](#), [pam\\_authenticate\(3PAM\)](#), [pam\\_chauthtok\(3PAM\)](#), [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\\_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.

**Name** pam\_dial\_auth – authentication management PAM module for dialups

**Synopsis** pam\_dial\_auth.so.1

**Description** The pam\_dial\_auth module implements [pam\\_sm\\_authenticate\(3PAM\)](#) which authenticates the user according to the [dialups\(4\)](#) and [d\\_passwd\(4\)](#) files configuration.

Authentication service modules must implement both `pam_sm_authenticate()` and `pam_sm_setcred()`. `pam_sm_setcred()` in this module always returns `PAM_IGNORE`.

The value of the `PAM_TTY` item is checked against entries in [dialups\(4\)](#). If there is a match, the user's shell is compared against entries in [d\\_passwd\(4\)](#). If there is a matching entry, the user is prompted for a password which is validated against the entry found.

The following option may be passed in to this service module:

`debug`     [syslog\(3C\)](#) debugging information at `LOG_DEBUG` level.

**Errors** If [dialups\(4\)](#) is not present, `PAM_IGNORE` is returned. Upon successful completion of `pam_sm_authenticate()`, `PAM_SUCCESS` is returned. The following error codes are returned upon error:

<code>PAM_AUTH_ERR</code>	Authentication failure.
<code>PAM_SERVICE_ERR</code>	Error in the calling service, <code>PAM_TTY</code> is not set.
<code>PAM_SYSTEM_ERR</code>	System error ( <a href="#">d_passwd(4)</a> is not present).
<code>PAM_USER_UNKNOWN</code>	No account is present for <i>user</i> .

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	MT-Safe with exceptions

**See Also** [pam\(3PAM\)](#), [pam\\_authenticate\(3PAM\)](#), [pam\\_sm\\_authenticate\(3PAM\)](#), [d\\_passwd\(4\)](#), [dialups\(4\)](#), [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.

**Name** pam\_krb5 – authentication, account, session, and password management PAM modules for Kerberos V5

**Synopsis** /usr/lib/security/pam\_krb5.so.1

**Description** The Kerberos V5 service module for PAM provides functionality for all four PAM modules: authentication, account management, session management, and password management. The service 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.

**Kerberos Authentication Module** The Kerberos V5 authentication component provides functions to verify the identity of a user, `pam_sm_authenticate()`, and to manage the Kerberos credentials cache, `pam_sm_setcred()`.

`pam_sm_authenticate()` 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 service module, which then verifies that the TGT came from a valid Key Distribution Center (KDC) by attempting to get a service ticket for the local host service. For this to succeed, the local host's keytab file (`/etc/krb5/krb5.keytab`) must contain the entry for the local host service. For example, in the file `host/hostname.com@REALM`, `hostname.com` is the fully qualified local hostname and `REALM` is the default realm of the local host as defined in `/etc/krb5/krb5.conf`. If the host entry is not found in the keytab file, the authentication fails. Administrators can optionally disable this “strict” verification by setting “`verify_ap_req_nofail = false`” in `/etc/krb5/krb5.conf`. See [krb5.conf\(4\)](#) for more details on this option. This allows TGT verification to succeed in the absence of a keytab host principal entry.

If `pam_sm_authenticate()` is called and the `pkinit` module option is set, the Kerberos V5 authentication module tries to do PKINIT authentication, assuming that both the system and the KDC are configured to support this type of authentication. This form of authentication uses a user's certificate and private key to acquire the user's initial Kerberos credential (TGT). One of the keystore formats supported is PKCS11 which supports use of any PKCS11 compatible keystore capable of storing the required credential and private key needed for PKINIT authentication (PKCS11 compatible smartcards are an example). See [krb5.conf\(4\)](#) for more details on PKINIT configuration. This form of authentication is typically useful for services where the system on which the auth stack is being processed has access to the user's certificate and private key.

If `pam_sm_authenticate()` is called and the `pkinit` module option is not set then the Kerberos V5 authentication module does password based authentication.

In either case, if the `PAM_AUTHTOK` password item has been set when `pam_sm_authenticate()` is called, which is the case when `pam_krb5` is stacked after `pam_authtok_get` in the auth stack, the Kerberos V5 authentication module uses that `PAM_AUTHTOK` password for either PKINIT or password based Kerberos authentication.

If the `PAM_USER` item is not set `pam_krb5` with the `pkinit` option prompts for and set that item.

If the `PAM_AUTHTOK` password item has not been set when `pam_sm_authenticate()` is called, which is the case when `pam_krb5` is stacked before `pam_authtok_get` in the auth stack, and the `pkinit` option is present the Kerberos V5 authentication module allows the Kerberos `pkinit` preauth plugin to prompt for whatever information is needed to perform PKINIT (typically this is for the user's PIN). No PAM items are set by way of this prompting. See [krb5.conf\(4\)](#) for more information on PKINIT configuration options.

The `pam_krb5` module sets the `KRB5CCNAME` shell environment variable upon successful authentication or password change to `FILE:/tmp/krb5cc_uid` where `uid` is the UID of the user that `pam_krb5` authenticated. `KRB5CCNAME` is documented in [krb5envvar\(5\)](#).

If it is desirable to initially have the Kerberos V5 authentication module try PKINIT Kerberos authentication and fall back to password based Kerberos authentication then either the sufficient or optional control flags must be provided for the instance of `pam_krb5` with the `pkinit` module option set and another instance of `pam_krb5` without the `pkinit` module option must be stacked below `pam_authtok_get`. If there are PAM modules other than `pam_krb5` that must be evaluated below `pam_authtok_get` then the control flag should be set to optional for the instance of `pam_krb5` with the `pkinit` module option set otherwise the control flag should be set to sufficient.

Only two instances of `pam_krb5` are supported in a auth stack.

`pam_sm_authenticate(3PAM)` can be passed the following flag:

**PAM\_DISALLOW\_NULL\_AUTHTOK**

This flag is ignored. The Kerberos authentication mechanism does not allow an empty password string by default.

`pam_sm_setcred()` creates and modifies the user's credential cache. This function initializes the user's credential cache, if it does not already exist, and stores the initial credentials for later use by Kerberos network applications. The following flags can be set in the flags field. They are best described by their effect on the user's credential cache.

**PAM\_ESTABLISH\_CRED**

Stores the initial credentials in the user's credential cache so that the user can access Kerberos network services. If a successful authentication pass was made, the new credentials are stored in the credential cache, overwriting any existing credentials that were previously stored. If an unsuccessful authentication pass was made, `PAM_CRED_UNAVAIL` is returned.

**PAM\_DELETE\_CRED**

This flag has no effect on the credential cache and always returns `PAM_SUCCESS`. The credential cache is not deleted because there is no accurate method to determine if the credentials are needed by another process. The credential cache can be deleted with the `kdestroy(1)` command.

**PAM\_REINITIALIZE\_CRED**

Deletes the user's existing credential cache, if it exists, and creates a new credential cache. The new credentials are stored in the new cache and the user's ticket lifetime and renewable life time values are reset.

**PAM\_REFRESH\_CRED**

Does not require a previous authentication pass, but if a successful one is made, the new credentials are stored in the credential cache. If a previous authentication pass was not made or was unsuccessful, an attempt to renew the existing credentials is made. This function fails if the user's renewable ticket lifetime is expired.

The following options can be passed to the Kerberos V5 authentication module:

- debug** Provides `syslog(3C)` debugging information at `LOG_DEBUG` level.
- nowarn** Turns off warning messages.
- pkinit** Indicates that the Kerberos V5 authentication module should try Kerberos PKINIT authentication instead of the default password based Kerberos authentication.

**Kerberos V5 Account Management Module**

The Kerberos account management component provides a function to perform account management, `pam_sm_acct_mgmt()`. This function checks to see if the `pam_krb5` authentication module has noted that the user's password has not expired. The following options can be passed in to the Kerberos V5 account management module:

- debug** Provides `syslog(3C)` debugging information at `LOG_DEBUG` level
- nowarn** Turns off warning messages. Also, does not query KDC for impending password expiration information used to warn the user.

**Kerberos V5 Session Management Module**

The Kerberos V5 session management component provides functions to initiate `pam_sm_open_session()` and terminate `pam_sm_close_session()` Kerberos sessions. For Kerberos V5, both `pam_sm_open_session` and `pam_sm_close_session()` are null functions, returning `PAM_IGNORE`.

**Kerberos V5 Password Management Module**

The Kerberos V5 password management component provides a function to change passwords, `pam_sm_chauthtok()`, in the Key Distribution Center (KDC) database.

If the Kerberos V5 authentication module used PKINIT authentication in the auth stack then the Kerberos V5 password management module returns `PAM_IGNORE` in the following cases:

- The new password is `NULL`.
- The old password is `NULL`.
- Verification of the old password fails.

The rationale behind this is that the KDC can not allow a PKINIT user to change/set a password since the user can be expected to use PKINIT only. If all of the cases above are false the Kerberos V5 password management module tries to change the user's password in the KDC database.

If the KDC only supports PKINIT authentication then the Kerberos V5 password management module should not be present in any password stacks.

Related to PKINIT the Kerberos V5 password management module does not support changing the key store PIN used to access a user's private key and certificate.

The following flags can be passed to `pam_sm_chauthtok(3PAM)`:

`PAM_CHANGE_EXPIRED_AUTH Tok`

The password service should only update the user's Kerberos password if it is expired.

Otherwise, this function returns `PAM_IGNORE`. The default behaviour is to always change the user's Kerberos password.

`PAM_PRELIM_CHECK`

This is a null function that always returns `PAM_IGNORE`.

`PAM_UPDATE_AUTH Tok`

This flag is necessary to change the user's Kerberos password. If this flag is not set, `pam_krb5` returns `PAM_SYSTEM_ERR`.

The following option can be passed to the Kerberos V5 password module:

`debug` Provides `syslog(3C)` debugging information at `LOG_DEBUG` level.

**Errors** The following error codes are returned for `pam_sm_authenticate()`:

<code>PAM_AUTH_ERR</code>	Authentication failure
<code>PAM_BUF_ERR</code>	Memory buffer error.
<code>PAM_IGNORE</code>	The user is "root" and the root key exists in the default keytab.
<code>PAM_SUCCESS</code>	Successfully obtained Kerberos credentials .
<code>PAM_SYSTEM_ERR</code>	System error.
<code>PAM_USER_UNKNOWN</code>	An unknown Kerberos principal was requested.

The following error codes are returned for `pam_sm_setcred()`:

<code>PAM_AUTH_ERR</code>	Authentication failure.
<code>PAM_BUF_ERR</code>	Memory buffer error.
<code>PAM_IGNORE</code>	The user is "root" and the root key exists in the default keytab.
<code>PAM_SYSTEM_ERR</code>	System error.
<code>PAM_SUCCESS</code>	Successfully modified the Kerberos credential cache.

The following error codes are returned for `pam_sm_acct_mgmt()`:

<code>PAM_AUTH_ERR</code>	Authentication failure.
---------------------------	-------------------------

PAM_IGNORE	Kerberos service module <code>pam_sm_authenticate()</code> was never called, or the user is “root” and the root key exists in the default keytab.
PAM_NEW_AUTHTOK_REQD	Obtain new authentication token from the user.
PAM_SERVICE_ERR	Error in underlying service module.
PAM_SUCCESS	Kerberos principal account is valid.
PAM_SYSTEM_ERR	System error.
PAM_USER_UNKNOWN	An unknown Kerberos principal was requested.

The following error code is returned for `pam_sm_open_session()` and `pam_sm_close_session()`:

PAM\_IGNORE These two functions are null functions in `pam_krb5`:

The following error codes are returned for `pam_sm_chauthtok()`:

PAM_AUTH_ERR	Authentication failure.
PAM_IGNORE	The user has not been authenticated by Kerberos service module <code>pam_sm_authenticate()</code> , or the user is “root” and the root key exists in the default keytab.
PAM_NEW_AUTHTOK_REQD	User's Kerberos password has expired.
PAM_SERVICE_ERR	Error in module. At least one input parameter is missing.
PAM_SYSTEM_ERR	System error.
PAM_USER_UNKNOWN	An unknown Kerberos principal was requested.
PAM_SUCCESS	Successfully changed the user's Kerberos password.

### Examples **EXAMPLE 1** Authenticating Users Through Kerberos as First Choice Using Password-based Authentication

The following is an excerpt of a sample `pam.conf` configuration file that authenticates users through the Kerberos authentication service and authenticates through the Unix login only if the Kerberos authentication fails. This arrangement is helpful when a majority of the users are networked by means of Kerberos and when there are only a few non-Kerberos type user accounts, such as root. The service illustrated below is for `gdm`.

```
gdm auth requisite      pam_authtok_get.so.1
gdm auth required      pam_dhkeys.so.1
gdm auth required      pam_unix_cred.so.1
gdm auth sufficient    pam_krb5.so.1
gdm auth required      pam_unix_auth.so.1
```

**EXAMPLE 1** Authenticating Users Through Kerberos as First Choice Using Password-based Authentication *(Continued)*

These changes should not be made to the existing `krlogin`, `krsh`, and `ktelnet` service entries. Those services require Kerberos authentication, so using a seemingly sufficient control flag would not provide the necessary functionality for privacy and integrity. There should be no need to change those entries.

The following entries check for password expiration when dealing with Kerberos and Unix password aging policies:

```
other account requisite    pam_roles.so.1
other account required     pam_unix_account.so.1
other account required     pam_krb5.so.1
```

The following entries would change the Kerberos password of the user and continue to change the Unix login password only if the Kerberos password change had failed:

```
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 sufficient  pam_krb5.so.1
other password required    pam_authtok_store.so.1
```

**EXAMPLE 2** Authenticating Users Through Kerberos Only Using Password-based Authentication

The following example allows authentication only to users that have Kerberos-based accounts.

```
gdm auth requisite        pam_authtok_get.so.1
gdm auth required         pam_dhkeys.so.1
gdm auth required         pam_unix_cred.so.1
gdm auth required         pam_krb5.so.1
```

Typically, you would have another service specified in the `pam.conf` file that would allow local users, such as database, web server, system administrator accounts, to log in to the host machine. For example, the service name “login” could be used for these users. These users should not belong to any roles.

The rest of the module types look similar to that shown in the previous example:

```
other account requisite    pam_roles.so.1
other account required     pam_unix_account.so.1
other account required     pam_krb5.so.1
```

With binding specified in the following, it is important that non-Kerberos users specify the repository in which they reside using the `-r` option with the `passwd(1)` command. This configuration is also based on the assumptions that:

- Kerberos users maintain only their Kerberos passwords;

**EXAMPLE 2** Authenticating Users Through Kerberos Only Using Password-based Authentication  
(Continued)

- changing their Unix password is not necessary, given that they are authenticated only through their Kerberos passwords when logging in.

```
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 binding       pam_krb5.so.1
```

**EXAMPLE 3** Authenticating Through Kerberos Optionally Using Password-based Authentication

This configuration is helpful when the majority of users are non-Kerberos users and would like to authenticate through Kerberos if they happened to exist in the Kerberos database. The effect of this is similar to users voluntarily executing `kinit(1)` after they have successfully logged in:

```
gdm auth requisite          pam_authtok_get.so.1
gdm auth required          pam_dhkeys.so.1
gdm auth required          pam_unix_cred.so.1
gdm auth required          pam_unix_auth.so.1
gdm auth optional          pam_krb5.so.1
```

The rest of the configuration is as follows:

```
other account requisite     pam_roles.so.1
other account required      pam_unix_account.so.1
other account optional      pam_krb5.so.1

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
other password optional     pam_krb5.so.1
```

Non-Kerberos users should specify their respective repositories by using the `-r` option when changing their password with the `passwd(1)` command.

**EXAMPLE 4** Authenticating Users Through Kerberos PKINIT as First Choice

The following is an excerpt of a sample `pam.conf` configuration file that authenticates users through the Kerberos authentication service and authenticates through the Unix login only if the Kerberos authentication (using PKINIT) fails. This arrangement is helpful when a majority of the users are networked by means of Kerberos and when there are only a few non-Kerberos type user accounts, such as `root`. The service illustrated below is for `login`. The user is prompted once for the PIN by `pam_krb5`.

**EXAMPLE 4** Authenticating Users Through Kerberos PKINIT as First Choice *(Continued)*

```
login auth required      pam_unix_cred.so.1
login auth sufficient    pam_krb5.so.1 pkinit
login auth requisite     pam_authtok_get.so.1
login auth required     pam_dhkeys.so.1
login auth required     pam_unix_auth.so.1
```

**EXAMPLE 5** Authenticating Users Through Kerberos PKINIT Only

The following example allows authentication only to users that have kerberos-based accounts requiring PKINIT authentication.

```
login auth required      pam_unix_cred.so.1
login auth required     pam_krb5.so.1 pkinit
```

**EXAMPLE 6** Authenticating Users Through Kerberos PKINIT Optionally

The following example allows users to acquire a Kerberos credential using PKINIT authentication if they have a Kerberos account. Whether pam\_krb5 succeeds or fails the user must provide their Unix password to login.

```
login auth required      pam_unix_cred.so.1
login auth optional     pam_krb5.so.1 pkinit
login auth requisite     pam_authtok_get.so.1
login auth required     pam_unix_auth.so.1
```

**EXAMPLE 7** Authenticating Users Through Kerberos PKINIT as a Requirement

The following example allows users to login if pam\_krb5 is able to acquire a Kerberos credential using PKINIT authentication and in addition must provide their Unix password to pam\_unix\_auth.

```
login auth required      pam_unix_cred.so.1
login auth required     pam_krb5.so.1 pkinit
login auth requisite     pam_authtok_get.so.1
login auth required     pam_unix_auth.so.1
```

**EXAMPLE 8** Authenticating Users Through Kerberos PKINIT as a Requirement

The following example allows users to login using their PAM\_AUTHTOK password acquired by pam\_authtok\_get. This password is used by pam\_krb5 to try PKINIT authentication and is also used by pam\_unix\_auth to authenticate the user using the user's Unix account. If PKINIT requires a password/PIN that differs from the user's Unix password then pam\_krb5 must be stacked above pam\_authtok\_get.

```
login auth required      pam_unix_cred.so.1
login auth requisite     pam_authtok_get.so.1
```

**EXAMPLE 8** Authenticating Users Through Kerberos PKINIT as a Requirement *(Continued)*

```
login auth required      pam_krb5.so.1 pkinit
login auth required      pam_unix_auth.so.1
```

**EXAMPLE 9** Authenticating Users Through Kerberos PKINIT with a Fall Back to Password-based krb auth

The following example allows users to acquire a Kerberos credential using PKINIT authentication or using password based authentication if PKINIT fails. If PKINIT succeeds the user is not prompted for their password. If pam\_krb5 PKINIT succeeds, the second instance of pam\_krb5 does not try password authentication and returns success. If PKINIT fails the user is prompted for their Kerberos password.

```
login auth required      pam_unix_cred.so.1
login auth sufficient     pam_krb5.so.1 pkinit
login auth requisite      pam_authtok_get.so.1
login auth required      pam_krb5.so.1
```

**EXAMPLE 10** Authenticating Users Through Kerberos Requiring Users to Authenticate Either through Kerberos PKINIT or Fall Back to Password-based krb auth

The following example allows users to acquire a Kerberos credential using PKINIT authentication or using password based authentication if PKINIT fails. If pam\_krb5 PKINIT succeeds, the second instance of pam\_krb5 does not try password authentication and returns ignore. If pam\_krb5 PKINIT fails the second instance of pam\_krb5 tries password based authentication and return success or failure.

```
login auth required      pam_unix_cred.so.1
login auth optional       pam_krb5.so.1 pkinit
login auth requisite      pam_authtok_get.so.1
login auth required       pam_krb5.so.1
login auth required       pam_dhkeys.so.1
login auth required       pam_unix_auth.so.1
```

**EXAMPLE 11** Authenticating Users Through Kerberos Requiring Users to Authenticate Either through Kerberos PKINIT or Fall Back to pam\_pkcs11

The following example allows users to acquire a Kerberos credential using PKINIT authentication or if that fails use pam\_pkcs11 to validate the user's PIN using their certificate and private key.

```
login auth required      pam_unix_cred.so.1
login auth sufficient     pam_krb5.so.1 pkinit
login auth sufficient     pam_pkcs11.so
```

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

**See Also** [kdestroy\(1\)](#), [kinit\(1\)](#), [passwd\(1\)](#), [kttkt\\_warnd\(1M\)](#), [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam\\_sm\(3PAM\)](#), [pam\\_sm\\_acct\\_mgmt\(3PAM\)](#), [pam\\_sm\\_authenticate\(3PAM\)](#), [pam\\_sm\\_chauthtok\(3PAM\)](#), [pam\\_sm\\_close\\_session\(3PAM\)](#), [pam\\_sm\\_open\\_session\(3PAM\)](#), [pam\\_sm\\_setcred\(3PAM\)](#), [syslog\(3C\)](#), [krb5.conf\(4\)](#), [pam.conf\(4\)](#), [attributes\(5\)](#), [kerberos\(5\)](#), [krb5envvar\(5\)](#), [pam\\_krb5\\_migrate\(5\)](#)

**Notes** The interfaces in [libpam\(3LIB\)](#) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

On successful acquisition of initial credentials (ticket-granting ticket), [kttkt\\_warnd\(1M\)](#) is notified, to alert the user when the initial credentials are about to expire.

**Name** pam\_krb5\_migrate – authentication PAM module for the KerberosV5 auto-migration of users feature

**Synopsis** /usr/lib/security/pam\_krb5\_migrate.so.1

**Description** The KerberosV5 auto-migrate service module for PAM provides functionality for the PAM authentication component. The service module helps in the automatic migration of PAM\_USER to the client's local Kerberos realm, using PAM\_AUTHTOK (the PAM authentication token associated with PAM\_USER) as the new Kerberos principal's password.

KerberosV5  
Auto-migrate  
Authentication Module

The KerberosV5 auto-migrate authentication component provides the `pam_sm_authenticate(3PAM)` function to migrate a user who does not have a corresponding krb5 principal account to the default Kerberos realm of the client.

`pam_sm_authenticate(3PAM)` uses a host-based client service principal, present in the local keytab (/etc/krb5/krb5.keytab) to authenticate to `kadmind(1M)` (defaults to the `host/nodename.fqdn` service principal), for the principal creation operation. Also, for successful creation of the krb5 user principal account, the host-based client service principal being used needs to be assigned the appropriate privilege on the master KDC's `kadm5.acl(4)` file. `kadmind(1M)` checks for the appropriate privilege and validates the user password using PAM by calling `pam_authenticate(3PAM)` and `pam_acct_mgmt(3PAM)` for the `k5migrate` service.

If migration of the user to the KerberosV5 infrastructure is successful, the module will inform users about it by means of a PAM\_TEXT\_INFO message, unless instructed otherwise by the presence of the quiet option.

The authentication component always returns PAM\_IGNORE and is meant to be stacked in the PAM configuration (see `pam.conf(4)`) with a requirement that it be listed below `pam_authtok_get(5)` in the authentication stack. Also, if `pam_krb5_migrate` is used in the authentication stack of a particular service, it is mandatory that `pam_krb5(5)` be listed in the PAM account stack of that service for proper operation (see EXAMPLES).

**Options** The following options can be passed to the KerberosV5 auto-migrate authentication module:

<code>debug</code>	Provides <code>syslog(3C)</code> debugging information at LOG_DEBUG level.
<code>client_service=&lt;service name&gt;</code>	Name of the service used to authenticate to <code>kadmind(1M)</code> defaults to <code>host</code> . This means that the module uses <code>host/&lt;nodename.fqdn&gt;</code> as its client service principal name, KerberosV5 user principal creation operation or <code>&lt;service&gt;/&lt;nodename.fqdn&gt;</code> if this option is provided.
<code>quiet</code>	Do not explain KerberosV5 migration to the user.

This has the same effect as passing the `PAM_SILENT` flag to `pam_sm_authenticate(3PAM)` and is useful where applications cannot handle `PAM_TEXT_INFO` messages.

If not set, the authentication component will issue a `PAM_TEXT_INFO` message after creation of the Kerberos V5 principal, indicating that it has done so.

`expire_pw`

Causes the creation of Kerberos V5 user principals with password expiration set to now (current time).

**Examples** EXAMPLE 1 Sample PAM entries for the `pam_krb5_migrate.so.1` module

The following `pam.conf(4)` entries demonstrate the use of the `pam_krb5_migrate.so.1` module:

```
login      auth requisite      pam_authtok_get.so.1
login      auth required       pam_dhkeys.so.1
login      auth required       pam_unix_cred.so.1
login      auth sufficient     pam_krb5.so.1
login      auth requisite     pam_unix_auth.so.1
login      auth optional      pam_krb5_migrate.so.1 expire_pw
login      auth required     pam_dial_auth.so.1

other     account requisite   pam_roles.so.1
other     account required   pam_krb5.so.1
other     account required   pam_unix_account.so.1
```

The equivalent PAM configuration in `/etc/pam.d/` would be the following entries in `/etc/pam.d/login`:

```
auth      requisite      pam_authtok_get.so.1
auth      required       pam_dhkeys.so.1
auth      required       pam_unix_cred.so.1
auth      sufficient     pam_krb5.so.1
auth      requisite     pam_unix_auth.so.1
auth      optional      pam_krb5_migrate.so.1 expire_pw
auth      required     pam_dial_auth.so.1
```

and the following entries in `/etc/pam.d/other`:

```
account  requisite      pam_roles.so.1
account  required       pam_krb5.so.1
account  required       pam_unix_account.so.1
```

The `pam_krb5_migrate` module can generally be present on the authentication stack of any service where the application calls `pam_sm_authenticate(3PAM)` and an authentication token (in the preceding example, the authentication token would be the user's Unix password) is available for use as a Kerberos V5 password.

**EXAMPLE 2** Sample Entries from `kadm5.acl`

The following entries from `kadm5.acl(4)` permit or deny privileges to the host client service principal:

```
host/*@ACME.COM U root
host/*@ACME.COM ui *
```

The preceding entries permit the `pam_krb5_migrate` add privilege to the host client service principal of any machine in the `ACME.COM` KerberosV5 realm, but denies the add privilege to all host service principals for addition of the root user account.

**EXAMPLE 3** Sample PAM entries for the Master KDC

The `/etc/pam.conf` entries below enable `kadmind(1M)` on the master KDC to use the `k5migrate` PAM service in order to validate Unix user passwords for accounts that require migration to the Kerberos realm.

```
k5migrate      auth    required      pam_unix_auth.so.1
k5migrate      account required    pam_unix_account.so.1
```

The equivalent PAM configuration in `/etc/pam.d/` would be the following entries in `/etc/pam.d/k5migrate`:

```
auth    required      pam_unix_auth.so.1
account required    pam_unix_account.so.1
```

**Attributes** See [attributes\(5\)](#) for a description of the following attribute:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

**See Also** [kadmind\(1M\)](#), [syslog\(3C\)](#), [pam\\_authenticate\(3PAM\)](#), [pam\\_acct\\_mgmt\(3PAM\)](#), [pam\\_sm\\_authenticate\(3PAM\)](#), [kadm5.acl\(4\)](#), [pam.conf\(4\)](#), [attributes\(5\)](#), [pam\\_authtok\\_get\(5\)](#), [pam\\_krb5\(5\)](#)

**Name** pam\_ldap – authentication and account management PAM module for LDAP

**Synopsis** /usr/lib/security/pam\_ldap.so.1

**Description** The pam\_ldap module implements [pam\\_sm\\_authenticate\(3PAM\)](#) and [pam\\_sm\\_acct\\_mgmt\(3PAM\)](#), the functions that provide functionality for the PAM authentication and account management stacks. The pam\_ldap module ties the authentication and account management 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 component of pam\_ldap will work only with the Sun Java System Directory Server. The server's user account management must be properly configured before it can be used by pam\_ldap. Refer to the *Sun Java System Directory Server Administration Guide* for information on how to configure user account management, including password and account lockout policy.

pam\_ldap must be used in conjunction with the modules that support the UNIX authentication, password, and account management, which are [pam\\_authok\\_get\(5\)](#), [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 to control network accounts, that is, LDAP users. For the stacks to work, pam\_unix\_auth, pam\_unix\_account, and pam\_passwd\_auth must be configured with the binding control flag and the server\_policy option. This configuration allows local account override of a network account.

**LDAP Authentication Module** The LDAP authentication module verifies the identity of a user. The [pam\\_sm\\_authenticate\(3PAM\)](#) function uses the password entered by the user to attempt to authenticate to the LDAP server. If successful, the user is authenticated. See NOTES for information on password prompting.

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.

The following options can 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 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. In the event that there is no user authentication token (`PAM_AUTHTOK`) available, the `pam_sm_acct_mgmt(3PAM)` function attempts to retrieve the user's account status without authenticating to the LDAP server as the user logging in. This procedure will succeed only if the LDAP server is Sun Java System Directory server 5.2 patch 4 or newer. The following options can 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

LDAP password management is no longer supported by `pam_ldap`. Use `pam_authtok_store(5)` instead of `pam_ldap` for password change. `pam_authtok_store(5)` handles both the local and LDAP accounts and updates the passwords in all the repositories configured by `nsswitch.conf(4)`.

**Errors** The authentication service returns the following error codes:

<code>PAM_SUCCESS</code>	The authentication was successful.
<code>PAM_MAXTRIES</code>	The maximum number of authentication attempts was exceeded.
<code>PAM_AUTH_ERR</code>	The authentication failed.
<code>PAM_USER_UNKNOWN</code>	No account is present for the user.
<code>PAM_BUF_ERR</code>	A memory buffer error occurred.
<code>PAM_SYSTEM_ERR</code>	A system error occurred.
<code>PAM_IGNORE</code>	The user's account was inactivated.

The account management service returns the following error codes:

<code>PAM_SUCCESS</code>	The user was allowed access to the account.
<code>PAM_NEW_AUTHTOK_REQD</code>	A new authentication token is required.
<code>PAM_ACCT_EXPIRED</code>	The user account has expired.
<code>PAM_PERM_DENIED</code>	The user was denied access to the account at this time.
<code>PAM_USER_UNKNOWN</code>	No account is present for the user.
<code>PAM_BUF_ERROR</code>	A memory buffer error occurred.
<code>PAM_SYSTEM_ERR</code>	A system error occurred.

**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 are 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 required   pam_unix_cred.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 comments and are ignored.

```
# 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 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. Lines that begin with the # symbol are comments and 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 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
```

**EXAMPLE 3** Using `pam_authtok_store` With Password Management For Both Local and LDAP Accounts *(Continued)*

```
# 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
```

**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 might not be human readable. Use `ldapclient(1M)` to update these files.

`/etc/pam.conf` PAM configuration file.

`/etc/pam.d/service` Alternate PAM configuration files.

**Attributes** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT-Level	MT-Safe with exceptions

**See Also** `idsconfig(1M)`, `ldap_cachemgr(1M)`, `ldapclient(1M)`, `libpam(3LIB)`, `pam(3PAM)`, `pam_sm_acct_mgmt(3PAM)`, `pam_sm_authenticate(3PAM)`, `pam_sm_chauthtok(3PAM)`, `pam_sm_close_session(3PAM)`, `pam_sm_open_session(3PAM)`, `pam_sm_setcred(3PAM)`, `syslog(3C)`, `pam.conf(4)`, `attributes(5)`, `ldap(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.

The previously supported `use_first_pass` and `try_first_pass` options are obsolete in this version, are no longer needed, can safely be removed from `pam.conf(4)`, and are silently ignored. They might be removed in a future release. Password prompting must be provided for by stacking `pam_authtok_get(5)` before `pam_ldap` in the `auth` and `password` module stacks and `pam_passwd_auth(5)` in the `passwd` service `auth` stack (as described in the `EXAMPLES` section). The previously supported password update function is replaced in this release by the previously recommended use of `pam_authtok_store` with the `server_policy` option (as described in the `EXAMPLES` section).

The functions: `pam_sm_setcred(3PAM)`, `pam_sm_chauthtok(3PAM)`, `pam_sm_open_session(3PAM)`, and `pam_sm_close_session(3PAM)` do nothing and return `PAM_IGNORE` in `pam_ldap`.

**Name** pam\_list – PAM account management module for UNIX

**Synopsis** pam\_list.so.1

**Description** The pam\_list module implements [pam\\_sm\\_acct\\_mgmt\(3PAM\)](#), which provides functionality to the PAM account management stack. The module provides functions to validate that the user's account is valid on this host based on a list of users and/or netgroups in the given file. The users and netgroups are separated by newline character. Netgroups are specified with character '@' as prefix before name of netgroup in the list. The maximum line length is 1023 characters.

The username is the value of PAM\_USER. The host is the value of PAM\_RHOST or, if PAM\_RHOST is not set, the value of the localhost as returned by [gethostname\(3C\)](#) is used.

If neither of the allow, deny, or compat options are specified, the module will look for +/- entries in the local /etc/passwd file. If this style is used, [nsswitch.conf\(4\)](#) must not be configured with compat for the passwd database. If no relevant +/- entry exists for the user, pam\_list is not participating in result.

If compat option is specified then the module will look for +/- entries in the local /etc/passwd file. Other entries in this file will be counted as + entries. If no relevant entry exists for the user, pam\_list will deny the access.

The following options can be passed to the module:

allow=	The full pathname to a file of allowed users and/or netgroups. Only one of allow= or deny= can be specified.
compat	Activate compat mode.
deny=	The full pathname to a file of denied users and/or netgroups. Only one of deny= or allow= can be specified.
debug	Provide <a href="#">syslog(3C)</a> debugging information at the LOG_AUTH   LOG_DEBUG level.
user	The module should only perform netgroup matches on the username. This is the default option.
nouser	The username should not be used in the netgroup match.
host	Only the host should be used in netgroup matches.
nohost	The hostname should not be used in netgroup matches.
norole	Return PAM_IGNORE if the account (PAM_USER) is a role. This is the default.
role	Evaluate the rules even if PAM_USER is a role account.
user_host_exact	The user and hostname must be in the same netgroup.

**Errors** The following error values are returned:

PAM_SERVICE_ERR	An invalid set of module options was specified in the PAM configuration (see <a href="#">pam.conf(4)</a> ) for this module, or the user/netgroup file could not be opened.
PAM_BUF_ERR	A memory buffer error occurred.
PAM_IGNORE	The module is ignored, as it is not participating in the result.
PAM_PERM_DENIED	The user is not on the allow list or is on the deny list.
PAM_SUCCESS	The account is valid for use at this time.
PAM_USER_UNKNOWN	No account is present for the user

**Examples** **EXAMPLE 1** Using pam\_list in default mode

The changes to /etc/pam.conf would be:

```
other account requisite    pam_roles.so.1
other account required    pam_unix_account.so.1
other account required    pam_list.so.1
```

The equivalent PAM configuration in /etc/pam.d/ would be the following entries in /etc/pam.d/other:

```
account requisite    pam_roles.so.1
account required    pam_unix_account.so.1
account required    pam_list.so.1
```

In the case of default mode or compat mode, the important lines in /etc/passwd appear as follows:

```
+loginname - user is approved
-loginname - user is disapproved
+@netgroup - netgroup members are approved
-@netgroup - netgroup members are disapproved
```

**EXAMPLE 2** Using pam\_list with allow file

The changes to /etc/pam.conf would be:

```
other account requisite    pam_roles.so.1
other account required    pam_unix_account.so.1
other account required    pam_list.so.1 allow=/etc/users.allow
```

The equivalent PAM configuration in /etc/pam.d/ would be the following entries in /etc/pam.d/other:

```
account requisite    pam_roles.so.1
account required    pam_unix_account.so.1
```

**EXAMPLE 2** Using pam\_list with allow file (Continued)

```
account required      pam_list.so.1 allow=/etc/users.allow
```

/etc/users.allow contains:

```
root
localloginname
remoteloginname
@netgroup
```

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT-Level	MT-Safe with exceptions

The interfaces in [libpam\(3LIB\)](#) are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

**See Also** [pam\(3PAM\)](#), [pam\\_authenticate\(3PAM\)](#), [pam\\_sm\\_acct\\_mgmt\(3PAM\)](#), [syslog\(3C\)](#), [libpam\(3LIB\)](#), [nsswitch.conf\(4\)](#), [pam.conf\(4\)](#), [attributes\(5\)](#)

**Name** pam\_passwd\_auth – authentication module for password

**Synopsis** pam\_passwd\_auth.so.1

**Description** 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. It should be the first module on the password service authentication stack.

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.

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 error codes are returned:

**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:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
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\\_authok\\_check\(5\)](#), [pam\\_authok\\_get\(5\)](#), [pam\\_authok\\_store\(5\)](#), [pam\\_dhkeys\(5\)](#), [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.

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\(3PAM\)](#).

**Name** pam\_pkcs11 – PAM Authentication Module for the PKCS#11 token libraries

**Synopsis** pam\_pkcs11.so [debug] [config\_file=filename]

**Description** The pam\_pkcs11 module implements `pam_sm_authenticate(3PAM)`, which provides functionality to the PAM authentication stack. This module allows a user to login a system, using a X.509 certificate and its dedicated private key stored in a PKCS#11 token. This module currently supports the RSA algorithm only.

To verify the dedicated private key is truly associated with the X.509 certificate, the following verification procedure is performed in this module by default:

- Generate 128 random byte data
- Sign the random data with the private key and get a signature. This step is done in the PKCS#11 token.
- Verify the signature using the public key extracted from the certificate.

For the verification of the users' certificates, locally stored CA certificates as well as either online or locally accessible CRLs are used.

**PAM CONFIGURATION** The pam\_pkcs11.so service module can be used in the <auth> PAM chain. The program that needs a PAM service should be configured in /etc/pam.conf or /etc/pam.d/service. For details on how to configure PAM services, see [pam.conf\(4\)](#).

The following example uses only pam\_pkcs11 for authentication:

```
login auth requisite pam_pkcs11.so.1
login autho required pam_unix_cred.so.1
```

The following example uses pam\_pkcs11 for authentication with fallback to standard UNIX authentication:

```
login auth sufficient pam_pkcs11.so.1
login auth requisite pam_authok_get.so.1
login auth required pam_dhkeys.so.1
login auth required pam_unix_cred.so.1
login auth required pam_unix_auth.so.1
```

**PAM\_PKCS11 CONFIGURATION** To configure the pam\_pkcs11 module, you must have the following information:

- Which PKCS #11 token you are going to use
- Which mapper(s) you need, and if needed, how to create and edit the related mapping files
- The root Certificate Authority files, and if required, the Certificate Revocation Lists files
- The list of authorized users to login, and their corresponding certificates

To configure the pam\_pkcs11 module, you need to modify the pam\_pkcs11.conf configuration file which is in the /etc/security/pam\_pkcs11 directory by default. For

detailed information on how to configure the `pam_pkcs11` module, see the *PAM-PKCS11 User Manual*, available at the <http://www.opensc-project.org/> web site, under the PAM PKCS#11 link.

The following example illustrates how to configure the `pam_pkcs11` module for a user whose certificate and private key are stored in the Solaris `pkcs11_softtoken` keystore. This example uses the default certificate verification policy.

- Set up the PKCS#11 module.  
On Solaris, the PKCS#11 module should be set to `/usr/lib/libpkcs11.so.1`, the PKCS#11 Cryptographic Framework library.
- Set up the `slot_description` entry.  
Specifies the slot to be used. For example, `slot_description = "Sun Crypto Softtoken"`. The default value for this entry is `none` which means to use the first slot with an available token.  
An administrator can use the `cryptoadm list -v` command to find all the available slots and their slot descriptions. For more information, see [libpkcs11\(3LIB\)](#) and [cryptoadm\(1M\)](#).
- Install or create user certificates and its dedicated private keys in the specific PKCS#11 token.
- Set up the certificate verification policy (`cert_policy`). If needed, set up CA certificate and CRL files.

The certificate verification policy includes:

<code>none</code>	Perform no verification
<code>ca</code>	Perform CA check
<code>signature</code>	Perform a signature check to ensure that private and public key matches
<code>crl_xxx</code>	Perform various certificate revocation checking

As this example uses the default policy, `cert_policy = ca, signature`, an administrator needs to set up the CA certificates.

- Copy the CA certificate to the `/etc/security/pam_pkcs11/cacerts` directory.  
A certificate that is self-signed is its own CA certificate. Therefore, in this example, the certificate is placed both in the Softtoken keystore and in the CA certificate directory.
- Make hash links for CA certificates  

```
$ /etc/security/pam_pkcs11/make_hash_link.sh \  
    /etc/security/pam_pkcs11/cacerts
```

- Set up the mappers and mapfiles.

When a X509 certificate is provided, there are no direct ways to map a certificate to a login. The `pam_pkcs11` module provides a configurable way with mappers to specify cert-to-user mapping.

Many mappers are provided by the `pam_pkcs11` module, for example, the common name (CN) mapper, the digest mapper, the Email mapper, or the LDAP mapper.

A user can configure a mapper list in the `pam_pkcs11.conf` file. The mappers in the list are used sequentially until the certificate is successfully matched with the user.

The default mapper list is as follows:

```
use_mappers = digest, cn, pwent, uid, mail, subject, null;
```

Some mappers do not require the specification of a mapfile, for example, the common name mapper. Other mappers require mapfiles, for example, the digest mapper. Some sample mapping files can be found in the `/etc/security/pam_pkcs11` directory.

**Options** The following options are supported:

`config_file=filename` Specify the configuration file. The default value is `/etc/security/pam_pkcs11/pam_pkcs11.conf`.

`debug` Enable debugging output.

**Files** `/usr/lib/security/pam_pkcs11.so`

`pam_pkcs11` module

`/usr/lib/pam_pkcs11/ldap_mapper.so`

Mapper module.

`/usr/lib/pam_pkcs11/opensc_mapper.so`

Mapper module.

`/usr/lib/pam_pkcs11/openssh_mapper.so`

Mapper module.

`/etc/security/pam_pkcs11/pam_pkcs11.conf`

Configuration file.

`/etc/security/pam_pkcs11/cacerts`

Configuration directory. Stores the CA certificates.

`/etc/security/pam_pkcs11/crls`

Configuration directory. Stores the CRL files.

`/etc/security/pam_pkcs11/digest_mapping.example`

Sample mapfile.

`/etc/security/pam_pkcs11/subject_mapping.example`

Sample mapfile.

`/etc/security/pam_pkcs11/mail_mapping.example`  
Sample mapfile.

`/etc/security/pam_pkcs11/make_hash_link.sh`  
Sample script.

**Authors** PAM-pkcs11 was originally written by MarioStrasser, mast@gmx.net.

Newer versions are from Juan Antonio Martinez, jonsito@teleline.es

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	library/security/pam/module/pam-pkcs11, SUNWpampkcs11r, SUNWpampkcs11-docs
Interface Stability	Uncommitted

**See Also** [pkcs11\\_inspect\(1\)](#), [pklogin\\_finder\(1\)](#), [cryptoadm\(1M\)](#),  
[libpkcs11\(3LIB\)](#), [libpkcs11\(3LIB\)](#), [pam\\_sm\\_authenticate\(3PAM\)](#), [pam.conf\(4\)](#),  
[attributes\(5\)](#), [pkcs11\\_softtoken\(5\)](#)

*PAM-PKCS11 User Manual*, available at the <http://www.opensc-project.org/> web site,  
under the PAM PKCS#11 link.

**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 – Solaris Roles account management module

**Synopsis** pam\_roles.so.1

**Description** The pam\_roles module implements [pam\\_sm\\_acct\\_mgmt\(3PAM\)](#). It provides functionality to verify that a user is authorized to assume a role. It also prevents direct logins to a role. The [user\\_attr\(4\)](#) database is used to determine which users can assume which roles.

The PAM items PAM\_USER and PAM\_AUSER, and PAM\_RHOST are used to determine the outcome of this module. PAM\_USER represents the new identity being verified. PAM\_AUSER, if set, represents the user asserting a new identity. If PAM\_AUSER is not set, the real user ID of the calling service implies that the user is asserting a new identity. Notice that root can never have roles.

This module is generally stacked above the [pam\\_unix\\_account\(5\)](#) module.

The following options are interpreted:

allow_remote	Allows a remote service to specify the user to enter as a role.
debug	Provides <a href="#">syslog(3C)</a> debugging information at the LOG_DEBUG level.

**Errors** The following values are returned:

PAM_IGNORE	If the type of the new user identity (PAM_USER) is “normal”. Or, if the type of the new user identity is “role” and the user asserting the new identity (PAM_AUSER) has the new identity name in its list of roles.
PAM_USER_UNKNOWN	No account is present for user.
PAM_PERM_DENIED	If the type of the new user identity (PAM_USER) is “role” and the user asserting the new identity (PAM_AUSER) does not have the new identity name in its list of roles.

**Examples** EXAMPLE 1 Using the pam\_roles.so.1 Module

The following example is a [pam.conf\(4\)](#) fragment that demonstrates the use of the pam\_roles.so.1 module:

```
cron account required pam_unix_account.so.1

other account requisite pam_roles.so.1
other account required pam_unix_account.so.1
```

The equivalent configuration in /etc/pam.d/ would be the following entry in /etc/pam.d/cron:

```
account required pam_unix_account.so.1
```

and the following entries in /etc/pam.d/other:

**EXAMPLE 1** Using the pam\_roles.so.1 Module *(Continued)*

```
account requisite pam_roles.so.1
account required pam_unix_account.so.1
```

The cron service does not invoke pam\_roles.so.1. Delayed jobs are independent of role assumption. All other services verify that roles cannot directly login. The “su” service (covered by the “other” service entry) verifies that if the new user is a role, the calling user is authorized for that role.

**EXAMPLE 2** Allowing Remote Roles

Remote roles should only be allowed from remote services that can be trusted to provide an accurate PAM\_AUSER name. This trust is a function of the protocol (such as sshd-hostbased).

The following example is a pam.conf(4) fragment that demonstrates the use of pam\_roles configuration for remote roles for the sshd-hostbased service.

```
sshd-hostbased account requisite pam_roles.so.1 allow_remote
sshd-hostbased account required pam_unix_account
```

The equivalent configuration in /etc/pam.d/ would be the following entries in /etc/pam.d/sshd-hostbased:

```
account requisite pam_roles.so.1 allow_remote
account required pam_unix_account
```

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT Level	MT-Safe with exceptions

**See Also** [roles\(1\)](#), [sshd\(1M\)](#), [su\(1M\)](#), [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam\\_acct\\_mgmt\(3PAM\)](#), [pam\\_setcred\(3PAM\)](#), [pam\\_set\\_item\(3PAM\)](#), [pam\\_sm\\_acct\\_mgmt\(3PAM\)](#), [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\\_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.

This module should never be stacked alone. It never returns PAM\_SUCCESS, as it never makes a positive decision.

The allow\_remote option should only be specified for services that are trusted to correctly identify the remote user (that is, sshd-hostbased).

PAM\_AUSER has replaced PAM\_RUSER whose definition is limited to the rlogin/rsh untrusted remote user name. See [pam\\_set\\_item\(3PAM\)](#).

**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** The SAMPLE authentication module 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 specified in the PAM configuration (see `pam.conf(4)`), or the string `test` if a default test password has not been set. The following options can be passed in to the SAMPLE Authentication module:

<code>debug</code>	Syslog debugging information at the LOG_DEBUG level.
<code>pass=newone</code>	Sets the password to be newone.
<code>first_pass_good</code>	The first password is always good when used with the <code>use_first_pass</code> or <code>try_first_pass</code> option.
<code>first_pass_bad</code>	The first password is always bad when used with the <code>use_first_pass</code> or <code>try_first_pass</code> option.
<code>always_fail</code>	Always returns PAM_AUTH_ERR.
<code>always_succeed</code>	Always returns PAM_SUCCESS.
<code>always_ignore</code>	Always returns PAM_IGNORE.
<code>use_first_pass</code>	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 configuration (see <code>pam.conf(4)</code> ).
<code>try_first_pass</code>	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.

**Sample Account Management Component** The SAMPLE Account Management Component 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 configuration (see [pam.conf\(4\)](#)). 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
gdm	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** The SAMPLE Password Management Component function ([pam\\_sm\\_chauthtok\(3PAM\)](#)), always returns `PAM_SUCCESS`.
- Sample Session Management Component** 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\(5\)](#)

**Warnings** This module should never be used outside of a closed debug environment. The examples of the `use_first_pass` and `try_first_pass` options are obsolete for all other Solaris delivered PAM service modules

**Notes** The interfaces in `libpam()` are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

**Name** pam\_smbfs\_login – PAM user credential authentication module for SMB/CIFS client login

**Synopsis** pam\_smb\_cred.so.1

**Description** The pam\_smbfs\_login module implements [pam\\_sm\\_setcred\(3PAM\)](#) to provide functions that act equivalently to the [smbadm\(1M\)](#) add-key command.

This optional functionality is meant to be used only in environments that do not run Active Directory or Kerberos, but which synchronize passwords between Solaris clients and their CIFS/SMB servers.

This module permits the login password to be stored as if the [smbadm\(1M\)](#) add-key command was used to store a password for PAM\_USER in the user or system default domain.

To use this functionality, add the following line to the /etc/pam.d/login file:

```
auth optional pam_smbfs_login.so.1
```

Authentication service modules must implement both [pam\\_sm\\_authenticate\(3PAM\)](#) and [pam\\_sm\\_setcred\(3PAM\)](#). In this module, [pam\\_sm\\_authenticate\(3PAM\)](#) always returns PAM\_IGNORE.

The [pam\\_sm\\_setcred\(3PAM\)](#) function accepts the following flags:

PAM\_REFRESH\_CRED  
Returns PAM\_IGNORE.

PAM\_SILENT  
Suppresses messages.

PAM\_ESTABLISH\_CRED  
PAM\_REINITIALIZE\_CRED  
Stores the authentication token for PAM\_USER in the same manner as the [smbadm\(1M\)](#) add-key command.

PAM\_DELETE\_CRED  
Deletes the stored password for PAM\_USER in the same manner as the [smbadm\(1M\)](#) remove-key command.

The following options can be passed to the pam\_smbfs\_login module:

debug  
Produces [syslog\(3C\)](#) debugging information at the LOG\_AUTH or LOG\_DEBUG level.

nowarn  
Suppresses warning messages.

**Errors** Upon successful completion of [pam\\_sm\\_setcred\(3PAM\)](#), PAM\_SUCCESS is returned. The following error codes are returned upon error:

PAM\_USER\_UNKNOWN  
User is unknown.

PAM\_AUTHOK\_ERR  
Password is bad.

PAM\_AUTH\_ERR  
Domain is bad.

PAM\_SYSTEM\_ERR  
System error.

**Attributes** See [attributes\(5\)](#) for descriptions of the following attribute:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT Level	MT-Safe with exceptions

**See Also** [smbadm\(1M\)](#), [syslog\(3C\)](#), [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam\\_setcred\(3PAM\)](#), [pam\\_sm\(3PAM\)](#), [pam\\_sm\\_authenticate\(3PAM\)](#), [pam\\_sm\\_chauthtok\(3PAM\)](#), [pam\\_sm\\_setcred\(3PAM\)](#), [pam.conf\(4\)](#), [attributes\(5\)](#), [smbfs\(7FS\)](#)

**Notes** The interfaces in [libpam\(3LIB\)](#) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

**Name** pam\_smb\_passwd – SMB password management module

**Synopsis** pam\_smb\_passwd.so.1

**Description** The pam\_smb\_passwd module enhances the PAM password management stack. This functionality supports the changing or adding of SMB passwords for local Solaris users. The Solaris SMB server uses SMB passwords to authenticate connected Solaris users. This module includes the [pam\\_sm\\_chauthtok\(3PAM\)](#) function.

The `pam_sm_chauthtok()` function accepts the following flags:

**PAM\_PRELIM\_CHECK**  
Always returns `PAM_IGNORE`.

**PAM\_SILENT**  
Suppresses messages.

**PAM\_UPDATE\_AUTHTOK**  
Updates or creates a new SMB local LM/NTLM hash for the user that is specified in `PAM_USER` by using the authentication information found in `PAM_AUTHTOK`. The LM hash is only created if the `smbd/lmauth_level` property value of the `smb/server` service is set to 3 or less. `PAM_IGNORE` is returned if the user is not in the `local/etc/passwd` repository.

The following options can be passed to the `pam_smb_passwd` module:

**debug**  
Produces [syslog\(3C\)](#) debugging information at the `LOG_AUTH` or `LOG_DEBUG` level.

**nowarn**  
Suppresses warning messages.

**Files** `/var/smb/smbpasswd`  
Stores SMB passwords for Solaris users.

**Errors** Upon successful completion of `pam_sm_chauthtok()`, `PAM_SUCCESS` is returned. The following error codes are returned upon error:

**PAM\_AUTHTOK\_ERR**  
Authentication token manipulation error

**PAM\_AUTHTOK\_LOCK\_BUSY**  
SMB password file is locked

**PAM\_PERM\_DENIED**  
Permissions are insufficient for accessing the SMB password file

**PAM\_SYSTEM\_ERR**  
System error

**PAM\_USER\_UNKNOWN**  
User is unknown

**Attributes** See the [attributes\(5\)](#) man page for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT Level	MT-Safe with exceptions

**See Also** [smbd\(1M\)](#), [syslog\(3C\)](#), [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam\\_chauthtok\(3PAM\)](#), [pam\\_sm\(3PAM\)](#), [pam\\_sm\\_chauthtok\(3PAM\)](#), [pam.conf\(4\)](#), [attributes\(5\)](#)

**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_smb_passwd.so.1` module should be stacked following all password qualification modules in the PAM password stack.

**Name** pam\_tsol\_account – PAM account management module for Trusted Extensions

**Synopsis** /usr/lib/security/pam\_tsol\_account.so.1

**Description** The Solaris Trusted Extensions service module for PAM, /usr/lib/security/pam\_tsol\_account.so.1, checks account limitations that are related to labels. The pam\_tsol\_account.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.

pam\_tsol\_account.so.1 contains a function to perform account management, pam\_sm\_acct\_mgmt(). The function checks for the allowed label range for the user. The allowable label range is set by the defaults in the [label\\_encodings\(4\)](#) file. These defaults can be overridden by entries in the [user\\_attr\(4\)](#) database.

By default, this module requires that remote hosts connecting to the global zone must have a CIPSO host type. To disable this policy, add the allow\_unlabeled keyword as an option to the entry in [pam.conf\(4\)](#), as in:

```
other account required pam_tsol_account allow_unlabeled
```

The equivalent PAM configuration in /etc/pam.d/ would be the following entry in /etc/pam.d/other:

```
account required pam_tsol_account allow_unlabeled
```

**Options** The following options can be passed to the module:

`allow_unlabeled` Allows remote connections from hosts with unlabeled template types.

`debug` Provides debugging information at the LOG\_DEBUG level. See [syslog\(3C\)](#).

**Return Values** The following values are returned:

PAM\_SUCCESS The account is valid for use at this time and label.

PAM\_PERM\_DENIED The current process label is outside the user's label range, or the label information for the process is unavailable, or the remote host type is not valid.

Other values Returns an error code that is consistent with typical PAM operations. For information on error-related return values, see the [pam\(3PAM\)](#) man page.

**Attributes** See [attributes\(5\)](#) for description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

---

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	MT-Safe with exceptions

The interfaces in [libpam\(3LIB\)](#) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

**See Also** [keylogin\(1\)](#), [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam\\_sm\\_acct\\_mgmt\(3PAM\)](#), [pam\\_start\(3PAM\)](#), [syslog\(3C\)](#), [label\\_encodings\(4\)](#), [pam.conf\(4\)](#), [user\\_attr\(4\)](#), [attributes\(5\)](#)

Chapter 14, “Using Pluggable Authentication Modules,” in *Oracle Solaris 11.1 Administration: Security Services*

**Notes** The functionality described on this manual page is available only if the system is configured with Trusted Extensions.

**Name** pam\_tty\_tickets – PAM authentication module

**Synopsis** pam\_tty\_tickets.so.1 [timeout=*minutes*] [*sudo-compat*] [debug]

**Description** The pam\_tty\_tickets module provides a mechanism for checking a ticket that was created by a prior successful authentication. Tickets by default validity of 5 minutes.

The default ticket location includes both the source (PAM\_AUSER) and destination (PAM\_USER) as well as the tty (PAM\_TTY) for which it is valid.

The module can be configured using the *sudo-compat* option to store the tickets in the same location as sudo, though use of sudo is not required to use this feature.

The pam\_sm\_setcred() function creates a ticket for the user in the tickets directory.

The pam\_sm\_authenticate() function checks the timestamp on the ticket is no older than the timeout value, if it is then it returns PAM\_SUCCESS. If it is older then the ticket is removed and the module returns PAM\_IGNORE.

This module is intended to be placed in the auth stack with the sufficient control flag.

No messages are produced by this module using the PAM conversation function. Some messages are sent to syslog for error conditions as well as messages at LOG\_INFO for ticket validity checking

The following options can be passed to the module:

*debug* Debugging information is sent to syslog LOG\_AUTH|LOG\_DEBUG.

*sudo-compat* Location of the per user (per tty) tickets, matches the sudo location. When this option is set PAM\_USER must be root other wise the module returns PAM\_IGNORE and tickets are not read or created.

*timeout* Validity time in minutes for a ticket. The default is 5 minutes.

### Examples **EXAMPLE 1** Using the Default Settings

The following is an excerpt of a sample pam.conf configuration file that has per tty tickets with the default time out (5 minutes) for users authenticating with *su(1M)*:

```
su auth required    pam_unix_cred.so.1
su auth sufficient  pam_tty_tickets.so.1
su auth requisite   pam_authtok_get.so.1
su auth required    pam_dhkeys.so.1
su auth required    pam_unix_auth.so.1
```

### **EXAMPLE 2** Changing the Default Settings

The following example changes the defaults so that tickets are valid for 10 minutes and uses the sudo location:

**EXAMPLE 2** Changing the Default Settings *(Continued)*

```

su auth required    pam_unix_cred.so.1
su auth sufficient pam_tty_tickets.so.1 sudo-compat timeout=10
su auth requisite   pam_authtok_get.so.1
su auth required    pam_dhkeys.so.1
su auth required    pam_unix_auth.so.1

```

**Errors** PAM\_SUCCESS Ticket is valid

PAM\_IGNORE All other cases

**Files** /system/volatile/tty\_tickets/<PAM\_AUSER>/<PAM\_USER>/<PAM\_TTY>  
 Default ticket location.

/system/volatile/sudo/<PAM\_AUSER>/<PAM\_TTY>

When used sudo-compat is set this file has the same format as those created by sudo.

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	See below.

The syslog messages are Volatile. The module name, module options, and ticket locations are Committed.

**See Also** [su\(1M\)](#), [sudo\(1M\)](#), [pam\(3PAM\)](#), [pam\\_sm\\_authenticate\(3PAM\)](#), [pam\\_sm\\_setcred\(3PAM\)](#), [attributes\(5\)](#)

**Name** pam\_unix\_account – PAM account management module for UNIX

**Synopsis** pam\_unix\_account.so.1

**Description** pam\_unix\_account module implements `pam_sm_acct_mgmt()`, which provides functionality to the PAM account management stack. The module provides functions to validate that the user's account is not locked or expired and that the user's password does not need to be changed. The module retrieves account information from the configured databases in `nsswitch.conf(4)`.

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\_UNIX\_ACCOUNT User account has expired

PAM\_AUTHTOK\_EXPIRED Password expired and no longer usable

PAM\_BUF\_ERR Memory buffer error

PAM\_IGNORE Ignore module, not participating in result

PAM\_NEW\_AUTHTOK\_REQD Obtain new authentication token from the user

PAM\_PERM\_DENIED The account is locked or has been inactive for too long

PAM\_SERVICE\_ERR Error in underlying service module

PAM\_SUCCESS The account is valid for use at this time

PAM\_USER\_UNKNOWN No account is present for the user

**Attributes** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
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)`

**Notes** The interfaces in [libpam\(3LIB\)](#) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

Attempts to validate locked accounts are logged via [syslog\(3C\)](#) to the LOG\_AUTH facility with a LOG\_NOTICE severity.

**Name** pam\_unix\_auth – PAM authentication module for UNIX

**Synopsis** pam\_unix\_auth.so.1

**Description** The pam\_unix\_auth module implements pam\_sm\_authenticate(), which provides functionality to the PAM authentication stack. It provides functions that use [crypt\(3C\)](#) to verify that the password contained in the PAM item PAM\_AUTHTOK is the correct password for the user specified in the item PAM\_USER.

If PAM\_AUSER and PAM\_USER are both specified and PAM\_USER is a role, the [user\\_attr\(4\)](#) keyword roleauth is checked to determine if the password that is checked is for the role (PAM\_USER) or the assuming user (PAM\_AUSER). If PAM\_REPOSITORY is specified, the user's password is fetched from that repository. Otherwise, the default [nsswitch.conf\(4\)](#) repository is searched for that user.

For accounts in the name services which support automatic account locking, the account can be configured to be automatically locked (see [user\\_attr\(4\)](#) and [policy.conf\(4\)](#)) after multiple failed login attempts. For accounts that are configured for automatic locking, if authentication failure is to be returned, the failed login counter is incremented upon each failure. If the number of successive failures equals or exceeds the configured value, the account is locked and PAM\_MAXTRIES is returned. The files (see [passwd\(4\)](#) and [shadow\(4\)](#)) and ldap (when configured with enableShadowUpdate true, see [ldapclient\(1M\)](#)), repositories support automatic account locking. A successful authentication by this module clears the failed login counter and reports the number of failed attempts since the last successful authentication.

Authentication service modules must implement both pam\_sm\_authenticate() and pam\_sm\_setcred(). To allow the authentication portion of UNIX authentication to be replaced, pam\_sm\_setcred() in this module always returns PAM\_IGNORE. This module should be stacked with [pam\\_unix\\_cred\(5\)](#) to ensure a successful return from [pam\\_setcred\(3PAM\)](#).

The following options can be passed to the module:

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.

noLock

Regardless of the automatic account locking setting for the account, do not lock the account, increment or clear the failed login count. The noLock option allows for exempting account locking on a per service basis.

**Errors** The following error codes are returned from `pam_sm_authenticate()`:

- PAM\_AUTH\_ERR  
Authentication failure.
- PAM\_BUF\_ERR  
Memory buffer error.
- PAM\_IGNORE  
Ignores module, not participating in result.
- PAM\_MAXTRIES  
Maximum number of retries exceeded.
- PAM\_PERM\_DENIED  
Permission denied.
- PAM\_SUCCESS  
Successfully obtains authentication token.
- PAM\_SYSTEM\_ERR  
System error.
- PAM\_USER\_UNKNOWN  
No account present for user.

The following error codes are returned from `pam_sm_setcred()`:

- PAM\_IGNORE  
Ignores this module regardless of the control flag.

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT Level	MT-Safe with exceptions

**See Also** [login\(1\)](#), [passwd\(1\)](#), [ldapclient\(1M\)](#), [useradd\(1M\)](#), [usermod\(1M\)](#), [roleadd\(1M\)](#), [rolemo\(1M\)](#), [crypt\(3C\)](#), [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam\\_authenticate\(3PAM\)](#), [pam\\_setcred\(3PAM\)](#), [syslog\(3C\)](#), [pam.conf\(4\)](#), [passwd\(4\)](#), [policy.conf\(4\)](#), [nsswitch.conf\(4\)](#), [shadow\(4\)](#), [user\\_attr\(4\)](#), [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\\_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.

If the `PAM_REPOSITORY item_type` is set and a service module does not recognize the type, the service module does not process any information, and returns `PAM_IGNORE`. If the `PAM_REPOSITORY item_type` is not set, a service module performs its default action.

**Name** pam\_unix\_cred – PAM user credential authentication module for UNIX

**Synopsis** pam\_unix\_cred.so.1

**Description** The pam\_unix\_cred module implements [pam\\_sm\\_setcred\(3PAM\)](#). It provides functions that establish user credential information. It is a module separate from the [pam\\_unix\\_auth\(5\)](#) module to allow replacement of the authentication functionality independently from the credential functionality.

The pam\_unix\_cred module must always be stacked along with whatever authentication module is used to ensure correct credential setting.

Authentication service modules must implement both `pam_sm_authenticate()` and `pam_sm_setcred()`.

`pam_sm_authenticate()` in this module always returns `PAM_IGNORE`.

`pam_sm_setcred()` initializes the user's project, privilege sets and initializes or updates the user's audit context if it hasn't already been initialized. The following flags may be set in the flags field:

`PAM_ESTABLISH_CRED`

`PAM_REFRESH_CRED`

`PAM_REINITIALIZE_CRED`

Initializes the user's project to the project specified in `PAM_RESOURCE`, or if `PAM_RESOURCE` is not specified, to the user's default project. Establishes the user's privilege sets.

If the audit context is not already initialized and auditing is configured, these flags cause the context to be initialized to that of the user specified in `PAM_AUSER` (if any) merged with the user specified in `PAM_USER` and host specified in `PAM_RHOST`. If `PAM_RHOST` is not specified, `PAM_TTY` specifies the local terminal name. Attributing audit to `PAM_AUSER` and merging `PAM_USER` is required for correctly attributing auditing when the system entry is performed by another user that can be identified as trustworthy.

If the audit context is already initialized, the `PAM_REINITIALIZE_CRED` flag merges the current audit context with that of the user specified in `PAM_USER`. `PAM_REINITIALIZE_CRED` is useful when a user is assuming a new identity, as with [su\(1M\)](#).

`PAM_DELETE_CRED`

This flag has no effect and always returns `PAM_SUCCESS`.

The following options are interpreted:

`debug` Provides [syslog\(3C\)](#) debugging information at the `LOG_DEBUG` level.

`nowarn` Disables any warning messages.

**Errors** Upon successful completion of `pam_sm_setcred()`, `PAM_SUCCESS` is returned. The following error codes are returned upon error:

<code>PAM_CRED_UNAVAIL</code>	Underlying authentication service cannot retrieve user credentials
<code>PAM_CRED_EXPIRED</code>	User credentials have expired
<code>PAM_USER_UNKNOWN</code>	User is unknown to the authentication service
<code>PAM_CRED_ERR</code>	Failure in setting user credentials
<code>PAM_BUF_ERR</code>	Memory buffer error
<code>PAM_SYSTEM_ERR</code>	System error

The following values are returned from `pam_sm_authenticate()`:

`PAM_IGNORE` Ignores this module regardless of the control flag

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT Level	MT-Safe with exceptions

**See Also** [ssh\(1\)](#), [su\(1M\)](#), [settaskid\(2\)](#), [libpam\(3LIB\)](#), [getprojent\(3PROJECT\)](#), [pam\(3PAM\)](#), [pam\\_set\\_item\(3PAM\)](#), [pam\\_sm\\_authenticate\(3PAM\)](#), [syslog\(3C\)](#), [setproject\(3PROJECT\)](#), [pam.conf\(4\)](#), [nsswitch.conf\(4\)](#), [project\(4\)](#), [attributes\(5\)](#), [pam\\_authtok\\_check\(5\)](#), [pam\\_authtok\\_get\(5\)](#), [pam\\_authtok\\_store\(5\)](#), [pam\\_dhkeys\(5\)](#), [pam\\_passwd\\_auth\(5\)](#), [pam\\_unix\\_auth\(5\)](#), [pam\\_unix\\_account\(5\)](#), [pam\\_unix\\_session\(5\)](#), [privileges\(5\)](#)

**Notes** The interfaces in [libpam\(3LIB\)](#) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

If this module is replaced, the audit context and credential may not be correctly configured.

**Name** pam\_unix\_session – session management PAM module for UNIX

**Synopsis** pam\_unix\_session.so.1

**Description** The pam\_unix\_session module implements [pam\\_sm\\_open\\_session\(3PAM\)](#) and [pam\\_sm\\_close\\_session\(3PAM\)](#).

[pam\\_sm\\_open\\_session\(\)](#) updates the `/var/adm/lastlog` file with the information contained in the `PAM_USER`, `PAM_TTY`, and `PAM_RHOST` items. [pam\\_unix\\_account\(5\)](#) uses this account to determine the previous time the user logged in.

[pam\\_sm\\_close\\_session\(\)](#) is a null function.

The following options can be passed to the module:

`debug`     [syslog\(3C\)](#) debugging information at the `LOG_DEBUG` level

**Errors** Upon successful completion, `PAM_SUCCESS` is returned. The following error codes are returned upon error:

`PAM_SESSION_ERR`     Cannot make or remove the entry for the specified session (`PAM_TTY` is not present).

`PAM_USER_UNKNOWN`     No account is present for *user*.

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
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\\_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.

**Name** pam\_user\_policy – PAM user authentication policy module

**Synopsis** pam\_user\_policy.so.1

**Description** The pam\_user\_policy module causes a user-specific PAM configuration to be evaluated and returns the result of evaluating such a configuration.

The pam\_user\_policy module implements all PAM service module functions.

The PAM configuration to evaluate is determined by looking for a pam\_policy key in a user's attributes (see [user\\_attr\(4\)](#)) or profiles (see [prof\\_attr\(4\)](#)), or failing that, by looking at the default profiles granted in [policy.conf\(4\)](#). If no PAM configuration is found, the “unix” policy is used.

Failure to obtain a user name is considered an error (see below).

This module should generally be stacked as the first module, possibly as the only module, in a PAM service configuration using a *control\_flag* value of 'sufficient' or 'binding', depending on the contents of the user-specific PAM configuration.

The pathname to the user-specific PAM configuration file passed to [pam\\_eval\(3PAM\)](#) must be absolute so pam\_user\_policy prepends “/etc/security/pam\_policy” to any non-absolute PAM configuration pathnames.

The following option can be passed to the module:

debug     [syslog\(3C\)](#) debugging information at the LOG\_DEBUG level

**Authentication Module** The [pam\\_get\\_user\(3PAM\)](#) function is used to retrieve the current user name and sets this to be the value of PAM\_USER if PAM\_USER was not already set. If no user name can be obtained, PAM\_USER\_UNKNOWN is returned.

The pam\_user\_policy authentication module then looks up the name of a PAM configuration file to use for that user as described above and evaluates the named configuration by calling [pam\\_eval\(\)](#) with the same flags as were passed to the pam\_user\_policy authentication module.

The configuration file name found or the default “unix” is saved as module data (see [pam\\_set\\_data\(3PAM\)](#)) for use by other pam\_user\_policy modules.

**Other Modules** If the PAM\_USER item is not set to a non-empty string then pam\_user\_policy returns PAM\_USER\_UNKNOWN immediately. If a PAM configuration file name was saved as module data by a previous call to a pam\_user\_policy module, then that configuration will be used; otherwise a PAM configuration will be looked up as described above. The service module then evaluates the named configuration by calling [pam\\_eval\(\)](#) with the same flags as were passed to the service module.

The configuration file name found or the default “unix” is saved as module data (see [pam\\_set\\_data\(3PAM\)](#)) for use by other pam\_user\_policy modules.

**Return Values** If PAM\_USER is not set or cannot be obtained, the module's service functions return PAM\_USER\_UNKNOWN. If module-specific data cannot be stored, PAM\_SERVICE\_ERR is returned. Failure to allocate resources causes the module to return PAM\_BUF\_ERR. Otherwise the value returned by pam\_eval() is returned.

**Files** A number of pam.conf files for inclusion by pam\_user\_policy can be found in /etc/security/pam\_policy:

unix	Use only Unix passwords for authentication, account management, and password management.
krb5_only	Use Kerberos V5 only for authentication, account management, and password management.
krb5_first	Use Kerberos V5 for authentication with fallback on Unix authentication, use Kerberos V5 for account management and password management for Kerberos users and Unix for account management and password management for Unix users.
krb5_optional	Use Unix for authentication, account management, and password management and then optionally using Kerberos V5 for authentication, account management and password management for Kerberos users.
ldap	Use pam_ldap(5) for authentication, account management, and password management for LDAP users and Unix for authentication, account management, and password management for Unix users.
any	Try Kerberos V, LDAP and Unix, in that order, and as sufficient, for authentication, account management, and password management.

**Examples** **EXAMPLE 1** Authenticate a user with Kerberos V5 for all PAM services.

In the following example, user 'larry' should only be authenticated with Kerberos V5 for all PAM services.

```
$ usermod -K pam_policy=krb5_only larry
```

**EXAMPLE 2** Use the PAM configuration /etc/security/pam\_policy/custom for a user.

In the following example, the PAM configuration /etc/security/pam\_policy/custom should be used for user 'curly'. This custom PAM configuration might have different configurations for different PAM services, such as requiring Unix authentication for console logins but Kerberos V5 for all other PAM services.

```
$ usermod -K pam_policy=custom curly
```

**EXAMPLE 3** Create a new profile.

The following example creates a new profile named “PAM Per-User Policy of LDAP” and assign it to user 'moe' indicating that [pam\\_ldap\(5\)](#) should be used for all PAM services. Alternatively the profile could be assigned to all users by adding it to PROFS\_GRANTED in [policy.conf\(4\)](#).

```
$ profiles -p "PAM Per-User Policy of LDAP" \
  'set desc="Profile which sets pam_policy=ldap";
  set pam_policy=ldap; exit;'
$ usermod -P "PAM Per-User Policy of LDAP" moe
```

**EXAMPLE 4** Add a new user.

The following example adds a new user named 'shemp' who uses the PAM configuration /usr/local/etc/pam.conf for all PAM services.

```
$ useradd -K pam_policy=/usr/local/etc/pam.conf shemp
```

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT-Level	MT-Safe with exceptions

**See Also** [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam\\_eval\(3PAM\)](#), [pam\\_get\\_user\(3PAM\)](#), [pam\\_set\\_data\(3PAM\)](#), [syslog\(3C\)](#), [pam.conf\(4\)](#), [policy.conf\(4\)](#), [prof\\_attr\(4\)](#), [user\\_attr\(4\)](#), [attributes\(5\)](#), [pam\\_ldap\(5\)](#)

**Notes** The interfaces in [libpam\(3LIB\)](#) are MT-Safe only if each thread within the multithreaded application uses its own PAM handle.

**Name** pam\_zfs\_key – PAM user credential module for ZFS

**Synopsis** pam\_zfs\_key.so.1 [create] [homes=]

**Description** The pam\_zfs\_key module implements pam\_sm\_setcred(3PAM) and pam\_sm\_chauthtok(3PAM).

The pam\_zfs\_key module provides functions that allow loading and changing of the ZFS encryption passphrase for encrypted file systems that are mounted at the users home directory location. Authentication service modules must implement both pam\_sm\_authenticate() and pam\_sm\_setcred().

pam\_sm\_authenticate() in this module always returns PAM\_IGNORE. If they are located at a different location then the module option, homes= can be used to specify that. It is the ZFS dataset name, not the mountpoint, which would usually be /export/home/.

The last component of the ZFS dataset name must match the value of PAM\_USER, that is, the users login name. If the users home directory is a local ZFS filesystem with encryption enabled and the ZFS keysource property is set to passphrase, prompt, on pam\_sm\_setcred() this module manages the keys as follows:

**PAM\_DELETE\_CRED** Attempts to umount the filesystem and unload the key. This often fails because there are still processes with the user's home directory as the current working directory. A force module option is provided to attempt a forced unmount first.

**PAM\_ESTABLISH\_CRED** Attempts to use the value of PAM\_AUTHTOK to load the key for the ZFS dataset and mount it.

If PAM\_AUTHTOK is not the correct passphrase, the user is prompted once for an alternate. This value is never stored in PAM\_AUTHTOK, even if it is the correct passphrase for the ZFS dataset that is the users home directory. This is equivalent to:

```
zfs key -l rpool/export/home/$USER
```

If no ZFS file system exists for the user and the create module option is provided, a new one is created. The ZFS encryption property defaults to on in this case unless the encryption= property is set for the module to override it.

If the create module option is not provided and no ZFS file system exists for the user, the module returns PAM\_IGNORE.

The newly created ZFS file system has the following ZFS delegations specified for the user for which it is created: key, keychange, mount. In these cases keysource is always set to passphrase, prompt.

When `pam_sm_chauthtok(3PAM)` is called, for example, on password change, this module attempts to change the passphrase for the ZFS dataset to match the value in `PAM_AUTHTOK`. This is equivalent to running:

```
zfs key -c rpool/export/home/$USER
```

This requires that the user have the keychange delegation, as password change usually runs as the user.

The following mount options are supported:

<code>create</code>	Create new ZFS datasets
<code>encryption</code>	Set the ZFS encryption property for create
<code>force</code>	Attempt a <code>umount2(2)</code> with a <code>MS_FORCE</code> of the dataset when doing <code>PAM_DELETE_CRED</code> .
<code>homes=</code>	Alternate location of ZFS datasets for user home directories. The default is <code>rpool/export/home</code> .
<code>nowarn</code>	Do not provide any error messages or warnings.

#### Examples EXAMPLE 1 Using pam\_zfs\_key in Default Mode

The following example uses `pam_zfs_key` in default mode.

```
gdm  auth requisite      pam_authtok_get.so.1
gdm  auth required      pam_dhkeys.so.1
gdm  auth required      pam_unix_cred.so.1
gdm  auth required      pam_unix_auth.so.1
gdm  auth optional      pam_zfs_key.so.1

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
other password optional  pam_zfs_key.so.1
```

#### EXAMPLE 2 Specifying an Alternate ZFS Dataset

The following example specifies an alternate ZFS dataset location for the home directory file systems. New entries should be created if they are not present using `aes-256-gcm` as the ZFS encryption property setting.

```
gdm auth requisite pam_authtok_get.so.1
gdm auth required  pam_dhkeys.so.1
gdm auth required  pam_unix_cred.so.1
gdm auth required  pam_unix_auth.so.1
gdm auth optional  pam_zfs_key.so.1 homes=tank/users \
create encryption=aes-256-gcm
```

**EXAMPLE 3** Making it Mandatory for the ZFS Dataset to Mount

The following example makes it mandatory for the ZFS dataset to mount and ensures the passphrase always stays in sync with the login password.

```
gdm    auth requisite      pam_authtok_get.so.1
gdm    auth required       pam_dhkeys.so.1
gdm    auth required       pam_unix_cred.so.1
gdm    auth required       pam_unix_auth.so.1
gdm    auth required       pam_zfs_key.so.1

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 requisite  pam_zfs_key.so.1
other  password required   pam_authtok_store.so.1
```

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT-Level	MT-Safe with exceptions. See below.

The interfaces in [libpam\(3LIB\)](#) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

**See Also** [zfs\(1M\)](#), [umount2\(2\)](#), [pam.conf\(4\)](#), [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam\\_sm\\_chauthtok\(3PAM\)](#), [pam\\_sm\\_setcred\(3PAM\)](#), [attributes\(5\)](#)

**Name** pkcs11\_kernel – PKCS#11 interface to Kernel Cryptographic Framework

**Synopsis** /usr/lib/security/pkcs11\_kernel.so  
/usr/lib/security/64/pkcs11\_kernel.so

**Description** The `pkcs11_kernel.so` object implements the RSA Security Inc. PKCS#11 Cryptographic Token Interface (Cryptoki), v2.20, specification by using a private interface to communicate with the Kernel Cryptographic Framework.

Each unique hardware provider is represented by a PKCS#11 slot. In a system with no hardware Kernel Cryptographic Framework providers, this PKCS#11 library presents no slots.

The PKCS#11 mechanisms provided by this library is determined by the available hardware providers.

Application developers should link to `libpkcs11.so` rather than link directly to `pkcs11_kernel.so`. See [libpkcs11\(3LIB\)](#).

All of the Standard PKCS#11 functions listed on [libpkcs11\(3LIB\)](#) are implemented except for the following:

- C\_DecryptDigestUpdate
- C\_DecryptVerifyUpdate
- C\_DigestEncryptUpdate
- C\_GetOperationState
- C\_InitToken
- C\_InitPIN
- C\_SetOperationState
- C\_SignEncryptUpdate
- C\_WaitForSlotEvent

A call to these functions returns `CKR_FUNCTION_NOT_SUPPORTED`.

Buffers cannot be greater than 2 megabytes. For example, `C_Encrypt()` can be called with a 2 megabyte buffer of plaintext and a 2 megabyte buffer for the ciphertext.

The maximum number of object handles that can be returned by a call to `C_FindObjects()` is 512.

The maximum amount of kernel memory that can be used for crypto operations is limited by the `project.max-crypto-memory` resource control. Allocations in the kernel for buffers and session-related structures are charged against this resource control.

**Return Values** The return values of each of the implemented functions are defined and listed in the RSA PKCS#11 v2.20 specification. See <http://www.rsasecurity.com>.

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT-Level	MT-Safe with exceptions. See section 6.6.2 of RSA PKCS#11 v2.20
Standard	PKCS#11 v2.20

**See Also** [cryptoadm\(1M\)](#), [rctladm\(1M\)](#), [libpkcs11\(3LIB\)](#), [attributes\(5\)](#), [pkcs11\\_softtoken\(5\)](#)

RSA PKCS#11 v2.20 <http://www.rsasecurity.com>

**Notes** Applications that have an open session to a PKCS#11 slot make the corresponding hardware provider driver not unloadable. An administrator must close the applications that have an PKCS#11 session open to the hardware provider to make the driver unloadable.

**Name** pkcs11\_kms – RSA PKCS#11 provider for the Oracle Key Manager

**Synopsis** /usr/lib/security/pkcs11\_kms.so  
/usr/lib/security/64/pkcs11\_kms.so

**Description** The `pkcs11_kms.so` object implements the RSA Security Inc. PKCS#11 Cryptographic Token Interface (Cryptoki), v2.20, specification using the Oracle Key Manager (OKM) KMS agent protocol to talk to an Oracle Key Manager appliance (KMA). This provider implements the PKCS#11 specification and communicates to a remote OKM using the (private) KMS client protocol.

The following PKCS#11 mechanisms are supported in this provider: `CKM_AES_KEY_GEN`, `CKM_AES_CBC_PAD`, and `CKM_AES_CBC`.

The following PKCS#11 interfaces are supported by this provider:

```
C_Initialize
C_Finalize
C_GetInfo
C_GetAttributeValue
C_SetAttributeValue
C_GetFunctionList
C_GetSlotList
C_GetSlotInfo
C_GetTokenInfo
C_GetMechanismList
C_GetMechanismInfo
C_InitToken
C_SetPIN
C_Login
C_Logout
C_FindObjectsInit/C_FindObjects/C_FindObjectsFinal
C_GenerateKey
C_EncryptInit/C_Encrypt/C_EncryptFinal
C_DecryptInit/C_Decrypt/C_DecryptFinal
C_DestroyObject
C_OpenSession
C_CloseSession
C_CloseAllSessions
C_GetSessionInfo
C_CreateObject
C_CopyObject
C_GetObjectSize
C_EncryptUpdate
C_DecryptUpdate
```

All other functions return `CKR_FUNCTION_NOT_SUPPORTED` when called.

**Prerequisites** The `pkcs11_kms` provider can only be used on a system that has access to an OKM. The OKM administrator must configure a an agent ID for each user (or application) that is accessing the OKM. This is done through the OKM utilities that are part of the OKM administrative tools and are not bundled in Oracle Solaris.

Once the OKM administrator has configured the KMA for use and communicated the parameters to the client, that is, an Oracle Solaris user or application, the Oracle Solaris PKCS#11 KMS provider can be initialized for use.

Initializing the KMS provider is done through the use of the `kmscfg(1M)` utility. At a minimum, `kmscfg` requires the user to enter the name of a profile, the OKM Agent ID, the initial password used to secure the profile, and the IP address of the KMA in order to initialize the local provider configuration files for further use. See the `kmscfg(1M)` manual page for details.

Once `kmscfg` has been run and the local token namespace has been configured, the user can then initialize the token for use. Initializing the token is done using the `pktool(1)` command as follows:

```
$ pktool inittoken currlabel=KMS
```

The user has to supply the default SO (security officer) PIN before being able to initialize the KMS provider for use. The default SO PIN is whatever was used by the OKM administrator when initially setting up the OKM Agent. The user initializing the token must know this passphrase in order to initialize the provider.

Once the provider is initialized, the user PIN can be changed from the default values. Again, `pktool(1)` is used to change the PIN value.

Use the following command to change the local PIN:

```
$ pktool setpin token=KMS
```

The PIN provided for the `pktool setpin` operation or by calling `C_Login()` and `C_SetPIN()` functions can be any string of characters with a length between 1 and 256 and no embedded nulls.

### Accessing the Token

After a user initializes their token, they can begin using it with `pktool(1)`, `decrypt(1)`, `encrypt(1)`, or by writing PKCS11 applications and specifying the KMS token.

**Examples** **EXAMPLE 1** Creating a Key on an Oracle Key Manager

The following command creates a key on an Oracle Key Manager:

```
$ pktool genkey token=KMS label=mykey1 keytype=aes keylen=256
```

**EXAMPLE 2** Encrypting a File Using a Key from an Oracle Key Manager

The following command encrypts a file using a key from an Oracle Key Manager:

```
$ encrypt -a aes -K mykey1 -T KMS -i input.txt -o output.enc
```

**EXAMPLE 3** Decrypting a File Using a Key From an Oracle Key Manager

The following command decrypts a file using a key from an Oracle Key Manger:

```
$ decrypt -a aes -K mykey1 -T KMS -i output.enc -o output.txt
```

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	/system/library/security/crypto/pkcs11_kms
Interface Stability	Committed
MT-Level	MT-Safe with Exceptions. See below.
Standard	PKCS#11 v2.20

Exceptions to MT-Safe attribute are documented in section 6.6.2 of RSA PKCS#11 v2.20.

**See Also** [decrypt\(1\)](#), [encrypt\(1\)](#), [pktool\(1\)](#), [cryptoadm\(1M\)](#), [kmscfg\(1M\)](#), [libpkcs11\(3LIB\)](#), [attributes\(5\)](#)

*KMS 2.2: Administration Guide*

*Oracle Key Manager (OKM) Administration Guide*

**Notes** `pkcs11_kms.so` uses a private directory for holding configuration files and other data needed to initialize the connection to a KMA. The private directory is local to the host on which it was first created. By default, the KMS token directory space is in `/var/user/$USERNAME/kms`. The default KMS directory can be overridden by setting the `KMSTOKEN_DIR` environment variable prior to using the [kmscfg\(1M\)](#), [decrypt\(1\)](#), [encrypt\(1\)](#), and [pktool\(1\)](#) commands.

PKCS#11 clients require that Oracle Key Manager Software Version 2.4 be installed.

If PKCS#11 clients use the same Agent ID from multiple systems, that agent should be created without the One Time Passphrase flag set. This option is not be available in OKM clusters with some members running versions of the OKM software prior to 2.4. Refer to the *Oracle Key Manager (OKM) Administration Guide* for assistance in creating Agents.

OKM Agents must have a Default Key Group assigned prior to being used to create keys with a PKCS#11 client. If a Default Key Group is not assigned to the Agent, operations fail with a `CKR_PIN_INCORRECT` error. Refer to the *Oracle Key Manager (OKM) Administration Guide* for assistance in assigning key groups to agents.

**Name** pkcs11\_softtoken – Software RSA PKCS#11 softtoken

**Synopsis** /usr/lib/security/pkcs11\_softtoken.so  
/usr/lib/security/64/pkcs11\_softtoken.so

**Description** The `pkcs11_softtoken.so` object implements the RSA Security Inc. PKCS#11 Cryptographic Token Interface (Cryptoki), v2.20, specification in software. Persistent storage for token objects is provided by this PKCS#11 implementation.

Application developers should link to `libpkcs11.so` rather than link directly to `pkcs11_softtoken.so`. See [libpkcs11\(3LIB\)](#).

The following cryptographic algorithms are implemented: DES, 3DES, AES, Blowfish, RC4, MD5, SHA1, SHA224, SHA256, SHA384, SHA512, RSA, DSA, DH, and ECC.

All of the Standard PKCS#11 functions listed on [libpkcs11\(3LIB\)](#) are implemented except for the following:

```
C_GetObjectSize
C_InitPIN
C_InitToken
C_WaitForSlotEvent
```

A call to these functions returns `CKR_FUNCTION_NOT_SUPPORTED`.

The following RSA PKCS#11 v2.20 mechanisms are supported:

```
CKM_RSA_PKCS_KEY_PAIR_GEN
CKM_RSA_PKCS
CKM_RSA_X_509
```

```
CKM_DSA_KEY_PAIR_GEN
CKM_DSA
CKM_DSA_SHA1
```

```
CKM_DH_PKCS_KEY_PAIR_GEN
CKM_DH_PKCS_DERIVE
```

```
CKM_EC_KEY_PAIR_GEN
CKM_ECDSA
CKM_ECDSA_SHA1
CKM_ECDH1_DERIVE
```

```
CKM_DES_KEY_GEN
CKM_DES_ECB
CKM_DES_CBC
CKM_DES_CBC_PAD
```

```
CKM_DES3_KEY_GEN
CKM_DES3_ECB
```

CKM\_DES3\_CBC  
CKM\_DES3\_CBC\_PAD

CKM\_AES\_KEY\_GEN  
CKM\_AES\_ECB  
CKM\_AES\_CBC  
CKM\_AES\_CBC\_PAD  
CKM\_AES\_CTR

CKM\_BLOWFISH\_KEY\_GEN  
CKM\_BLOWFISH\_CBC

CKM\_RC4\_KEY\_GEN  
CKM\_RC4

CKM\_MD5\_RSA\_PKCS  
CKM\_SHA1\_RSA\_PKCS  
CKM\_SHA224\_RSA\_PKCS  
CKM\_SHA256\_RSA\_PKCS  
CKM\_SHA384\_RSA\_PKCS  
CKM\_SHA512\_RSA\_PKCS

CKM\_MD5  
CKM\_SHA\_1  
CKM\_SHA224  
CKM\_SHA256  
CKM\_SHA384  
CKM\_SHA512

CKM\_MD5\_HMAC  
CKM\_MD5\_HMAC\_GENERAL  
CKM\_SHA\_1\_HMAC  
CKM\_SHA\_1\_HMAC\_GENERAL  
CKM\_SHA224\_HMAC  
CKM\_SHA256\_HMAC  
CKM\_SHA224\_HMAC\_GENERAL  
CKM\_SHA256\_HMAC\_GENERAL  
CKM\_SHA384\_HMAC  
CKM\_SHA384\_HMAC\_GENERAL

CKM\_MD5\_KEY\_DERIVATION  
CKM\_SHA1\_KEY\_DERIVATION  
CKM\_SHA224\_KEY\_DERIVATION  
CKM\_SHA256\_KEY\_DERIVATION  
CKM\_SHA384\_KEY\_DERIVATION  
CKM\_SHA512\_KEY\_DERIVATION

CKM\_SSL3\_PRE\_MASTER\_KEY\_GEN

```

CKM_SSL3_MASTER_KEY_DERIVE
CKM_SSL3_KEY_AND_MAC_DERIVE
CKM_SSL3_MASTER_KEY_DERIVE_DH
CKM_TLS_PRE_MASTER_KEY_GEN
CKM_TLS_MASTER_KEY_DERIVE
CKM_TLS_KEY_AND_MAC_DERIVE
CKM_TLS_MASTER_KEY_DERIVE_DH

```

Each of the following types of key objects has certain token-specific attributes that are set to true by default as a result of object creation, key/key pair generation, and key derivation.

Public key object	CKA_ENCRYPT, CKA_VERIFY, CKA_VERIFY_RECOVER
Private key object	CKA_DECRYPT, CKA_SIGN, CKA_SIGN_RECOVER, CKA_EXTRACTABLE
Secret key object	CKA_ENCRYPT, CKA_DECRYPT, CKA_SIGN, CKA_VERIFY, CKA_EXTRACTABLE

The following certificate objects are supported:

CKC_X_509	For CKC_X_509 certificate objects, the following attributes are supported: CKA_SUBJECT, CKA_VALUE, CKA_LABEL, CKA_ID, CKA_ISSUER, CKA_SERIAL_NUMBER, and CKA_CERTIFICATE_TYPE.
CKC_X_509_ATTR_CERT	For CKC_X_509_ATTR_CERT certificate objects, the following attributes are supported: CKA_OWNER, CKA_VALUE, CKA_LABEL, CKA_SERIAL_NUMBER, CKA_AC_ISSUER, CKA_ATTR_TYPES, and CKA_CERTIFICATE_TYPE.

The search operation of objects matching the template is performed at `C_FindObjectsInit`. The matched objects are cached for subsequent `C_FindObjects` operations.

The `pkcs11_softtoken.so` object provides a filesystem-based persistent token object store for storing token objects. The default location of the token object store is the user's home directory returned by `getpwuid_r()`. The user can override the default location by using the `${SOFTTOKEN_DIR}` environment variable.

If the token object store has never been initialized, the `C_Login()` function might return `CKR_OK` but the user is not able to create, generate, derive or find any private token object and receives `CKR_PIN_EXPIRED`.

The user must use the `pktool(1) setpin` command with the default passphrase “changeme” as the old passphrase to change the passphrase of the object store. This action is needed to initialize and set the passphrase to a newly created token object store.

After logging into object store with the new passphrase that was set by the `pktool setpin` command, the user can create and store the private token object in this newly created object store. Until the token object store is initialized by `setpin`, the `C_Login()` function is allowed, but all attempts by the user to create, generate, derive or find any private token object fails with a `CKR_PIN_EXPIRED` error.

The PIN provided for `C_Login()` and `C_SetPIN()` functions can be any string of characters with lengths between 1 and 256 and no embedded nulls.

The default location of the token object store is `/var/user/$USERNAME/pkcs11_softtoken`.

The user can override the default location by using the `${SOFTTOKEN_DIR}` environment variable. The location for the alternate token object store is `${SOFTTOKEN_DIR}/pkcs11_softtoken/`.

**Return Values** The return values for each of the implemented functions are defined and listed in the RSA PKCS#11 v2.20 specification. See <http://www.rsasecurity.com>

**Files**

<code>/var/user/\$USERNAME/pkcs11_softtoken</code>	user's default token object store
<code>\${SOFTTOKEN_DIR}/pkcs11_softtoken</code>	alternate token object store

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT-Level	MT-Safe with exceptions. See section 6.6.2 of RSA PKCS#11 v2.20.
Standard	PKCS#11 v2.20

**See Also** [pktool\(1\)](#), [cryptoadm\(1M\)](#), [libpkcs11\(3LIB\)](#), [attributes\(5\)](#), [pkcs11\\_kernel\(5\)](#)  
 RSA PKCS#11 v2.20 <http://www.rsasecurity.com>

**Name** pkcs11\_tpm – RSA PKCS#11 token for Trusted Platform Modules (TPM)

**Synopsis** /usr/lib/security/pkcs11\_tpm.so  
 /usr/lib/security/64/pkcs11\_tpm.so

**Description** The `pkcs11_tpm.so` object implements the RSA Security Inc. PKCS#11 Cryptographic Token Interface (Cryptoki), v2.20, specification using Trusted Computing Group protocols to talk to a TPM security device. This provider implements the PKCS#11 specification and uses the TCG Software Stack (TSS) APIs in the `SUNWtss` package.

Application developers should link to `libpkcs11.so.1` rather than link directly with `pkcs11_tpm.so`. See [libpkcs11\(3LIB\)](#).

The following cryptographic algorithms are implemented: RSA, SHA1, and MD5.

All of the standard PKCS#11 functions listed in [libpkcs11\(3LIB\)](#) are implemented except for the following:

```
C_EncryptUpdate
C_EncryptFinal
C_DecryptUpdate
C_DecryptFinal
C_DigestEncryptUpdate
C_DecryptDigestUpdate
C_SignEncryptUpdate
C_DecryptVerifyUpdate
C_GetFunctionStatus
C_CancelFunction
C_WaitForSlotEvent
C_GenerateKey
C_DeriveKey
```

The following RSA PKCS#11 v2.20 mechanisms are supported:

```
CKM_RSA_PKCS_KEY_PAIR_GEN
CKM_RSA_PKCS
CKM_RSA_PKCS_OAEP
CKM_RSA_X_509
CKM_MD5_RSA_PKCS
CKM_SHA1_RSA_PKCS
CKM_SHA_1
CKM_SHA_1_HMAC
CKM_SHA_1_HMAC_GENERAL
CKM_MD5
CKM_MD5_HMAC
CKM_MD5_HMAC_GENERAL
```

**Per-User Initialization** The `pkcs11_tpm` provider can only be used on a system which has a TPM device and which also has the `SUNWtss` package installed. If those prerequisites are met, users can create their own private tokens using `pktool(1)`, which will allow them to perform operations using the TPM device and protect their private data with TPM-protected keys.

To prepare and initialize a user's TPM token, the following steps must be performed:

1. Initialize the token.
2. Set the SO (security officer) PIN.
3. Set the user's unique PIN.

Initializing the token is done using the `pktool(1)` command as follows:

```
$ pktool inittoken currLabel=TPM newLabel=tpm/myname
```

- By default, an uninitialized TPM is recognized by the name `TPM`. When a user initializes their own private token, it can either be renamed to something else (for example, `tpm/joeuser`) or kept as `TPM` (in which case the `newLabel` argument would be omitted).
- The user will have to supply the default SO PIN before being able to initialize his or her token. The default SO PIN is 87654321. It is changed in step 2, above.

Once the token is initialized, the SO and user PINs must be changed from the default values. Again, `pktool(1)` is used to change these PIN values.

Changing the SO PIN:

```
$ pktool setpin token=tpm/joeuser so
```

The `so` option indicates that this “setpin” operation is to change the SO PIN and must be present. The user must then enter the default SO PIN (87654321) and then enter (and confirm) a new PIN.

Once the SO PIN is reset from the default, the user's unique PIN must also be changed.

Changing the user's PIN:

```
$ pktool setpin token=tpm/joeuser
```

The default PIN for a non-SO user is 12345678. The user must enter the default PIN and then enter (and confirm) a new, unique PIN.

The PIN provided for the `pktool setpin` operation or by calling `C_Login()` and `C_SetPIN()` functions can be any string of characters with a length between 1 and 256 and no embedded nulls.

**Accessing the Token** After a user initializes their token, they can begin using it with `pktool(1)` or by writing PKCS11 applications and locating the token using the name created above (`tpm/joeuser` in the examples above).

Examples:

```
$ pkctool gencert token=tpm/joeuser -i
$ pkctool list token=tpm/joeuser
```

**Notes** `pkcs11_tpm.so` provides object storage in a filesystem-specific token object storage area. Private objects are protected by encryption with private keys and can only be decrypted by loading the token's private key into the TPM and performing the decryption entirely in the TPM. The user's private key is generated by the TPM when the user sets their personal PIN (see above). The keys for both the SO and users are stored in the TSS persistent storage database and are referenced by a unique UUID value. All user tokens have a unique SO key and unique user key so that the PINs for one user's token will not unlock private data in another user's token on the same machine.

Each TPM is unique and the token keys created on one TPM may not be used on another TPM. The `pkcs11_tpm.so` token data is all managed on the system where the TPM resides and may not be moved to other systems. If the TPM is reset and the SRK (Storage Root Key) is changed, all of the keys previously generated for that TPM will no longer be valid.

`pkcs11_tpm.so` creates a private workspace to manage administrative files for each token created. By default, this area is created as `/var/user/$USERNAME/tpm/`. However, users may override this by setting the `PKCS11_TPM_DIR` environment variable prior to initializing or using the token.

**Return Values** The return values for each of the implemented functions are defined and listed in the RSA PKCS#11 v2.20 specification. See <http://www.rsasecurity.com>.

**Files** `/var/user/$USERNAME/tpm/`  
User's default token object store.

`${PKCS11_TPM_DIR}`  
Alternate token object store.

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Interface Stability	Committed
MT-Level	MT-Safe with Exceptions (see below)
Standard	PKCS#11 v2.20

Exceptions to MT-Safe attribute are documented in section 6.6.2 of RSA PKCS#11 v2.20.

**See Also** [pkctool\(1\)](#), [cryptoadm\(1M\)](#), [libpkcs11\(3LIB\)](#), [attributes\(5\)](#)

TCG Software Stack (TSS) Specifications, <https://www.trustedcomputinggroup.org/specs/TSS> (as of the date of publication)

RSA PKCS#11 v2.20, <http://www.rsasecurity.com>

**Name** pkg – Image Packaging System

**Description** The image packaging system, pkg(5), is a framework that provides for software lifecycle management (installation, upgrade, and removal). Image packaging manages software in units of packages, which are collections of actions, defined by a set of key/value pairs and possibly a data payload. In many cases, actions are files found in a file system, but they also represent other installable objects, such as drivers, services, and users.

**Package Fmris and Versions** Each package is represented by a fault management resource identifier (FMRI) with the scheme `pkg:.` The full FMRI for a package consists of the scheme, a publisher, the package name, and a version string in the following format:

```
pkg://solaris/system/library/c++-runtime@0.5.11,5.11-0.175.0.0.0.2.1:20120921T190358Z
```

`solaris` is the publisher. `system/library/c++-runtime` is the package name. Although the namespace is hierarchical and arbitrarily deep, there is no enforced containment; the name is essentially arbitrary. The publisher information is optional, but must be preceded by `pkg://` if present. An FMRI that includes the publisher is often referred to as being “fully qualified.” If publisher information is not present, then the package name should generally be preceded by `pkg:/.`

Packaging clients often allow the scheme of an FMRI to be omitted if it does not contain publisher information. For example, `pkg:/system/library/c++-runtime` can be written as `system/library/c++-runtime`. If the scheme is omitted, clients also allow omission of all but the last component of a package name for matching purposes. For example, `system/library/c++-runtime` could be written as `library/c++-runtime` or `c++-runtime`, which would then match packages named `c++-runtime` or package names ending in `/c++-runtime`.

A publisher name identifies a person, group of persons, or an organization as the source of one or more packages. To avoid publisher name collisions and help identify the publisher, a best practice is to use a domain name that represents the entity publishing the packages as a publisher name.

The version follows the package name, separated by an at sign (`@`). The version consists of four sequences of numbers, separated by punctuation. The elements in the first three sequences are separated by dots, and the sequences are arbitrarily long. Leading zeros in version components (for example, `01.1` or `1.01`) are not permitted. Trailing zeros (for example, `1.10`) are permitted.

The first part of the version is the component version. For components tightly bound to the operating system, this is usually the value of `uname -r` for that version of the operating system. For a component with its own development lifecycle, this sequence is a dotted release number, such as `2.4.10`.

The second part of the version, which if present must follow a comma (,), is the build version. The build version specifies what version of the operating system the contents of the package were built on, providing a minimum bound on which operating system version the contents can be expected to run successfully.

The third part of the version, which if present must follow a hyphen (-), is the branch version. The branch version is a versioning component that provides vendor-specific information. The branch version can be incremented when the packaging metadata is changed, independently of the component version. The branch version might contain a build number or other information.

The fourth part of the version, which if present must follow a colon (:), is a timestamp. The timestamp represents when the package was published.

When performing comparisons between versions, no component of the full version is considered unless the components to its left are the same. Thus, “4.3-1” is greater than “4.2-7” because “4.3” is greater than “4.2”, and “4.3-3” is greater than “4.3-1” because “3” is greater than “1”.

Many parts of the system, when appropriate, abridge FMRI when displaying them, and accept input in shorter forms to reduce the volume of information displayed or required. Typically, the scheme, publisher, build version, and timestamp can be elided. Sometimes all of the versioning information can be omitted.

**Actions** Actions represent the installable objects on a system. Actions are described in the manifest of a package. Every action consists primarily of its name and a key attribute. Together, these refer to a unique object as it follows a version history. Actions can have other attributes. Some attributes are interpreted directly by the packaging system. Other attributes might be useful only to the system administrator or the end-user.

Actions have a simple text representation:

```
action_name attribute1=value1 attribute2=value2 . . .
```

Names of attributes cannot have whitespace, quotation marks, or equals signs (=) in them. All characters after the first equals sign belong to the value. Values can have all of those, though spaces must be enclosed in single or double quotation marks. Single quotation marks do not need to be escaped inside a string that is enclosed in double quotation marks, and double quotation marks do not need to be escaped inside a string that is enclosed in single quotation marks. A quotation mark can be prefixed with a backslash (\) character to avoid terminating the quoted string. A backslash can be escaped with a backslash.

Attributes can be named more than once with multiple values. These are treated as unordered lists.

Actions with many attributes can create long lines in a manifest file. Such lines can be wrapped by terminating each incomplete line with a backslash. Note that this continuation character must occur between attribute/value pairs. Neither attributes nor their values nor the combination can be split.

The attributes listed below are not an exhaustive set. In fact, the attributes that can be attached to an action are arbitrary, and the standard sets of attributes are easy to augment to incorporate future developments.

Certain action attributes cause additional operations to be executed outside of the packaging context. These attributes are documented in the “Actuators” section below.

**File Actions** The `file` action represents an ordinary file. The `file` action references a payload, and has four standard attributes:

- `path` The file system path where the file is installed. This is a `file` action's key attribute.
- `mode` The access permissions (in numeric form) of the file. These are simple permissions only, not ACLs.
- `owner` The name of the user that owns the file.
- `group` The name of the group that owns the file.

The payload is a positional attribute in that it is not named. It is the first word after the action name. In a published manifest, it is the SHA-1 hash of the file contents. If present in a manifest that has yet to be published, it represents the path where the payload can be found. See `pkgsend(1)`. The hash attribute can be used instead of the positional attribute, should the value include an equals sign. Both can be used in the same action. However, the hashes must be identical.

Other attributes include:

`preserve`

This specifies that the file's contents should not be overwritten on upgrade if the contents are determined to have changed since the file was installed or last upgraded. On initial installs, if an existing file is found, the file is salvaged (stored in `/var/pkg/lost+found`).

If the value of `preserve` is `renameold`, then the existing file is renamed with the extension `.old`, and the new file is put in its place.

If the value of `preserve` is `renamew`, then the existing file is left alone, and the new file is installed with the extension `.new`.

If the value of `preserve` is `legacy`, then this file is not installed for initial package installs. On upgrades, any existing file is renamed with the extension `.legacy`, and then the new file is put in its place.

If the value of `preserve` is `true` (or a value not listed above, such as `strawberry`), then the existing file is left alone, and the new file is not installed.

### overlay

This specifies whether the action allows other packages to deliver a file at the same location or whether it delivers a file intended to overlay another. This functionality is intended for use with configuration files that do not participate in any self-assembly (for example, `/etc/motd`) and that can be safely overwritten.

If `overlay` is not specified, multiple packages cannot deliver files to the same location.

If the value of `overlay` is `allow`, one other package is allowed to deliver a file to the same location. This value has no effect unless the `preserve` attribute is also set.

If the value of `overlay` is `true`, the file delivered by the action overwrites any other action that has specified `allow`. Changes to the installed file are preserved based on the value of the `preserve` attribute of the overlaying file. On removal, the contents of the file are preserved if the action being overlaid is still installed, regardless of whether the `preserve` attribute was specified. Only one action can overlay another, and the mode, owner, and group attributes must match.

Files can also be “tasted,” and depending on the flavor, can have additional attributes. For ELF files, the following attributes are recognized:

### elfarch

The architecture of the ELF file. This is the output of `uname -p` on the architecture for which the file is built.

### elfbits

This is 32 or 64.

### elfhash

This is the hash of the “interesting” ELF sections in the file. These are the sections that are mapped into memory when the binary is loaded. These are the only sections necessary to consider when determining whether the executable behavior of two binaries will differ.

### original\_name

This attribute is used to handle editable files moving from package to package or from place to place, or both. The form this takes is the name of the originating package, followed by a colon and the original path to the file. Any file being deleted is recorded either with its package and path, or with the value of the `original_name` attribute if specified. Any editable file being installed that has the `original_name` attribute set uses the file of that name if it is deleted as part of the same packaging operation.

### release-note

This attribute is used to indicate that this file contains release note text. The value of this attribute is a package FMRI. If the FMRI specifies a package name that is present in the original image and a version that is newer than the version of the package in the original image, this file will be part of the release notes. A special FMRI of `feature/pkg/self` refers to the containing package. If the version of `feature/pkg/self` is 0, this file will only be part of the release notes on initial installation.

**revert - tag**

This attribute is used to tag editable files that should be reverted as a set. Multiple `revert - tag` values can be specified. The file reverts to its manifest-defined state when `pkg revert` is invoked with any of those tags specified. See `pkg(1)`.

**Directory Actions** The `dir` action is like the `file` action in that it represents a file system object. The `dir` action represents a directory instead of an ordinary file. The `dir` action has the same four standard attributes as the `file` action, and `path` is the key attribute.

Directories are reference counted in IPS. When the last package that either explicitly or implicitly references a directory no longer does so, that directory is removed. If that directory contains unpackaged file system objects, those items are moved into `$IMAGE_META/lost+found`. See the “Files” section for more information about `$IMAGE_META`.

To move unpackaged contents into a new directory, the following attribute might be useful:

**salvage - from**

This names a directory of salvaged items. A directory with such an attribute inherits on creation the salvaged directory contents if they exist.

**Link Actions** The `link` action represents a symbolic link. The `link` action has the following standard attributes:

**path**

The file system path where the symbolic link is installed. This is a `link` action's key attribute.

**target**

The target of the symbolic link. The file system object to which the link resolves.

**mediator**

Specifies the entry in the mediation namespace shared by all path names participating in a given mediation group (for example, `python`). Link mediation can be performed based on `mediator - version` and/or `mediator - implementation`. All mediated links for a given path name must specify the same mediator. However, not all mediator versions and implementations need to provide a link at a given path. If a mediation does not provide a link, then the link is removed when that mediation is selected. A `mediator`, in combination with a specific version and/or implementation represents a mediation that can be selected for use by the packaging system.

**mediator - version**

Specifies the version (expressed as a dot-separated sequence of nonnegative integers) of the interface described by the `mediator` attribute. This attribute is required if `mediator` is specified and `mediator - implementation` is not. A local system administrator can set the version to use explicitly. The value specified should generally match the version of the package delivering the link (for example, `runtime/python - 26` should use `mediator - version=2.6`), although this is not required.

`mediator-implementation`

Specifies the implementation of the mediator for use in addition to or instead of the `mediator-version`. Implementation strings are not considered to be ordered and a string is arbitrary selected by `pkg(5)` if not explicitly specified by a system administrator.

The value can be a string of arbitrary length composed of alphanumeric characters and spaces. If the implementation itself can be versioned or is versioned, then the version should be specified at the end of the string, after a `@` (expressed as a dot-separated sequence of nonnegative integers). If multiple versions of an implementation exist, the default behavior is to select the implementation with the greatest version.

If only one instance of an implementation mediation link at a particular path is installed on a system, then that one is chosen automatically. If future links at the path are installed, the link is not switched unless a vendor, site, or local override applies, or if one of the links is version mediated.

`mediator-priority`

When resolving conflicts in mediated links, `pkg(5)` normally chooses the link with the greatest value of `mediator-version` or based on `mediator-implementation` if that is not possible. This attribute is used to specify an override for the normal conflict resolution process.

If this attribute is not specified, the default mediator selection logic is applied.

If the value is `vendor`, the link is preferred over those that do not have a `mediator-priority` specified.

If the value is `site`, the link is preferred over those that have a value of `vendor` or that do not have a `mediator-priority` specified.

A local system administrator can override the selection logic described above.

**Hardlink Actions** The `hardlink` action represents a hard link. It has the same attributes as the `link` action, and `path` is also its key attribute.

**Driver Actions** The `driver` action represents a device driver. The `driver` action does not reference a payload. The driver files themselves must be installed as `file` actions. The following attributes are recognized (see `add_drv(1M)` for more information):

`name`

The name of the driver. This is usually, but not always, the file name of the driver binary. This is the `driver` action's key attribute.

`alias`

This represents an alias for the driver. A given driver can have more than one `alias` attribute. No special quoting rules are necessary.

`class`

This represents a driver class. A given driver can have more than one `class` attribute.

**perms**

This represents the file system permissions for the driver's device nodes.

**clone\_perms**

This represents the file system permissions for the clone driver's minor nodes for this driver.

**policy**

This specifies additional security policy for the device. A given driver can have more than one **policy** attribute, but no minor device specification can be present in more than one attribute.

**privs**

This specifies privileges used by the driver. A given driver can have more than one **privs** attribute.

**devlink**

This specifies an entry in `/etc/devlink.tab`. The value is the exact line to go into the file, with tabs denoted by `\t`. See `devlinks(1M)` for more information. A given driver can have more than one **devlink** attribute.

**Depend Actions** The **depend** action represents an inter-package dependency. A package can depend on another package because the first requires functionality in the second for the functionality in the first to work, or even to install. Dependencies can be optional. If a dependency is not met at the time of installation, the packaging system attempts to install or update the dependent package to a sufficiently new version, subject to other constraints.

The following attributes are recognized:

**fMRI**

The FMRI representing the dependent package. This is the dependency action's key attribute. The **fMRI** value must not include the publisher. The package name is assumed to be complete. Dependencies of type **require-any** can have multiple **fMRI** attributes. A version is optional on the **fMRI** value, though for some types of dependencies, an **fMRI** with no version has no meaning.

**type**

The type of the dependency.

**require**

The dependency is required and must have a version equal to or greater than the version specified in the **fMRI** attribute. If the version is not specified, any version satisfies the dependency. A package cannot be installed if any of its required dependencies cannot be satisfied.

**optional**

The dependency, if present, must be at the specified version level or greater.

**exclude**

The containing package cannot be installed if the dependency is present at the specified version level or greater. If no version is specified, the dependent package cannot be installed concurrently with the package specifying the dependency.

**incorporate**

The dependency is optional, but the version of the dependent package is constrained. See “Constraints and Freezing” below.

**require-any**

Any one of multiple dependent packages as specified by multiple `fMRI` attributes can satisfy the dependency, following the same rules as the `require` dependency type.

**conditional**

The dependency is required only if the package defined by the `predicate` attribute is present on the system.

**origin**

The dependency must, if present, be at the specified value or better on the image to be modified prior to installation. If the value of the `root-image` attribute is `true`, the dependency must be present on the image rooted at `/` in order to install this package.

**group**

The dependency is required unless the package is on the image avoid list. Note that obsolete packages silently satisfy the group dependency. See the `avoid` subcommand in `pkg(1)`.

**parent**

The dependency is ignored if the image is not a child image. If the image is a child image then the dependency is required to be present in the parent image. The package version matching for a parent dependency is the same as that used for `incorporate` dependencies.

**predicate**

The FMRI representing the predicate for conditional dependencies.

**root-image**

Has an effect only for `origin` dependencies as mentioned above.

**License Actions** The `license` action represents a license or other informational file associated with the package contents. A package can deliver licenses, disclaimers, or other guidance to the package installer through the use of the `license` action.

The payload of the `license` action is delivered into the image metadata directory related to the package, and should only contain human-readable text data. It should not contain HTML or any other form of markup. Through attributes, `license` actions can indicate to clients that the related payload must be displayed and/or require acceptance of it. The method of display and/or acceptance is at the discretion of clients.

The following attributes are recognized:

#### `license`

This is a `license` action's key attribute. This attribute provides a meaningful description for the license to assist users in determining the contents without reading the license text itself. Some example values include:

- ABC Co. Copyright Notice
- ABC Co. Custom License
- Common Development and Distribution License 1.0 (CDDL)
- GNU General Public License 2.0 (GPL)
- GNU General Public License 2.0 (GPL) Only
- MIT License
- Mozilla Public License 1.1 (MPL)
- Simplified BSD License

The `license` value must be unique within a package. Including the version of the license in the description, as shown in several of the examples above, is recommended. If a package has code under multiple licenses, use multiple `license` actions. The length of the license attribute value should not be more than 64 characters.

#### `must-accept`

When `true`, this license must be accepted by a user before the related package can be installed or updated. Omission of this attribute is equivalent to `false`. The method of acceptance (interactive or configuration-based, for example) is at the discretion of clients.

#### `must-display`

When `true`, the action's payload must be displayed by clients during packaging operations. Omission of this value is equivalent to `false`. This attribute should not be used for copyright notices, only actual licenses or other material that must be displayed during operations. The method of display is at the discretion of clients.

**Legacy Actions** The `legacy` action represents package data used by a legacy packaging system. The attributes associated with this action are added into the legacy system's databases so that the tools querying those databases can operate as if the legacy package were actually installed. In particular, this should be sufficient to convince the legacy system that the package named by the `pkg` attribute is installed on the system, so that the package can be used to satisfy dependencies.

The following attributes, named in accordance with the parameters on `pkginfo(4)`, are recognized:

#### `category`

The value for the `CATEGORY` parameter. The default value is `system`.

#### `desc`

The value for the `DESC` parameter.

**hotline**

The value for the `HOTLINE` parameter.

**name**

The value for the `NAME` parameter. The default value is none provided.

**pkg**

The abbreviation for the package being installed. The default value is the name from the FMRI of the package. This is a legacy action's key attribute.

**vendor**

The value for the `VENDOR` parameter.

**version**

The value for the `VERSION` parameter. The default value is the version from the FMRI of the package.

**Set Actions** The set action represents a package-level attribute, or metadata, such as the package description.

The following attributes are recognized:

**name** The name of the attribute.

**value** The value given to the attribute.

The set action can deliver any metadata the package author chooses. However, there are a number of well defined attribute names that have specific meaning to the packaging system.

**pkg.fmri**

See “Package FMRI and Versions” in the “Description” section.

**info.classification**

One or more tokens that a `pkg(5)` client can use to classify the package. The value should have a scheme (such as “`org.opensolaris.category.2008`” or “`org.acm.class.1998`”) and the actual classification, such as “Applications/Games”, separated by a colon (:).

**pkg.description**

A detailed description of the contents and functionality of the package, typically a paragraph or so in length.

**pkg.obsolete**

When `true`, the package is marked obsolete. An obsolete package can have no actions other than more set actions, and must not be marked renamed.

**pkg.renamed**

When `true`, the package has been renamed. There must be one or more `depend` actions in the package as well that point to the package versions to which this package has been renamed. A package cannot be marked both renamed and obsolete, but otherwise can have any number of set actions.

`pkg.summary`

A short, one-line description of the package.

**Group Actions** The group action defines a UNIX group as defined in `group(4)`. No support is present for group passwords. Groups defined with this action initially have no user list. Users can be added with the user action. The following attributes are recognized:

`groupname`

The value for the name of the group.

`gid`

The group's unique numerical id. The default value is the first free group under 100.

**User Actions** The user action defines a UNIX user as defined in `/etc/passwd`, `/etc/shadow`, `/etc/group`, and `/etc/ftpd/ftpusers` files. Users defined with this attribute have entries added to the appropriate files.

The following attributes are recognized:

`username`

The unique name of the user

`password`

The encrypted password of the user. Default value is `*LK*`. See `shadow(4)`.

`uid`

The unique uid of the user. Default value is first free value under 100.

`group`

The name of the user's primary group. Must be found in `/etc/group`.

`gcos-field`

The value of the `gcos` field in `/etc/passwd`. Default value is `username`.

`home-dir`

The user's home directory. Default value is `/`.

`login-shell`

The user's default shell. Default value is empty.

`group-list`

Secondary groups to which the user belongs. See `group(4)`.

`ftpuser`

Can be set to `true` or `false`. The default value of `true` indicates that the user is permitted to login via FTP. See `ftpusers(4)`.

`lastchg`

The number of days between January 1, 1970, and the date that the password was last modified. Default value is empty. See `shadow(4)`.

**min**

The minimum number of days required between password changes. This field must be set to 0 or above to enable password aging. Default value is empty. See `shadow(4)`.

**max**

The maximum number of days the password is valid. Default value is empty. See `shadow(4)`.

**warn**

The number of days before password expires that the user is warned. See `shadow(4)`.

**inactive**

The number of days of inactivity allowed for that user. This is counted on a per-machine basis. The information about the last login is taken from the machine's `last log` file. See `shadow(4)`.

**expire**

An absolute date expressed as the number of days since the UNIX Epoch (January 1, 1970). When this number is reached, the login can no longer be used. For example, an expire value of 13514 specifies a login expiration of January 1, 2007. See `shadow(4)`.

**flag**

Set to empty. See `shadow(4)`.

**Actuators** In certain contexts, additional operations can be appropriate to execute in preparation for or following the introduction of a particular action. These additional operations are generally needed only on a live system image, and are operating system specific. When multiple actions involved in a package installation or removal have identical actuators, then the operation corresponding to actuator presence is executed once for that installation or removal.

Incorrectly specified actuators can result in package installation failure, if the actuator cannot determine a means of making safe installation progress.

The following actuators are defined:

**reboot - needed**

Can be set to `true` or `false`. This actuator declares that update or removal of the tagged action must be performed in a new boot environment if the package system is operating on a live image. Creation of a new boot environment is controlled by the `be-policy` image property. See the “Image Properties” section in the `pkg(1)` man page for more information about the `be-policy` property.

**disable\_fmri, refresh\_fmri, restart\_fmri, suspend\_fmri**

Each of these actuators takes the value of an FMRI of a service instance to operate on during the package installation or removal. `disable_fmri` causes the given FMRI to be disabled prior to action removal, per the `disable` subcommand to `svcadm(1M)`. `refresh_fmri` and `restart_fmri` cause the given FMRI to be refreshed or restarted after action installation, update, or removal per the respective subcommands of `svcadm(1M)`.

Finally, `suspend_fmri` causes the given FMRI to be disabled temporarily prior to the action install phase, and then enabled after the completion of that phase.

The value can contain a pattern that matches multiple service instances. However, it must do so explicitly with a glob as accepted by `svcs(1)`, rather than doing so implicitly by not indicating any instances.

**Mediations** A mediator is a name that represents a set of related symbolic or hard links. If two or more link actions have the same path and mediator name, the user or the package system selects the link target based on version, implementation, or priority. See “Link Actions” for information about mediator attributes.

The following example shows two different instances of a mediator named `java` where the link choices are between versions. These two link actions would appear in two different packages.

```
link mediator=java mediator-version=1.6 path=usr/java target=jdk/jdk1.6.0_31
link mediator=java mediator-version=1.7 path=usr/java target=jdk/jdk1.7.0_02
```

See the `set-mediator` subcommand in the `pkg(1)` man page for information about how to select the version you want for this link path. To have a choice of versions, both packages must be installed.

**Constraints and Freezing** When a package is transitioned to a new version, or when it is added to or removed from the system, the version that is chosen, or whether removal is allowed, is determined by a variety of constraints put on the package. Those constraints can be defined by other packages in the form of dependencies, or by the administrator in the form of freezes.

The most common form of constraint is delivered by the `require` dependency, as described in “Depend Actions” above. Such a constraint prevents the package from being downgraded or removed.

Most parts of the operating system are encapsulated by packages called *incorporations*. These packages primarily deliver constraints represented by the `incorporate` dependency.

As described above, an incorporated package need not be present on the system, but if it is, then it specifies both an inclusive minimum version and an exclusive maximum version. For example, if the dependent FMRI has a version of 1.4.3, then no version less than 1.4.3 would satisfy the dependency, and neither would any version greater than or equal to 1.4.4. However, versions that merely extended the dotted sequence, such as 1.4.3.7, would be allowed.

Incorporations are used to force parts of the system to upgrade synchronously. For some components, such as the C library and the kernel, this is a basic requirement. For others, such as a simple userland component on which nothing else has a dependency, the synchronous upgrade is used merely to provide a known and tested set of package versions that can be referred to by a particular version of the incorporation.

---

Since an incorporation is simply a package, it can be removed, and all the constraints it delivers are therefore relaxed. However, many of the incorporations delivered by Oracle Solaris are required by the packages they incorporate because that relaxation would not be safe.

Attempting an upgrade of a package to a version that is not allowed by an installed incorporation will not attempt to find a newer version of the incorporation in order to satisfy the request, but will instead fail. If the constraint itself must be moved, and the incorporation specifying it cannot be removed, then the incorporation must be upgraded to a version that specifies a desired version of the constraint. Upgrading an incorporation causes all of the incorporated packages that would not satisfy the constraints delivered by the new version to be upgraded as well.

A system administrator can constrain a package by using the `pkg freeze` command. The named package is constrained to the version installed on the system if no version is provided. If a versioned package is provided, then this administrative constraint, or freeze, acts as if an incorporate dependency were installed where the `fMRI` attribute had the value of the provided package version.

A freeze is never lifted automatically by the packaging system. To relax a constraint, use the `pkg unfreeze` command.

## Publishers and Repositories

As detailed above, a publisher is simply a name that package clients use to identify the provider of packages. Publishers can distribute their packages using package repositories and/or package archives. There are two types of repositories currently supported by the package system: origin repositories and mirror repositories.

An *origin* is a package repository that contains all of the metadata (such as catalogs, manifests, and search indexes) and content (files) for one or more packages. If multiple origins are configured for a given publisher in an image, the package client API attempts to choose the best origin to retrieve package data from. This is the most common type of repository, and is implicitly created whenever `pkg send` or `pkg recv` is used on a package repository.

A *mirror* is a package repository that contains only package content (files). If one or more mirrors are configured for a given publisher in an image, the client API prefers the mirrors for package content retrieval and attempts to choose the best one to retrieve package content from. If the mirror is unreachable, does not have the required content, or is slower, the client API retrieves the content from any configured origin repositories. Mirrors are intended for content sharing among a trusted set of clients using the dynamic mirror functionality of `pkg.depotd(1M)`. Mirrors are also intended to be used to authenticate access to package metadata, but distribute the package content without authentication. For example, a client might be configured with an `https` origin that requires an SSL key and certificate pair to access, and with an `http` mirror that provides the package content. In this way, only authorized clients can install or update the packages, while the overhead of authentication for package content retrieval is avoided. A mirror can be created by removing all subdirectories of a repository except those named `file` and their parents. An origin repository can be also be provisioned as a mirror by using the mirror mode of `pkg.depotd(1M)`.

## Global and Non-Global Zone Update

The pkg system forces non-global zones to be kept in sync with the global zone. This means that certain packages must be at the same version in the global zone and all non-global zones to ensure the same kernel is run. To do this, pkg uses parent dependencies to impose certain constraints on non-global zones. See “Depend Actions” above for more information about parent dependencies.

Because of restrictions that the global zone imposes on non-global zones, the non-global zones must have access to the packages of the global zone and must have a similar publisher configuration. Both of these objectives are achieved by using a **system repository** (see the `pkg.sysrepo(1M)` man page). The system repository provides access to the publishers configured in the global zone and information about how those publishers are configured. To prevent non-global zones from choosing different packages during installation or update, system publishers are ranked higher in the publisher search order than publishers configured in the non-global zone. See the `pkg set-publisher` command in the `pkg(1)` man page for information about publisher search order.

To update all non-global zones on the system, use the `pkg update` command with no arguments in the global zone. This command operates on the global zone and on each non-global zone recursively. The minimal changes necessary are made to non-global zones to bring them in sync with the changes made in the global zone. For example, suppose package `foo` is installed at version 1 in both the global zone and non-global zones, and suppose version 2 is available in a system repository. If `foo` has a parent dependency, then `pkg update foo` updates `foo` to version 2 in both the global zone and the non-global zones because the parent dependency forces the package to stay in sync. If `foo` does not have a parent dependency, then `foo` is updated to version 2 in the global zone but remains at version 1 in the non-global zones.

## Facets and Variants

Software can have components that are optional and components that are mutually exclusive. Examples of optional components include locales and documentation. Examples of mutually exclusive components include SPARC or x86 and debug or non-debug binaries.

In IPS, optional components are called *facets* and mutually exclusive components are called *variants*. Facets and variants are specified as tags on package actions. Each facet and variant tag has a name and a value. A single action can have multiple facet and variant tags. Examples of components with multiple facet and variant tags include an architecture-specific header file that is used by developers, or a component that is only for a SPARC global zone.

An example of a variant tag is `variant.arch=sparc`. An example of a facet tag is `facet.devel=true`. Facets and variants are often referred to without the leading `facet.` and `variant.`.

Facets and variants are special properties of the image and cannot be set on individual packages. To view the current values of the facets and variants set on the image, use the `pkg facet` and `pkg variant` commands as shown in the `pkg(1)` man page. To modify the values of the facets and variants set on the image, use the `pkg change-facet` and `pkg change-variant` commands.

Facets are boolean: They can be set only to `true` (enabled) or `false` (disabled). By default, all facets are considered to be set to `true` in the image. A facet tag on an action should only have the value `true`; other values have undefined behavior. A facet set on the image can be a full facet such as `doc.man` or a pattern such as `locale.*`. This is useful when you want to disable a portion of the facet namespace, and only enable individual facets within it. For example, you could disable all locales and then only enable one or two specific locales, as shown in the following example:

```
# pkg change-facet locale.*=false
[output about packages being updated]
# pkg change-facet locale.en_US=true
[output about packages being updated]
```

Most variants can have any number of values. For example, the `arch` variant can be set to `i386`, `sparc`, `ppc`, `arm`, or whatever architectures the distribution supports. (Only `i386` and `sparc` are used in Oracle Solaris.) The exception are the debug variants. The debug variants can only be set to `true` or `false`; other values have undefined behavior. If a file action has both non-debug and debug versions, both versions must have the applicable debug variant explicitly set, as shown in the following example:

```
file group=sys mode=0644 overlay=allow owner=root \
  path=etc/motd pkg.csize=115 pkg.size=103 preserve=true \
  variant.debug.osnet=true

file group=sys mode=0644 overlay=allow owner=root \
  path=etc/motd pkg.csize=68 pkg.size=48 preserve=true \
  variant.debug.osnet=false
```

The variant value must be set on the image in order for a package using the variant to be installed. The `arch` and `zone` variants are set by the program that creates the image and installs its initial contents. The `debug.*` variants are `false` in the image by default.

The facets and variants set on the image affect whether a particular action is installed.

- Actions with no facet or variant tags are always installed.
- Actions with facet tags are installed unless all of the facets or facet patterns matching the tags are set to `false` on the image. If any facet is set to `true` or is not explicitly set (`true` is the default), then the action is installed.
- Actions with variant tags are installed only if the values of all the variant tags are the same as the values set in the image.
- Actions with both facet and variant tags are installed if both the facets and the variants allow the action to be installed.

You can create your own facet and variant tags. The following tags are in common use in Oracle Solaris.

Variant Name	Possible Values
variant.arch	sparc, i386
variant.opensolaris.zone	global, nonglobal
variant.debug.*	true, false

The following list shows a small sample of the facet tags that are used in Oracle Solaris:

```

facet.devel          facet.doc
facet.doc.html       facet.doc.info
facet.doc.man        facet.doc.pdf
facet.locale.de      facet.locale.en_GB
facet.locale.en_US   facet.locale.fr
facet.locale.ja_JP   facet.locale.zh_CN

```

**Image Policies** Image policies are defined by image properties with boolean values. See “Image Properties” in the pkg(1) man page for descriptions of the flush-content-cache-on-success and send-uuid properties and information about how to view and modify their values.

**Files** Since pkg(5) images can be located arbitrarily within a larger file system, the token \$IMAGE\_ROOT is used to distinguish relative paths. For a typical system installation, \$IMAGE\_ROOT is equivalent to /.

\$IMAGE\_ROOT/var/pkg

Metadata directory for a full or partial image.

\$IMAGE\_ROOT/.org.opensolaris,pkg

Metadata directory for a user image.

Within the metadata of a particular image, certain files and directories can contain information useful during repair and recovery. The token \$IMAGE\_META is used to refer to the top-level directory that contains the metadata. \$IMAGE\_META is typically one of the two paths given above.

\$IMAGE\_META/lost+found

Location of conflicting directories and files moved during a package operation.

\$IMAGE\_META/publisher

Contains a directory for each publisher. Each directory stores publisher-specific metadata.

Other paths within the \$IMAGE\_META directory hierarchy are Private, and are subject to change.

**Attributes** See attributes(5) for descriptions of the following attributes:

---

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	package/pkg
Interface Stability	Uncommitted

**See Also** [pkg\(1\)](#), [pkgsend\(1\)](#), [pkg.depotd\(1M\)](#), [pkg.sysrepo\(1M\)](#), [svcs\(1\)](#), [svcadm\(1M\)](#)

*[Adding and Updating Oracle Solaris 11.1 Software Packages](#)*

*[Copying and Creating Oracle Solaris 11.1 Package Repositories](#)*

*[Packaging and Delivering Software With the Image Packaging System in Oracle Solaris 11.1](#)*

<http://hub.opensolaris.org/bin/view/Project+pkg/>

**Name** privileges – process privilege model

**Description** Oracle Solaris software implements a set of privileges that provide fine-grained control over the actions of processes. The possession of a certain privilege allows a process to perform a specific set of restricted operations.

The change to a primarily privilege-based security model in the Oracle Solaris operating system gives developers an opportunity to restrict processes to those privileged operations actually needed instead of all (super-user) or no privileges (non-zero UIDs). Additionally, a set of previously unrestricted operations now requires a privilege; these privileges are dubbed the “basic” privileges and are by default given to all processes.

Taken together, all defined privileges with the exception of the “basic” privileges compose the set of privileges that are traditionally associated with the root user. The “basic” privileges are “privileges” unprivileged processes were accustomed to having.

The defined privileges are:

**PRIV\_CONTRACT\_EVENT**

Allow a process to request reliable delivery of events to an event endpoint.

Allow a process to include events in the critical event set term of a template which could be generated in volume by the user.

**PRIV\_CONTRACT\_IDENTITY**

Allows a process to set the service FMRI value of a process contract template.

**PRIV\_CONTRACT\_OBSERVER**

Allow a process to observe contract events generated by contracts created and owned by users other than the process's effective user ID.

Allow a process to open contract event endpoints belonging to contracts created and owned by users other than the process's effective user ID.

**PRIV\_CPC\_CPU**

Allow a process to access per-CPU hardware performance counters.

**PRIV\_DTRACE\_KERNEL**

Allow DTrace kernel-level tracing.

**PRIV\_DTRACE\_PROC**

Allow DTrace process-level tracing. Allow process-level tracing probes to be placed and enabled in processes to which the user has permissions.

**PRIV\_DTRACE\_USER**

Allow DTrace user-level tracing. Allow use of the syscall and profile DTrace providers to examine processes to which the user has permissions.

**PRIV\_FILE\_CHOWN**

Allow a process to change a file's owner user ID. Allow a process to change a file's group ID to one other than the process's effective group ID or one of the process's supplemental group IDs.

**PRIV\_FILE\_CHOWN\_SELF**

Allow a process to give away its files. A process with this privilege runs as if `{_POSIX_CHOWN_RESTRICTED}` is not in effect.

**PRIV\_FILE\_DAC\_EXECUTE**

Allow a process to execute an executable file whose permission bits or ACL would otherwise disallow the process execute permission.

**PRIV\_FILE\_DAC\_READ**

Allow a process to read a file or directory whose permission bits or ACL would otherwise disallow the process read permission.

**PRIV\_FILE\_DAC\_SEARCH**

Allow a process to search a directory whose permission bits or ACL would not otherwise allow the process search permission.

**PRIV\_FILE\_DAC\_WRITE**

Allow a process to write a file or directory whose permission bits or ACL do not allow the process write permission. All privileges are required to write files owned by UID 0 in the absence of an effective UID of 0.

**PRIV\_FILE\_DOWNGRADE\_SL**

Allow a process to set the sensitivity label of a file or directory to a sensitivity label that does not dominate the existing sensitivity label.

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_FILE\_FLAG\_SET**

Allows a process to set `immutable`, `nolink` or `appendonly` file attributes.

**PRIV\_FILE\_LINK\_ANY**

Allow a process to create hardlinks to files owned by a UID different from the process's effective UID.

**PRIV\_FILE\_OWNER**

Allow a process that is not the owner of a file to modify that file's access and modification times. Allow a process that is not the owner of a directory to modify that directory's access and modification times. Allow a process that is not the owner of a file or directory to remove or rename a file or directory whose parent directory has the "save text image after execution" (sticky) bit set. Allow a process that is not the owner of a file to mount a `namefs` upon that file. Allow a process that is not the owner of a file or directory to modify that file's or directory's permission bits or ACL.

**PRIV\_FILE\_READ**

Allow a process to read a file or directory whose permission or ACL allow the process read permission.

**PRIV\_FILE\_SETID**

Allow a process to change the ownership of a file or write to a file without the set-user-ID and set-group-ID bits being cleared. Allow a process to set the set-group-ID bit on a file or directory whose group is not the process's effective group or one of the process's supplemental groups. Allow a process to set the set-user-ID bit on a file with different ownership in the presence of PRIV\_FILE\_OWNER. Additional restrictions apply when creating or modifying a `setuid 0` file.

**PRIV\_FILE\_UPGRADE\_SL**

Allow a process to set the sensitivity label of a file or directory to a sensitivity label that dominates the existing sensitivity label.

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_FILE\_WRITE**

Allow a process to write a file or directory whose permission or ACL allow the process write permission. +

**PRIV\_GRAPHICS\_ACCESS**

Allow a process to make privileged ioctls to graphics devices. Typically only an xserver process needs to have this privilege. A process with this privilege is also allowed to perform privileged graphics device mappings.

**PRIV\_GRAPHICS\_MAP**

Allow a process to perform privileged mappings through a graphics device.

**PRIV\_IPC\_DAC\_READ**

Allow a process to read a System V IPC Message Queue, Semaphore Set, or Shared Memory Segment whose permission bits would not otherwise allow the process read permission.

**PRIV\_IPC\_DAC\_WRITE**

Allow a process to write a System V IPC Message Queue, Semaphore Set, or Shared Memory Segment whose permission bits would not otherwise allow the process write permission.

**PRIV\_IPC\_OWNER**

Allow a process that is not the owner of a System V IPC Message Queue, Semaphore Set, or Shared Memory Segment to remove, change ownership of, or change permission bits of the Message Queue, Semaphore Set, or Shared Memory Segment.

**PRIV\_NET\_ACCESS**

Allow a process to open a TCP, UDP, SDP or SCTP network endpoint.

**PRIV\_NET\_BINDMLP**

Allow a process to bind to a port that is configured as a multi-level port (MLP) for the process's zone. This privilege applies to both shared address and zone-specific address MLPs. See `tnzonecfg(4)` from the Trusted Extensions manual pages for information on configuring MLP ports.

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_NET\_ICMPACCESS**

Allow a process to send and receive ICMP packets.

**PRIV\_NET\_MAC\_AWARE**

Allow a process to set the `NET_MAC_AWARE` process flag by using `setpflags(2)`. This privilege also allows a process to set the `SO_MAC_EXEMPT` socket option by using `setsockopt(3SOCKET)`. The `NET_MAC_AWARE` process flag and the `SO_MAC_EXEMPT` socket option both allow a local process to communicate with an unlabeled peer if the local process's label dominates the peer's default label, or if the local process runs in the global zone.

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_NET\_OBSERVABILITY**

Allow a process to open a device for just receiving network traffic, sending traffic is disallowed.

**PRIV\_NET\_PRIVADDR**

Allow a process to bind to a privileged port number. The privilege port numbers are 1-1023 (the traditional UNIX privileged ports) as well as those ports marked as “`udp/tcp_extra_priv_ports`” with the exception of the ports reserved for use by NFS and SMB.

**PRIV\_NET\_RAWACCESS**

Allow a process to have direct access to the network layer.

**PRIV\_PROC\_AUDIT**

Allow a process to generate audit records. Allow a process to get its own audit pre-selection information.

**PRIV\_PROC\_CHROOT**

Allow a process to change its root directory.

**PRIV\_PROC\_CLOCK\_HIGHRES**

Allow a process to use high resolution timers.

**PRIV\_PROC\_EXEC**

Allow a process to call `exec(2)`.

**PRIV\_PROC\_FORK**

Allow a process to call `fork(2)`, `fork1(2)`, or `vfork(2)`.

**PRIV\_PROC\_INFO**

Allow a process to examine the status of processes other than those to which it can send signals. Processes that cannot be examined cannot be seen in `/proc` and appear not to exist.

**PRIV\_PROC\_LOCK\_MEMORY**

Allow a process to lock pages in physical memory.

**PRIV\_PROC\_OWNER**

Allow a process to send signals to other processes and inspect and modify the process state in other processes, regardless of ownership. When modifying another process, additional restrictions apply: the effective privilege set of the attaching process must be a superset of the target process's effective, permitted, and inheritable sets; the limit set must be a superset of the target's limit set; if the target process has any UID set to 0 all privilege must be asserted unless the effective UID is 0. Allow a process to bind arbitrary processes to CPUs.

**PRIV\_PROC\_PRIOCNTRL**

Allow a process to elevate its priority above its current level. Allow a process to change its scheduling class to any scheduling class, including the RT class.

**PRIV\_PROC\_SESSION**

Allow a process to send signals or trace processes outside its session.

**PRIV\_PROC\_SETID**

Allow a process to set its UIDs at will, assuming UID 0 requires all privileges to be asserted.

**PRIV\_PROC\_TASKID**

Allow a process to assign a new task ID to the calling process.

**PRIV\_PROC\_ZONE**

Allow a process to trace or send signals to processes in other zones. See [zones\(5\)](#).

**PRIV\_SYS\_ACCT**

Allow a process to enable and disable and manage accounting through [acct\(2\)](#).

**PRIV\_SYS\_ADMIN**

Allow a process to perform system administration tasks such as setting node and domain name and specifying [coreadm\(1M\)](#) and [nscd\(1M\)](#) settings

**PRIV\_SYS\_AUDIT**

Allow a process to start the (kernel) audit daemon. Allow a process to view and set audit state (audit user ID, audit terminal ID, audit sessions ID, audit pre-selection mask). Allow a process to turn off and on auditing. Allow a process to configure the audit parameters (cache and queue sizes, event to class mappings, and policy options).

**PRIV\_SYS\_CONFIG**

Allow a process to perform various system configuration tasks. Allow filesystem-specific administrative procedures, such as filesystem configuration ioctls, quota calls, creation and deletion of snapshots, and manipulating the PCFS bootsector.

**PRIV\_SYS\_DEVICES**

Allow a process to create device special files. Allow a process to successfully call a kernel module that calls the kernel `drv_priv(9F)` function to check for allowed access. Allow a process to open the real console device directly. Allow a process to open devices that have been exclusively opened.

**PRIV\_SYS\_DL\_CONFIG**

Allow a process to configure a system's datalink interfaces.

**PRIV\_SYS\_IP\_CONFIG**

Allow a process to configure a system's IP interfaces and routes. Allow a process to configure TCP/IP parameters. Allow a process to pop anchored STREAMS modules with matching zoneid.

**PRIV\_SYS\_IPC\_CONFIG**

Allow a process to increase the size of a System V IPC Message Queue buffer.

**PRIV\_SYS\_LINKDIR**

Allow a process to unlink and link directories.

**PRIV\_SYS\_MOUNT**

Allow a process to mount and unmount filesystems that would otherwise be restricted (that is, most filesystems except `namefs`). Allow a process to add and remove swap devices.

**PRIV\_SYS\_NET\_CONFIG**

Allow a process to do all that `PRIV_SYS_IP_CONFIG`, `PRIV_SYS_DL_CONFIG`, and `PRIV_SYS_PPP_CONFIG` allow, plus the following: use the `rpcmod` STREAMS module and insert/remove STREAMS modules on locations other than the top of the module stack.

**PRIV\_SYS\_NFS**

Allow a process to provide NFS service: start NFS kernel threads, perform NFS locking operations, bind to NFS reserved ports: ports 2049 (`nfs`) and port 4045 (`lockd`).

**PRIV\_SYS\_PPP\_CONFIG**

Allow a process to create, configure, and destroy PPP instances with `pppd(1M)` `pppd(1M)` and control PPPoE plumbing with `spptun(1M)``spptun(1M)`. This privilege is granted by default to exclusive IP stack instance zones.

**PRIV\_SYS\_RES\_BIND**

Allow a process to bind processes to processor sets.

**PRIV\_SYS\_RES\_CONFIG**

Allow a process to bind processes to processor sets, as `PRIV_SYS_RES_BIND`, in addition to the following outlined in this paragraph. Allow a process to create and delete processor sets, assign CPUs to processor sets and override the `PSET_NOESCAPE` property. Allow a process to change the operational status of CPUs in the system using `p_online(2)`. Allow a process to configure filesystem quotas. Allow a process to configure resource pools and bind processes to pools.

**PRIV\_SYS\_RESOURCE**

Allow a process to exceed the resource limits imposed on it by `setrlimit(2)` and `setrctl(2)`.

**PRIV\_SYS\_SHARE**

Allow a process to share and unshare filesystems.

**PRIV\_SYS\_SMB**

Allow a process to provide NetBIOS or SMB services: start SMB kernel threads or bind to NetBIOS or SMB reserved ports: ports 137, 138, 139 (NetBIOS) and 445 (SMB).

**PRIV\_SYS\_SUSER\_COMPAT**

Allow a process to successfully call a third party loadable module that calls the kernel `user()` function to check for allowed access. This privilege exists only for third party loadable module compatibility and is not used by Oracle Solaris proper.

**PRIV\_SYS\_TIME**

Allow a process to manipulate system time using any of the appropriate system calls: `stime(2)`, `adjtime(2)`, and `ntp_adjtime(2)`.

**PRIV\_SYS\_TRANS\_LABEL**

Allow a process to translate labels that are not dominated by the process's sensitivity label to and from an external string form.

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_VIRT\_MANAGE**

Allows a process to manage virtualized environments.

**PRIV\_WIN\_COLORMAP**

Allow a process to override colormap restrictions.

Allow a process to install or remove colormaps.

Allow a process to retrieve colormap cell entries allocated by other processes.

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_WIN\_CONFIG**

Allow a process to configure or destroy resources that are permanently retained by the X server.

Allow a process to use `SetScreenSaver` to set the screen saver timeout value

Allow a process to use `ChangeHosts` to modify the display access control list.

Allow a process to use `GrabServer`.

Allow a process to use the `SetCloseDownMode` request that can retain window, pixmap, colormap, property, cursor, font, or graphic context resources.

---

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_WIN\_DAC\_READ**

Allow a process to read from a window resource that it does not own (has a different user ID).

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_WIN\_DAC\_WRITE**

Allow a process to write to or create a window resource that it does not own (has a different user ID). A newly created window property is created with the window's user ID.

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_WIN\_DEVICES**

Allow a process to perform operations on window input devices.

Allow a process to get and set keyboard and pointer controls.

Allow a process to modify pointer button and key mappings.

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_WIN\_DGA**

Allow a process to use the direct graphics access (DGA) X protocol extensions. Direct process access to the frame buffer is still required. Thus the process must have MAC and DAC privileges that allow access to the frame buffer, or the frame buffer must be allocated to the process.

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_WIN\_DOWNGRADE\_SL**

Allow a process to set the sensitivity label of a window resource to a sensitivity label that does not dominate the existing sensitivity label.

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_WIN\_FONTPATH**

Allow a process to set a font path.

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_WIN\_MAC\_READ**

Allow a process to read from a window resource whose sensitivity label is not equal to the process sensitivity label.

This privilege is interpreted only if the system is configured with Trusted Extensions.

**PRIV\_WIN\_MAC\_WRITE**

Allow a process to create a window resource whose sensitivity label is not equal to the process sensitivity label. A newly created window property is created with the window's sensitivity label.

This privilege is interpreted only if the system is configured with Trusted Extensions.

#### PRIV\_WIN\_SELECTION

Allow a process to request inter-window data moves without the intervention of the selection confirmer.

This privilege is interpreted only if the system is configured with Trusted Extensions.

#### PRIV\_WIN\_UPGRADE\_SL

Allow a process to set the sensitivity label of a window resource to a sensitivity label that dominates the existing sensitivity label.

This privilege is interpreted only if the system is configured with Trusted Extensions.

Of the privileges listed above, the privileges `PRIV_FILE_LINK_ANY`, `PRIV_FILE_READ`, `PRIV_FILE_WRITE`, `PRIV_PROC_INFO`, `PRIV_PROC_SESSION`, `PRIV_NET_ACCESS`, `PRIV_PROC_FORK`, and `PRIV_PROC_EXEC` are considered “basic” privileges. These are privileges that used to be always available to unprivileged processes. By default, processes still have the basic privileges.

The privileges `PRIV_PROC_SETID`, `PRIV_PROC_AUDIT`, and `PRIV_SYS_RESOURCE` must be present in the Limit set (see below) of a process in order for `setuid root` execs to be successful; that is, get an effective UID of 0 and additional privileges.

The privilege implementation in Oracle Solaris extends the process credential with four privilege sets:

I, the inheritable set	The privileges inherited on exec.
P, the permitted set	The maximum set of privileges for the process.
E, the effective set	The privileges currently in effect.
L, the limit set	The upper bound of the privileges a process and its offspring can obtain. Changes to L take effect on the next exec.

The sets I, P and E are typically identical to the basic set of privileges for unprivileged processes. The limit set is typically the full set of privileges.

Each process has a Privilege Awareness State (PAS) that can take the value PA (privilege-aware) and NPA (not-PA). PAS is a transitional mechanism that allows a choice between full compatibility with the old superuser model and completely ignoring the effective UID.

To facilitate the discussion, we introduce the notion of “observed effective set” (oE) and “observed permitted set” (oP) and the implementation sets iE and iP.

A process becomes privilege-aware either by manipulating the effective, permitted, or limit privilege sets through `setppriv(2)` or by using `setpflags(2)`. In all cases, oE and oP are

invariant in the process of becoming privilege-aware. In the process of becoming privilege-aware, the following assignments take place:

```
iE = oE
iP = oP
```

When a process is privilege-aware, oE and oP are invariant under UID changes. When a process is not privilege-aware, oE and oP are observed as follows:

```
oE = euid == 0 ? L : iE
oP = (euid == 0 || ruid == 0 || suid == 0) ? L : iP
```

When a non-privilege-aware process has an effective UID of 0, it can exercise the privileges contained in its limit set, the upper bound of its privileges. If a non-privilege-aware process has any of the UIDs 0, it appears to be capable of potentially exercising all privileges in L.

It is possible for a process to return to the non-privilege aware state using `setpflags()`. The kernel always attempts this on `exec(2)`. This operation is permitted only if the following conditions are met:

- If any of the UIDs is equal to 0, P must be equal to L.
- If the effective UID is equal to 0, E must be equal to L.

When a process gives up privilege awareness, the following assignments take place:

```
if (euid == 0) iE = L & I
if (any uid == 0) iP = L & I
```

The privileges obtained when not having a UID of 0 are the inheritable set of the process restricted by the limit set.

Only privileges in the process's (observed) effective privilege set allow the process to perform restricted operations. A process can use any of the privilege manipulation functions to add or remove privileges from the privilege sets. Privileges can be removed always. Only privileges found in the permitted set can be added to the effective and inheritable set. The limit set cannot grow. The inheritable set can be larger than the permitted set.

When a process performs an `exec(2)`, the kernel first tries to relinquish privilege awareness before making the following privilege set modifications:

```
E' = P' = I' = L & I
L is unchanged
```

If a process has not manipulated its privileges, the privilege sets effectively remain the same, as E, P and I are already identical.

The limit set is enforced at `exec` time.

To run a non-privilege-aware application in a backward-compatible manner, a privilege-aware application should start the non-privilege-aware application with `I=basic`.

For most privileges, absence of the privilege simply results in a failure. In some instances, the absence of a privilege can cause system calls to behave differently. In other instances, the removal of a privilege can force a `setuid` root application to seriously malfunction. Privileges of this type are considered “unsafe”. When a process is lacking any of the unsafe privileges from its limit set, the system does not honor the `setuid` bit of `setuid` root applications. An `exec` of a `setuid` root application would proceed without the change in effective user ID or increase in privilege. The following unsafe privileges have been identified: `PRIV_PROC_SETID`, `PRIV_SYS_RESOURCE` and `PRIV_PROC_AUDIT`.

**Privilege Escalation** In certain circumstances, a single privilege could lead to a process gaining one or more additional privileges that were not explicitly granted to that process. To prevent such an escalation of privileges, the security policy requires explicit permission for those additional privileges.

Common examples of escalation are those mechanisms that allow modification of system resources through “raw” interfaces; for example, changing kernel data structures through `/dev/kmem` or changing files through `/dev/dsk/*`. Escalation also occurs when a process controls processes with more privileges than the controlling process. A special case of this is manipulating or creating objects owned by UID 0 or trying to obtain UID 0 using `setuid(2)`. The special treatment of UID 0 is needed because the UID 0 owns all system configuration files and ordinary file protection mechanisms allow processes with UID 0 to modify the system configuration. With appropriate file modifications, a given process running with an effective UID of 0 can gain all privileges.

In situations where a process might obtain UID 0, the security policy requires additional privileges, up to the full set of privileges. Such restrictions could be relaxed or removed at such time as additional mechanisms for protection of system files became available. There are no such mechanisms in the current Oracle Solaris release.

The use of UID 0 processes should be limited as much as possible. They should be replaced with programs running under a different UID but with exactly the privileges they need.

Daemons that never need to `exec` subprocesses should remove the `PRIV_PROC_EXEC` privilege from their permitted and limit sets.

**Assigned Privileges and Safeguards** When privileges are assigned to a user, the system administrator could give that user more powers than intended. The administrator should consider whether safeguards are needed. For example, if the `PRIV_PROC_LOCK_MEMORY` privilege is given to a user, the administrator should consider setting the `project.max-locked-memory` resource control as well, to prevent that user from locking all memory.

Extended Policy When privileges are listed in configuration files or on the command line, it is sometimes possible to use the Extended Policy syntax. An Extended Policy is a privilege set enclosed in braces followed by a colon and an object. An extended policy adds the ability to use that list of privileges on the specified object. Currently we support extended policies for network ports, UIDs and file objects.

For example,

```
{file_dac_read}:/var/core/*
```

Allows the using of the privilege `file_dac_read` when accessing files under `/var/core`.

```
{net_privaddr}:80/tcp, {net_privaddr}:443/tcp
```

Allows a process to bind a network endpoint to TCP port 80 and 443.

```
{proc_setid}:80-100
```

Allows a process to change UID to UID 80 through 100, inclusive.

```
{proc_setid}:casper
```

Allows a process to change the UID to the UID of the username “casper”.

A privilege set listed in an extended policy will be removed from the inheritable set and consequently from the permitted and the effective set when the Extended Policy is installed unless the privilege set in the policy includes all privileges available in the zone. For example, when installing an Extended Policy of the form `{zone}:/etc/shadow`, no privileges are dropped; if, on the other hand, the Extended Policy has the form `{file_dac_read}:/etc/shadow`, the `PRIV_FILE_DAC_READ` privilege is removed.

The Extended Policy is in effect only when a privilege is missing from the effective set.

While it is possible to specify an Extended Policy such as `{all}:/some/file`, the system will still restrict some applications such as adding a `setuid` bit.

All privileges listed in an Extended Policy need to be effective in the process when that process installs that policy. For example, when `ppriv(1)` is used to install a policy, it needs to have all privileges listed in its effective set. No such restrictions apply to the process that is the object of `ppriv`. However, its Limit set overrides any privileges in the Extended Policy.

In some contexts it is required to escape part of Extended Policy syntax. For example, in `exec_attr(4)`, the colon (`:`) needs to be escaped using a backslash (`\`). Some characters in filenames may also needed to be escaped using a backslash, depending on the context.

The Extended Policy is evaluated at every layer in the filesystem; in the case of `lofs(7FS)` file systems, the specified policy needs to takes this into account: the policy needs to specify both the `lofs` filesystem and the underlying filesystem.

The following list contains types of objects and relevant privileges.

Object	Syntax	Privilege
Username	name	proc_setid
Uid	uid	proc_setid
Range of uids	uid1-uid2	proc_setid
Network port <sup>1</sup>	port/udp, port/tcp, port/sctp, port/*	net_privaddr
Range of ports	port1-port2/<proto>	net_privaddr
Filename	pathname	file privileges, proc_exec
Wildcard <sup>2</sup>	pathname*	file privileges, proc_exec

<sup>1</sup> numeric as defined in [services\(4\)](#)

<sup>2</sup> matches all filenames starting with the specified pathname

**Privilege Debugging** When a system call fails with a permission error, it is not always immediately obvious what caused the problem. To debug such a problem, you can use a tool called *privilege debugging*. When privilege debugging is enabled for a process, the kernel reports missing privileges on the controlling terminal of the process. (Enable debugging for a process with the `-D` option of [ppriv\(1\)](#).) Additionally, the administrator can enable system-wide privilege debugging by setting the [system\(4\)](#) variable `priv_debug` using:

```
set priv_debug = 1
```

On a running system, you can use [mdb\(1\)](#) to change this variable.

**Privilege Administration** Use [usermod\(1M\)](#) or [rolemod\(1M\)](#) to assign privileges to assign or modify privileges to, respectively, a user or a role. Use [ppriv\(1\)](#) to enumerate the privileges supported on a system and [truss\(1\)](#) to determine which privileges a program requires.

**See Also** [mdb\(1\)](#), [ppriv\(1\)](#), [add\\_drv\(1M\)](#), [ifconfig\(1M\)](#), [lockd\(1M\)](#), [nfsd\(1M\)](#), [pppd\(1M\)](#), [rem\\_drv\(1M\)](#), [smbd\(1M\)](#), [spptun\(1M\)](#), [update\\_drv\(1M\)](#), [Intro\(2\)](#), [access\(2\)](#), [acct\(2\)](#), [acl\(2\)](#), [adjtime\(2\)](#), [chmod\(2\)](#), [chown\(2\)](#), [chroot\(2\)](#), [creat\(2\)](#), [exec\(2\)](#), [fcntl\(2\)](#), [fork\(2\)](#), [fpathconf\(2\)](#), [getacct\(2\)](#), [getpflags\(2\)](#), [getppriv\(2\)](#), [getsid\(2\)](#), [kill\(2\)](#), [link\(2\)](#), [memcntl\(2\)](#), [mknod\(2\)](#), [mount\(2\)](#), [msgctl\(2\)](#), [nice\(2\)](#), [ntp\\_adjtime\(2\)](#), [open\(2\)](#), [p\\_online\(2\)](#), [prioctl\(2\)](#), [prioctlset\(2\)](#), [processor\\_bind\(2\)](#), [pset\\_bind\(2\)](#), [pset\\_create\(2\)](#), [readlink\(2\)](#), [resolvepath\(2\)](#), [rmdir\(2\)](#), [semctl\(2\)](#), [setegid\(2\)](#), [seteuid\(2\)](#), [setgid\(2\)](#), [setgroups\(2\)](#), [setpflags\(2\)](#), [setppriv\(2\)](#), [setrctl\(2\)](#), [setregid\(2\)](#), [setreuid\(2\)](#), [setrlimit\(2\)](#), [settaskid\(2\)](#), [setuid\(2\)](#), [shmctl\(2\)](#), [shmget\(2\)](#), [shmop\(2\)](#), [sigsend\(2\)](#), [stat\(2\)](#), [statvfs\(2\)](#), [stime\(2\)](#), [swapctl\(2\)](#), [sysinfo\(2\)](#), [uadmin\(2\)](#), [ulimit\(2\)](#), [umount\(2\)](#), [unlink\(2\)](#), [utime\(2\)](#), [utimes\(2\)](#), [bind\(3SOCKET\)](#), [door\\_ucred\(3C\)](#), [priv\\_addset\(3C\)](#), [priv\\_set\(3C\)](#), [priv\\_getbyname\(3C\)](#), [priv\\_getbynum\(3C\)](#), [priv\\_set\\_to\\_str\(3C\)](#), [priv\\_str\\_to\\_set\(3C\)](#), [socket\(3SOCKET\)](#), [t\\_bind\(3NSL\)](#), [timer\\_create\(3C\)](#),

ucred\_get(3C), exec\_attr(4), proc(4), services(4), system(4), user\_attr(4),  
lofs(7FS) ddi\_cred(9F), drv\_priv(9F), priv\_getbyname(9F), priv\_policy(9F),  
priv\_policy\_choice(9F), priv\_policy\_only(9F)

*Oracle Solaris 11.1 Administration: Security Services*

**Notes** Removal of any of the basic privileges from a process leaves it in a non-standards compliant state, may cause unexpected application failures, and should only be performed with full knowledge of the potential side effects.

**Name** prof – profile within a function

**Synopsis**

```
#define MARK
#include <prof.h>

void MARK(name);
```

**Description** MARK introduces a mark called *name* 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.

*name* may be any combination of letters, numbers, or underscores. Each *name* in a single compilation must be unique, but may be the same as any ordinary program symbol.

For marks to be effective, the symbol MARK must be defined before the header `prof.h` is included, either by a preprocessor directive as in the synopsis, or by a command line argument:

```
cc -p -DMARK work.c
```

If MARK is not defined, the `MARK(name)` statements may be left in the source files containing them and are ignored. `prof -g` must be used to get information on all labels.

**Examples** 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.

```
#include <prof.h>
work( )
{
    int i, j;
    . . .
    MARK(loop1);
    for (i = 0; i < 2000; i++) {
        . . .
    }
    MARK(loop2);
    for (j = 0; j < 2000; j++) {
        . . .
    }
}
```

**See Also** [profil\(2\)](#), [monitor\(3C\)](#)

**Name** rbac, RBAC – role-based access control

**Description** Role-based access control allows system administrators to delegate the administrative control of parts of the system to users. Users can be given the ability to run commands with additional privileges in two ways:

- by assigning a profile directly to the user, in which case no additional authentication is required
- by creating a role and assigning the profiles to the role. It can also be used to build restrictive environments for users by removing their ability to run commands they would normally be allowed to run.

**Profiles** Profiles are named collections of commands and authorizations that are run with additional privilege and/or a specific real and effective UID and GID. For example, most of the printer system can be managed by having the `lp` commands run with the UID or `lp`. Some commands need privileges as defined in [privileges\(5\)](#) to run. For example, the “Process Management” profile allows a user to run the `kill` command with the `proc_owner` privilege so that it can send signals to processes it does not own.

See [exec\\_attr\(4\)](#) and [prof\\_attr\(4\)](#) for information about how the administrator can extend the system-provided profiles and create their own. Profile configuration can be stored in any of the currently supported name services (files, NIS, LDAP).

Profiles can also be used with the Service Management Facility (SMF) to control the privileges and UID/GID with which a service runs. See [smf\\_security\(5\)](#) for more information.

**Roles** A role is a special shared account that cannot directly login to the system that can only be accessed by authorized users with the `su(1M)` command or over the network with `ssh(1)` when using host-based authentication or GSS-API authentication. It can not login with `rlogin(1)`, `telnet(1)`, or `gdm`.

A role has a UID, a password, and a home directory just like a normal user. Authentication to the role can be either with the user's own password or with the per-role password (the `roleauth` keyword in [user\\_attr\(4\)](#) controls that behavior on a per-role basis). Usually a role's login shell is one of the profile shells ([pfs\(1\)](#), [pfsk\(1\)](#), [pfcsh\(1\)](#)) that are granted one or more Profiles, allowing the role to always execute commands with privilege.

A role is normally needed only if a shared account environment is required. Usually assigning profiles directly to the user is sufficient.

The root user can be configured to be a role using the `usermod(1M)` command. This ensures that only authorized users can become root even when the root password is more widely known.

```
# usermod -K type=role root
```

Making root a role does not restrict access to single user mode. The system console should be protected using other means, such as setting a security password with [eeprom\(1M\)](#).

**Authorizations** An authorization is a unique string that represents a user's right to perform some operation or class of operations. Authorizations are normally only checked by programs that always run with some privilege, for example the [setuid\(2\)](#) programs such as [cdrw\(1\)](#) or the system [cron\(1M\)](#) daemon.

Authorization definitions are stored in the [auth\\_attr\(4\)](#) database. For programming authorization checks, only the authorization name is significant.

Some typical values in an `auth_attr` database are as follows:

```
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
```

Authorization name strings ending with the `grant` suffix are special authorizations that give a user the ability to delegate authorizations with the same prefix and functional area to other users.

All authorization names starting with `solaris` are reserved for allocation by the operating system vendor. Developers and administrators may create their own top level namespace; use of a unique identifier such as the company name, DNS domain name, or application name is suggested.

**Authorization Checks** To check authorizations from C code, developers should use the [chkauthattr\(3C\)](#) library function, which verifies whether or not a user has a given authorization.

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
```

Authorizations are also used by the Service Management Facility (SMF) to control which users can change the state of a service or reconfigure a service. See [smf\\_security\(5\)](#) for more information.

Comparison with `sudo(1M)` RBAC in Solaris provides a similar set of functionality to `sudo(1M)` that is often provided with UNIX or UNIX-like systems. It is provided on the Companion CD for Solaris.

One of the most obvious differences between Solaris RBAC and `sudo` is the authentication model. In `sudo`, users reauthenticate as themselves. In Solaris RBAC, either no additional authentication is needed (when profiles are assigned directly to the user) or the user authenticates to a shared account called a role.

Using the `NOPASSWD` functionality in `sudo` is similar to assigning the profile to the user and having the user execute the command using `pfexec(1)`. For example, if `sudoers(4)` allows the user to run `kill(1)` as UID 0 but without authentication (`NOPASSWD`), the user would run:

```
$ sudo kill -HUP 1235
```

In Solaris RBAC, if the user has a normal (that is, no profile) login shell, the user would execute the equivalent operation by being assigned the “Process Management” profile and would use `pfexec` as follows:

```
$ pfexec kill -HUP 1235
```

If the user has a profile shell (such as `pfsh`) as the login shell, then `kill` will always run with the additional privilege without the need of a “prefix”. For example,

```
$ kill -HUP 1235
```

An RBAC role is similar in concept to the `User_Alias` in `sudoers(4)`, except that the role password rather than the user password is required.

Execution profiles `exec_attr(4)` entries) in RBAC are similar to the `Cmnd_Alias` in `sudoers`.

There is currently no equivalent of the `Host_Alias` `sudo(1M)` functionality in Solaris RBAC.

**See Also** [auths\(1\)](#), [ld.so.1\(1\)](#), [pfcsh\(1\)](#), [pfexec\(1\)](#), [pfksh\(1\)](#), [pfsh\(1\)](#), [roles\(1\)](#), [sudo\(1M\)](#), [exec\\_attr\(4\)](#), [prof\\_attr\(4\)](#), [user\\_attr\(4\)](#), [smf\\_security\(5\)](#)

*Oracle Solaris 11.1 Administration: Security Services*

**Name** regex – internationalized basic and extended regular expression matching

**Description** 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 [regex\(5\)](#) manual page in the following ways:

- both Basic and Extended Regular Expressions are supported
- the Internationalization features—character class, equivalence class, and multi-character collation—are supported.

The Basic Regular Expression (BRE) notation and construction rules described in the BASIC REGULAR EXPRESSIONS section apply to most utilities supporting regular expressions. Some utilities, instead, support the Extended Regular Expressions (ERE) described in the EXTENDED REGULAR EXPRESSIONS section; any exceptions for both cases are noted in the descriptions of the specific utilities using regular expressions. Both BREs and EREs are supported by the Regular Expression Matching interfaces [regcomp\(3C\)](#) and [regex\(3C\)](#).

### Basic Regular Expressions

**BREs Matching a Single Character** 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 RE Bracket Expression, below.

**BRE Ordinary Characters** An ordinary character is a BRE that matches itself: any character in the supported character set, except for the BRE special characters listed in BRE Special Characters, below.

The interpretation of an ordinary character preceded by a backslash (\) is undefined, except for:

1. the characters ), (, {, and }
2. the digits 1 to 9 inclusive (see BREs Matching Multiple Characters, below)
3. a character inside a bracket expression.

**BRE Special Characters** A BRE *special character* 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:

- . [ \ The period, left-bracket, and backslash are special except when used in a bracket expression (see RE Bracket Expression, below). An expression containing a [ that is not preceded by a backslash and is not part of a bracket expression produces undefined results.
- \* The asterisk is special except when used:
  - in a bracket expression
  - as the first character of an entire BRE (after an initial ^, if any)

- as the first character of a subexpression (after an initial `^`, if any); see BREs Matching Multiple Characters, below.
  - ^ The circumflex is special when used:
    - as an anchor (see BRE Expression Anchoring, below).
    - as the first character of a bracket expression (see RE Bracket Expression, below).
  - \$ The dollar sign is special when used as an anchor.
- Periods in BREs A period (`.`), when used outside a bracket expression, is a BRE that matches any character in the supported character set except NUL.
- RE Bracket Expression 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.

The following rules and definitions apply to bracket expressions:

1. A *bracket expression* 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:

`. * [ \`

(period, asterisk, left-bracket and backslash, respectively) lose their special meaning within a bracket expression.

The character sequences:

`[. [= [:`

(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 *matching list* 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 *non-matching list* 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 `z`, 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 `abababcccccd`, 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>\&lt;special character&gt;</code>
bracket expression	<code>[]</code>
subexpressions/back-references	<code>\( \) \n</code>

single-character-BRE duplication	* $\{m,n\}$
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.

**ERE Special Characters** An *ERE special character* 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:

- . \ (     The period, left-bracket, backslash, and left-parenthesis are special except when used in a bracket expression (see *RE Bracket Expression*, above). Outside a bracket expression, a left-parenthesis immediately followed by a right-parenthesis produces undefined results.
- )         The right-parenthesis is special when matched with a preceding left-parenthesis, both outside a bracket expression.
- \* + ? {    The asterisk, plus-sign, question-mark, and left-brace are special except when used in a bracket expression (see *RE Bracket Expression*, above). Any of the following uses produce undefined results:
  - 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.
- |         The vertical-line is special except when used in a bracket expression (see *RE Bracket Expression*, 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.
- ^         The circumflex is special when used:
  - as an anchor (see *ERE Expression Anchoring*, below).
  - as the first character of a bracket expression (see *RE Bracket Expression*, above).
- \$         The dollar sign is special when used as an anchor.

**Periods in EREs** A period (.), when used outside a bracket expression, is an ERE that matches any character in the supported character set except NUL.

**ERE Bracket Expression** The rules for ERE Bracket Expressions are the same as for Basic Regular Expressions; see *RE Bracket Expression*, above).

**EREs Matching Multiple Characters** The following rules will be used to construct EREs matching multiple characters from EREs matching a single character:

1. A *concatenation of EREs* 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 `cd` and the ERE `( cd )` are matched by the third and fourth character of the string `abcdefabcdef`.

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 `cabbbbcdebbbbbcdbc`. 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 `abababcccccd` 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>\&lt;special character&gt;</code>
bracket expression	<code>[]</code>

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 `abade` or `abcde`, 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**

```
#define INIT declarations
#define GETC(void) getc code
#define PEEKC(void) peekc code
#define UNGETC(void) ungetc code
#define RETURN(ptr) return code
#define ERROR(val) error code

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);
```

**Description** 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:

- only Basic Regular Expressions are supported
- the Internationalization features—character class, equivalence class, and multi-character collation—are not supported.

The functions `step()`, `advance()`, and `compile()` are general purpose regular expression matching routines to be used in programs that perform regular expression matching. These functions are defined by the `<regex.h>` header.

The functions `step()` and `advance()` do pattern matching given a character string and a compiled regular expression as input.

The function `compile()` takes as input a regular expression as defined below and produces a compiled expression that can be used with `step()` or `advance()`.

**Basic Regular Expressions** 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.

The following *one-character REs* match a *single* character:

- 1.1 An ordinary character (*not* one of those discussed in 1.2 below) is a one-character RE that matches itself.
- 1.2 A backslash ( `\` ) 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 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 <regexp.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 <regexp.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().

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 `loc2` 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 Using Regular Expression Macros and Calls

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\(5\)](#)

**Name** resource\_controls – resource controls available through project database

**Description** The resource controls facility is configured through the project database. See [project\(4\)](#). You can set and modify resource controls through the following utilities:

- [prctl\(1\)](#)
- [projadd\(1M\)](#)
- [projmod\(1M\)](#)
- [rctladm\(1M\)](#)

In a program, you use [setrctl\(2\)](#) to set resource control values.

In addition to the preceding resource controls, there are resource pools, accessible through the [pooladm\(1M\)](#) and [poolcfg\(1M\)](#) utilities. In a program, resource pools can be manipulated through the [libpool\(3LIB\)](#) library.

The following are the resource controls are available:

`process.max-address-space`

Maximum amount of address space, as summed over segment sizes, that is available to this process, expressed as a number of bytes.

`process.max-core-size`

Maximum size of a core file created by this process, expressed as a number of bytes.

`process.max-cpu-time`

Maximum CPU time that is available to this process, expressed as a number of seconds.

`process.max-data-size`

Maximum heap memory available to this process, expressed as a number of bytes.

`process.max-file-descriptor`

Maximum file descriptor index available to this process, expressed as an integer.

`process.max-file-size`

Maximum file offset available for writing by this process, expressed as a number of bytes.

`process.max-itimers`

Maximum allowable number of interval timers, expressed as an integer.

`process.max-msg-messages`

Maximum number of messages on a message queue (value copied from the resource control at `msgget()` time), expressed as an integer.

`process.max-msg-qbytes`

Maximum number of bytes of messages on a message queue (value copied from the resource control at `msgget()` time), expressed as a number of bytes.

`process.max-port-events`

Maximum allowable number of events per event port, expressed as an integer.

`process.max-sem-nsems`

Maximum number of semaphores allowed per semaphore set, expressed as an integer.

`process.max-sem-ops`

Maximum number of semaphore operations allowed per `semop` call (value copied from the resource control at `semget()` time). Expressed as an integer, specifying the number of operations.

`process.max-sigqueue-size`

Maximum number of outstanding queued signals that a process can have.

`process.max-stack-size`

Maximum stack memory segment available to this process, expressed as a number of bytes.

`project.cpu-caps`

Maximum amount of CPU resources that a project can use. The unit used is the percentage of a single CPU that can be used by all user threads in a project. Expressed as an integer. The cap does not apply to threads running in real-time scheduling class. This resource control does not support the `syslog` action.

`project.cpu-shares`

Number of CPU shares granted to a project for use with the fair share scheduler (see [FSS\(7\)](#)). The unit used is the number of shares (an integer). This resource control does not support the `syslog` action.

`project.max-contracts`

Maximum number of contracts allowed in a project, expressed as an integer.

`project.max-crypto-memory`

Maximum amount of kernel memory that can be used for crypto operations. Allocations in the kernel for buffers and session-related structures are charged against this resource control.

`project.max-locked-memory`

Total amount of physical memory locked by device drivers and user processes (including D/ISM), expressed as a number of bytes.

`project.max-lwps`

Maximum number of LWPs simultaneously available to a project, expressed as an integer.

`project.max-msg-ids`

Maximum number of message queue IDs allowed for a project, expressed as an integer.

`project.max-port-ids`

Maximum allowable number of event ports, expressed as an integer.

`project.max-processes`

Maximum number of processes simultaneously available to a project, expressed as an integer.

`project.max-sem-ids`

Maximum number of semaphore IDs allowed for a project, expressed as an integer.

`project.max-shm-ids`

Maximum number of shared memory IDs allowed for a project, expressed as an integer.

`project.max-shm-memory`

Total amount of shared memory allowed for a project, expressed as a number of bytes.

`project.max-tasks`

Maximum number of tasks allowable in a project, expressed as an integer.

`project.pool`

Binds a specified resource pool with a project.

`rcap.max-rss`

The total amount of physical memory, in bytes, that is available to processes in a project.

`task.max-cpu-time`

Maximum CPU time that is available to this task's processes, expressed as a number of seconds.

`task.max-lwps`

Maximum number of LWPs simultaneously available to this task's processes, expressed as an integer.

`task.max-processes`

Maximum number of processes simultaneously available to a task, expressed as an integer.

The following zone-wide resource controls are available:

`zone.cpu-cap`

Sets a limit on the amount of CPU time that can be used by a zone. The unit used is the percentage of a single CPU that can be used by all user threads in a zone. Expressed as an integer. When projects within the capped zone have their own caps, the minimum value takes precedence. This resource control does not support the `syslog` action.

`zone.cpu-shares`

Sets a limit on the number of fair share scheduler (FSS) CPU shares for a zone. CPU shares are first allocated to the zone, and then further subdivided among projects within the zone as specified in the `project.cpu-shares` entries. Expressed as an integer. This resource control does not support the `syslog` action.

`zone.max-locked-memory`

Total amount of physical locked memory available to a zone.

`zone.max-lofi`

Maximum number of `lofi(7D)` devices available to a zone.

`zone.max-lwps`

Enhances resource isolation by preventing too many LWPs in one zone from affecting other zones. A zone's total LWPs can be further subdivided among projects within the zone within the zone by using `project.max-lwps` entries. Expressed as an integer.

`zone.max-msg-ids`

Maximum number of message queue IDs allowed for a zone, expressed as an integer.

`zone.max-processes`

Maximum number of processes simultaneously available to a zone, expressed as an integer.

`zone.max-sem-ids`

Maximum number of semaphore IDs allowed for a zone, expressed as an integer.

`zone.max-shm-ids`

Maximum number of shared memory IDs allowed for a zone, expressed as an integer.

`zone.max-shm-memory`

Total amount of shared memory allowed for a zone, expressed as a number of bytes.

`zone.max-swap`

Total amount of swap that can be consumed by user process address space mappings and `tmpfs` mounts for this zone.

See [zones\(5\)](#).

Units Used in Resource Controls Resource controls can be expressed as in units of size (bytes), time (seconds), or as a count (integer). These units use the strings specified below.

Category	Res Ctrl Type String	Modifier	Scale
-----	-----	-----	-----
Size	bytes	B KB MB GB TB PB EB	1 2 <sup>10</sup> 2 <sup>20</sup> 2 <sup>30</sup> 2 <sup>40</sup> 2 <sup>50</sup> 2 <sup>60</sup>
Time	seconds	s Ks Ms Gs Ts Ps Es	1 10 <sup>3</sup> 10 <sup>6</sup> 10 <sup>9</sup> 10 <sup>12</sup> 10 <sup>15</sup> 10 <sup>18</sup>
Count	integer	none K	1 10 <sup>3</sup>

M	10 <sup>6</sup>
G	10 <sup>9</sup>
T	10 <sup>12</sup>
P	10 <sup>15</sup>
Es	10 <sup>18</sup>

Scaled values can be used with resource controls. The following example shows a scaled threshold value:

```
task.max-lwps=(priv,1K,deny)
```

In the project file, the value 1K is expanded to 1000:

```
task.max-lwps=(priv,1000,deny)
```

A second example uses a larger scaled value:

```
process.max-file-size=(priv,5G,deny)
```

In the project file, the value 5G is expanded to 5368709120:

```
process.max-file-size=(priv,5368709120,deny)
```

The preceding examples use the scaling factors specified in the table above.

Note that unit modifiers (for example, 5G) are accepted by the `prctl(1)`, `projadd(1M)`, and `projmod(1M)` commands. You cannot use unit modifiers in the project database itself.

#### Resource Control Values and Privilege Levels

A threshold value on a resource control constitutes a point at which local actions can be triggered or global actions, such as logging, can occur.

Each threshold value on a resource control must be associated with a privilege level. The privilege level must be one of the following three types:

##### basic

Can be modified by the owner of the calling process.

##### privileged

Can be modified by the current process (requiring `sys_resource` privilege) or by `prctl(1)` (requiring `proc_owner` privilege).

##### system

Fixed for the duration of the operating system instance.

A resource control is guaranteed to have one `system` value, which is defined by the system, or resource provider. The `system` value represents how much of the resource the current implementation of the operating system is capable of providing.

Any number of privileged values can be defined, and only one basic value is allowed. Operations that are performed without specifying a privilege value are assigned a basic privilege by default.

The privilege level for a resource control value is defined in the `privilege` field of the resource control block as `RCTL_BASIC`, `RCTL_PRIVILEGED`, or `RCTL_SYSTEM`. See [setrctl\(2\)](#) for more information. You can use the `prctl` command to modify values that are associated with basic and privileged levels.

In specifying the privilege level of `privileged`, you can use the abbreviation `priv`. For example:

```
task.max-lwps=(priv,1K,deny)
```

#### Global and Local Actions on Resource Control Values

There are two categories of actions on resource control values: global and local.

Global actions apply to resource control values for every resource control on the system. You can use [rctladm\(1M\)](#) to perform the following actions:

- Display the global state of active system resource controls.
- Set global logging actions.

You can disable or enable the global logging action on resource controls. You can set the `syslog` action to a specific degree by assigning a severity level, `syslog=level`. The possible settings for `level` are as follows:

- `debug`
- `info`
- `notice`
- `warning`
- `err`
- `crit`
- `alert`
- `emerg`

By default, there is no global logging of resource control violations.

Local actions are taken on a process that attempts to exceed the control value. For each threshold value that is placed on a resource control, you can associate one or more actions. There are three types of local actions: `none`, `deny`, and `signal=`. These three actions are used as follows:

#### `none`

No action is taken on resource requests for an amount that is greater than the threshold. This action is useful for monitoring resource usage without affecting the progress of applications. You can also enable a global message that displays when the resource control is exceeded, while, at the same time, the process exceeding the threshold is not affected.

#### `deny`

You can deny resource requests for an amount that is greater than the threshold. For example, a `task.max-lwps` resource control with action `deny` causes a `fork()` system call to fail if the new process would exceed the control value. See the [fork\(2\)](#).

signal=

You can enable a global signal message action when the resource control is exceeded. A signal is sent to the process when the threshold value is exceeded. Additional signals are not sent if the process consumes additional resources. Available signals are listed below.

Not all of the actions can be applied to every resource control. For example, a process cannot exceed the number of CPU shares assigned to the project of which it is a member. Therefore, a deny action is not allowed on the `project.cpu-shares` resource control.

Due to implementation restrictions, the global properties of each control can restrict the range of available actions that can be set on the threshold value. (See [rctladm\(1M\)](#).) A list of available signal actions is presented in the following list. For additional information about signals, see [signal\(3HEAD\)](#).

The following are the signals available to resource control values:

SIGABRT

Terminate the process.

SIGHUP

Send a hangup signal. Occurs when carrier drops on an open line. Signal sent to the process group that controls the terminal.

SIGTERM

Terminate the process. Termination signal sent by software.

SIGKILL

Terminate the process and kill the program.

SIGSTOP

Stop the process. Job control signal.

SIGXRES

Resource control limit exceeded. Generated by resource control facility.

SIGXFSZ

Terminate the process. File size limit exceeded. Available only to resource controls with the `RCTL_GLOBAL_FILE_SIZE` property (`process.max-file-size`). See [rctlblk\\_set\\_value\(3C\)](#).

SIGXCPU

Terminate the process. CPU time limit exceeded. Available only to resource controls with the `RCTL_GLOBAL_CPU_TIME` property (`process.max-cpu-time`). See [rctlblk\\_set\\_value\(3C\)](#).

#### Resource Control Flags and Properties

Each resource control on the system has a certain set of associated properties. This set of properties is defined as a set of flags, which are associated with all controlled instances of that resource. Global flags cannot be modified, but the flags can be retrieved by using either [rctladm\(1M\)](#) or the [setrctl\(2\)](#) system call.

Local flags define the default behavior and configuration for a specific threshold value of that resource control on a specific process or process collective. The local flags for one threshold value do not affect the behavior of other defined threshold values for the same resource control. However, the global flags affect the behavior for every value associated with a particular control. Local flags can be modified, within the constraints supplied by their corresponding global flags, by the `prctl` command or the `setrctl` system call. See [setrctl\(2\)](#).

For the complete list of local flags, global flags, and their definitions, see [rctlblk\\_set\\_value\(3C\)](#).

To determine system behavior when a threshold value for a particular resource control is reached, use `rctladm` to display the global flags for the resource control. For example, to display the values for `process.max-cpu-time`, enter:

```
$ rctladm process.max-cpu-time
process.max-cpu-time  syslog=off [ lowerable no-deny cpu-time inf seconds ]
```

The global flags indicate the following:

`lowerable`

Superuser privileges are not required to lower the privileged values for this control.

`no-deny`

Even when threshold values are exceeded, access to the resource is never denied.

`cpu-time`

SIGXCPU is available to be sent when threshold values of this resource are reached.

`seconds`

The time value for the resource control.

Use the `prctl` command to display local values and actions for the resource control. For example:

```
$ prctl -n process.max-cpu-time $$
process 353939: -ksh
      NAME  PRIVILEGE  VALUE  FLAG  ACTION  RECIPIENT
process.max-cpu-time
      privileged  18.4Es  inf  signal=XCPU  -
      system      18.4Es  inf  none
```

The `max` (`RCTL_LOCAL_MAXIMAL`) flag is set for both threshold values, and the `inf` (`RCTL_GLOBAL_INFINITE`) flag is defined for this resource control. An `inf` value has an infinite quantity. The value is never enforced. Hence, as configured, both threshold quantities represent infinite values that are never exceeded.

**Resource Control Enforcement** More than one resource control can exist on a resource. A resource control can exist at each containment level in the process model. If resource controls are active on the same resource at different container levels, the smallest container's control is enforced first. Thus, action is taken on `process.max-cpu-time` before `task.max-cpu-time` if both controls are encountered simultaneously.

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

**See Also** [prctl\(1\)](#), [pooladm\(1M\)](#), [poolcfg\(1M\)](#), [projadd\(1M\)](#), [projmod\(1M\)](#), [rctladm\(1M\)](#), [setrctl\(2\)](#), [rctlblk\\_set\\_value\(3C\)](#), [libpool\(3LIB\)](#), [project\(4\)](#), [attributes\(5\)](#), [FSS\(7\)](#), [lofi\(7D\)](#)

*Oracle Solaris Administration: Oracle Solaris Zones, Oracle Solaris 10 Zones, and Resource Management*

**Name** sgml, solbook – Standard Generalized Markup Language

**Description** 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.

In an SGML document the main structural components are called elements. The organization and structure of a document and the meaning of elements are described in the Document Type Definition ( DTD ). Elements are the *tags* that identify the content. Element names may be descriptive of the content for ease of use. For example `<para>` for paragraphs. Elements can have attributes 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.

Another component of the DTD is entities. 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.

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.

**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 `<!DOCTYPE` string is what the SGML -aware `man(1)` command uses to identify that a file is formatted in SGML rather than `nroff(1)`.

**RefEntry** The top layer element that contains a reference page is `<refentry>`. All of the text and other tags must be contained within this tag.

**RefMeta** The next tag in a reference page is `<refmeta>`, which is a container for several other tags. They are:

`<refentrytitle>` 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.

`<manvolnum>` This is the section number that the reference page resides in. The contents may be a text entity reference.

<code>&lt;refmiscinfo&gt;</code>	There are one or more <code>&lt;refmiscinfo&gt;</code> tags which contain <i>meta</i> information. Meta information is information about the reference page. The <code>&lt;refmiscinfo&gt;</code> tag has the <code>class</code> attribute. There are four classes that are routinely used.
date	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.
sectdesc	This is the section title of the reference page; for example User Commands. The value of this attribute may be a text entity reference.
software	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.
arch	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> .
copyright	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>&lt;refmiscinfo&gt;</code> entries. The value of this attribute may be a text entity reference.
RefNameDiv	This tag contains the equivalent information to the <code>.TH</code> macro line in an <code>nroff(1)</code> reference page. <code>&lt;refnamediv&gt;</code> contains three tags. These tags contain the text that is before and after the <code>'-</code> (dash) on the <code>NAME</code> line.
<code>&lt;refname&gt;</code>	These are the names of the topics that are discussed in the file. There may be more than one <code>&lt;refname&gt;</code> for a page. The first <code>&lt;refname&gt;</code> must match the name of the file and the <code>&lt;refentrytitle&gt;</code> . If there are more than one <code>&lt;refname&gt;</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.
<code>&lt;refpurpose&gt;</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.
<code>&lt;refdescriptor&gt;</code>	In some cases the <code>&lt;refentrytitle&gt;</code> is a general topic descriptor of a group of related objects that are discussed on the same page. In this case the first tag after the <code>&lt;refnamediv&gt;</code> is a <code>&lt;refdescriptor&gt;</code> . The <code>&lt;refname&gt;</code> tags follow. Only one <code>&lt;refdescriptor&gt;</code> is allowed, and it

should match the `<refentrytitle>`.

**RefSynopsisDiv** The SYNOPSIS line of the reference page is contained by this tag. There is a `<title>` that usually contains an entity reference. The text is the word SYNOPSIS. There are several tags within `<refsynopsisdiv>` that are designed specifically for the type of synopsis that is used in the different reference page sections. The three types are:

- `<cmdsynopsis>` Used for commands and utilities pages.
- `<funcsynopsis>` Used for programming interface pages.
- `<synopsis>` Used for pages that do not fall into the other two categories.

**RefSect1** This tag is equivalent to the `.SH nroff` macro. It contains a `<title>` 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 `<title>` may be a text entity reference.

**RefSect2** This tag is equivalent to the `.SS nroff` macro. It contains a `<title>` element that contains the text of the sub-section heading. `<refsect2>` tags may also be used within a `<refsynopsisdiv>` as a sub-section heading for the SYNOPSIS section.

**Block Elements** There are a number of block elements that are used for grouping text. This is a list of some of these elements.

- `<para>` This tag is used to contain a paragraph of text.
- `<variablelist>` 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.
- `<orderedlist>` An list of items in a specific order.
- `<itemizedlist>` A list of items that are marked with a character such as a bullet or a dash.
- `<literallayout>` 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.
- `<programlisting>` A segment of program code. Line breaks and leading white space are significant.
- `<table>` This tag contains the layout and content for tabular formatting of information. `<table>` has a required `<title>`.
- `<informaltable>` This tag is the same as the `<table>` tag except the `<title>` is not required.

- `<example>` This tag contains examples of source code or usage of commands. It contains a required `<title>`.
- `<informalexample>` This tag is the same as the `<example>` tag except the `<title>` is not required.

**Inline Elements** The inline elements are used for tagging text.

- `<command>` An executable program or the entry a user makes to execute a command.
- `<function>` A subroutine in a program or external library.
- `<literal>` Contains any literal string.
- `<parameter>` An argument passed to a computer program by a function or routine.
- `<inlineequation>` An untitled mathematical equation occurring in-line.
- `<link>` 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.
- `<olink>` A hypertext link used to create cross references to books other than the reference manual.
- `<xref>` A cross reference to another part of the same reference page.

**See Also** [man\(1\)](#), [nroff\(1\)](#), [man\(5\)](#)

**Name** smf – service management facility

**Description** The Solaris service management facility defines a programming model for providing persistently running applications called *services*. The facility also provides the infrastructure in which to run services. A service can represent a running application, the software state of a device, or a set of other services. Services are represented in the framework by *service instance* objects, which are children of service objects. Instance objects can inherit or override the configuration of the parent service object, which allows multiple service instances to share configuration information. All service and instance objects are contained in a *scope* that represents a collection of configuration information. The configuration of the local Solaris instance is called the “localhost” scope, and is the only currently supported scope.

Each service instance is named with a fault management resource identifier (FMRI) with the scheme `svc:.` For example, the `syslogd(1M)` daemon started at system startup is the default service instance named:

```
svc://localhost/system/system-log:default
svc:/system/system-log:default
system/system-log:default
```

Many commands also allow FMRI abbreviations. See the `svcs(1)` man page for one such example.

In the above example, `default` is the name of the instance and `system/system-log` is the service name. Service names can comprise multiple components separated by slashes (`/`). All components, except the last, compose the *category* of the service. Site-specific services should be named with a category beginning with `site`.

A service instance is either enabled or disabled. All services can be enabled or disabled with the `svcadm(1M)` command.

The list of managed service instances on a system can be displayed with the `svcs(1)` command.

When an administrator deletes an entity backed by a manifest or profile in a standard location, that entity is masked, and is not seen by normal queries to SMF. Masked entity can be explored using `svccfg listcust`, and removed using the `delcust` subcommand to `svccfg`. See `svccfg(1M)` for details.

**Dependencies** Service instances can have dependencies on a set of *entities* which can include services, instances, and files. Dependencies govern when the service is started and automatically stopped. When the dependencies of an enabled service are not satisfied, the service is kept in the offline state. When its dependencies are satisfied, the service is started. If the start is successful, the service is transitioned to the online state. Unlike services and instances, file dependencies are not evaluated dynamically as files are created or deleted. They are evaluated only one time.

Whether a dependency is satisfied is determined by its grouping:

<code>require_all</code>	Satisfied when all cited services are running (online or degraded), or when all indicated files are present.
<code>require_any</code>	Satisfied when one of the cited services is running (online or degraded), or when at least one of the indicated files is present.
<code>optional_all</code>	Satisfied if the cited services are running (online or degraded) or do not run without administrative action (disabled, maintenance, not present, or offline waiting for dependencies which do not start without administrative action). Incomplete services also satisfy optional dependencies.
<code>exclude_all</code>	Satisfied when all of the cited services are disabled, in the maintenance state, or when cited services or files are not present.

Once running (online or degraded), if a service cited by a `require_all`, `require_any`, or `optional_all` dependency is stopped or refreshed, the SMF considers why the service was stopped and the `restart_on` attribute of the dependency to decide whether to stop the service.

event	restart_on value			
	none	error	restart	refresh
stop due to error	no	yes	yes	yes
non-error stop	no	no	yes	yes
refresh	no	no	no	yes

A service is considered to have stopped due to an error if the service has encountered a hardware error or a software error such as a core dump. For `exclude_all` dependencies, the service is stopped if the cited service is started and the `restart_on` attribute is not none.

The dependencies on a service can be listed with `svcs(1)` or `svccfg(1M)`, and modified with `svccfg(1M)`.

**Restarters** Each service is managed by a restarter. The master restarter, `svc.startd(1M)` manages states for the entire set of service instances and their dependencies. The master restarter acts on behalf of its services and on delegated restarters that can provide specific execution environments for certain application classes. For instance, `inetd(1M)` is a delegated restarter that provides its service instances with an initial environment composed of a network connection as input and output file descriptors. Each instance delegated to `inetd(1M)` is in the online state. While the daemon of a particular instance might not be running, the instance is available to run.

As dependencies are satisfied when instances move to the online state, `svc.startd(1M)` invokes start methods of other instances or directs the delegated restarter to do so. These operations might overlap.

The current set of services and associated restarters can be examined using `svcs(1)`. A description of the common configuration used by all restarters is given in `smf_restarter(5)`.

**Methods** Each service or service instance must define a set of methods that start, stop, and, optionally, refresh the service. See `smf_method(5)` for a more complete description of the method conventions for `svc.startd(1M)` and similar `fork(2)-exec(2)` restarters.

Administrative methods, such as for the capture of legacy configuration information into the repository, are discussed on the `svccfg(1M)` manual page.

The methods for a service can be listed and modified using the `svccfg(1M)` command.

**States** Each service instance is always in a well-defined state based on its dependencies, the results of the execution of its methods, and its potential contracts events. The following states are defined:

UNINITIALIZED	This is the initial state for all service instances. Instances are moved to maintenance, offline, or a disabled state upon evaluation by <code>svc.startd(1M)</code> or the appropriate restarter.
OFFLINE	The instance is enabled, but not yet running or available to run. If restarter execution of the service start method or the equivalent method is successful, the instance moves to the online state. Failures might lead to a degraded or maintenance state. Administrative action can lead to the uninitialized state.
ONLINE	The instance is enabled and running or is available to run. The specific nature of the online state is application-model specific and is defined by the restarter responsible for the service instance. Online is the expected operating state for a properly configured service with all dependencies satisfied. Failures of the instance can lead to a degraded or maintenance state. Failures of services on which the instance depends can lead to offline or degraded states.
DEGRADED	The instance is enabled and running or available to run. The instance, however, is functioning at a limited capacity in comparison to normal operation. Failures of the instance can lead to the maintenance state. Failures of services on which the instance depends can lead to offline or degraded states. Restoration of capacity should result in a transition to the online state.
MAINTENANCE	The instance is enabled, but not able to run. Administrative action (through <code>svcadm clear</code> ) is required to move the instance out of the maintenance state. The maintenance state might be a temporarily reached state if an administrative operation is underway.

DISABLED	The instance is disabled. Enabling the service results in a transition to the offline state and eventually to the online state with all dependencies satisfied.
LEGACY - RUN	This state represents a legacy instance that is not managed by the service management facility. Instances in this state have been started at some point, but might or might not be running. Instances can only be observed using the facility and are not transferred into other states.

States can also have transitions that result in a return to the originating state.

Events Notification SMF allows notification by using SNMP or SMTP of state transitions. It publishes Information Events for state transitions which are consumed by notification daemons like `snmp-notify(1M)` and `smtp-notify(1M)`. SMF state transitions of disabled services do not generate notifications unless the final state for the transition is disabled and there exist notification parameters for that transition. Notification is not generated for transitions that have the same initial and final state.

Notification Parameters Notification parameters for FMA Events are stored in `svc:/system/fm/notify-params:default` except for Information Events generated by SMF state transitions. Those are stored in the service or in the instance of the transitioning service. Notification parameters for SMF state transition generated events can be set system wide in `svc:/system/svc/global:default`. The system wide notification parameters are used when a composed lookup, as in `scf_instance_get_pg_composed(3SCF)`, in the transitioning instance cannot be found. Notification parameters can be manipulated using `svccfg(1M)`. Notification parameters can be configured in a service manifest or profile using the `notification_parameters` element described in the DTD. An example is provided below:

```
<notification_parameters>
  <event value='from-online' />
  <type name='smtp' active="false">
    <parameter name='to'>
      <value_node value='root@local' />
      <value_node value='admin-alias@eng' />
    </parameter>
  </type>
  <type name='snmp' />
</notification_parameters>
```

*events* is a comma separated list of SMF state transition sets or a comma separated list of FMA event classes. *events* cannot have a mix of SMF state transition sets and FMA event classes.

For convenience, the tags `problem- {diagnosed, updated, repaired, resolved}` describe the lifecycle of a problem diagnosed by the FMA subsystem - from initial diagnosis to interim updates and finally problem closure. These tags are aliases for underlying FMA protocol event classes (all in the `list.*` hierarchy), but the latter should not be used in configuring notification preferences.

**problem-diagnosed**

A new problem has been diagnosed by the FMA subsystem. The diagnosis includes a list of one or more suspects, which (where appropriate) might have been automatically isolated to prevent further errors occurring. The problem is identified by a UUID in the event payload, and further events describing the resolution lifecycle of this problem quote a matching UUID.

**problem-updated**

One or more of the suspect resources in a problem diagnosis has been repaired, replaced or acquitted (or has been faulted again), but there remains at least one faulted resource in the list. A repair could be the result of an `fmadm` command line (`fmadm repaired`, `fmadm acquit`, `fmadm replaced`) or might have been detected automatically such as through detection of a part serial number change.

**problem-repaired**

All of the suspect resources in a problem diagnosis have been repaired, resolved or acquitted. Some or all of the resources might still be isolated at this stage.

**problem-resolved**

All of the suspect resources in a problem diagnosis have been repaired resolved or acquitted *and* are no longer isolated (for example, a cpu that was a suspect and offlined is now back online again; this un-isolate action is usually automatic).

State Transition Sets are defined as:

<code>to-&lt;state&gt;</code>	Set of all transitions that have <code>&lt;state&gt;</code> as the final state of the transition.
<code>from-&lt;state&gt;</code>	Set of all transitions that have <code>&lt;state&gt;</code> as the initial state of the transition.
<code>&lt;state&gt;</code>	Set of all transitions that have <code>&lt;state&gt;</code> as the initial state of the transition.
<code>all</code>	Set of all transitions.

Valid values of state are maintenance, offline, disabled, online and degraded. An example of a transitions set definition: `maintenance, from-online, to-degraded`.

**Properties and Property Groups**

The dependencies, methods, delegated restarter, and instance state mentioned above are represented as properties or property groups of the service or service instance. A service or service instance has an arbitrary number of property groups in which to store application data. Using property groups in this way allows the configuration of the application to derive the attributes that the repository provides for all data in the facility. The application can also use the appropriate subset of the [service\\_bundle\(4\)](#) DTD to represent its configuration data within the framework.

Property lookups are composed. If a property group-property combination is not found on the service instance, most commands and the high-level interfaces of [libscf\(3LIB\)](#) search for the same property group-property combination on the service that contains that instance. This allows common configuration among service instances to be shared. Composition can be viewed as an inheritance relationship between the service instance and its parent service.

Properties are protected from modification by unauthorized processes. See [smf\\_security\(5\)](#).

General Property Group	The general property group applies to all service instances. It includes the following properties:
enabled (boolean)	Specifies whether the instance is enabled. If this property is not present on an instance, SMF does not tell the instance's restarter about the existence of the instance.
restarter (fmri)	The restarter for this service. See the Restarters section for more information. If this property is unset, the default system restarter is used.
complete (astring)	Whether this service is complete or is a partial definition that should not be started. This property is automatically set on manifest import. Alternatively, an instance without this property that successfully validates against the template definitions (see <a href="#">scf_tmpl_validate_fmri(3SCF)</a> ) will have this property created by <a href="#">svcadm(1M)</a> on enable.

**Layers** The repository is assembled from a combination of administrative customization, current state, and default values from files in standard locations. Services, instances, property groups, and properties defined by manifests in SMF-managed filesystem locations are always accurately represented in the repository. Customizations made during runtime by administrators or other programs are captured and stored in the repository.

A property can have different values in the repository which reflect different settings from manifests, profiles, and administrative customizations. Which one is presented to the user and service by default is arbitrated by a simple priority scheme called **layers**.

Four layers are tracked by SMF. In decreasing priority order, they are:

admin	Any change made by interactive use of SMF commands or libraries. This layer has the highest priority.
site-profile	Any values from the files in the <code>/etc/svc/profile/site</code> directory, or the legacy <code>/etc/svc/profile/site.xml</code> and <code>/var/svc/profile/site.xml</code> files.
system-profile	Any values from the system profile locations <code>/etc/svc/profile/generic.xml</code> and <code>/etc/svc/profile/platform.xml</code>
manifest	Any values from the system manifest locations <code>/lib/svc/manifest</code> or <code>/var/svc/manifest</code> .

Property conflicts are not permitted within any individual layer. A conflicting property in the `admin` layer simply overwrites the previous property. If the same property is delivered by multiple files in any other layer, and is not set at a higher layer, the entire instance is tagged as

in-conflict, and are not started by `svc.startd(1M)` until the conflicting definition is removed or the property is set at a higher layer. Other `libscf` consumers requesting a single value, including `svccfg` and `svccprop`, see a random property setting from amongst all appropriate values. We do not guarantee which of the conflicting values are returned.

**Snapshots** Historical data about each instance in the repository is maintained by the service management facility. This data is made available as read-only snapshots for administrative inspection and rollback. The following set of snapshot types might be available:

<code>initial</code>	Initial configuration of the instance created by the administrator or produced during package installation.
<code>previous</code>	Current configuration captured when an administrative undo operation is performed.
<code>running</code>	The running configuration of the instance.
<code>start</code>	Configuration captured during a successful transition to the online state.

The `svccfg(1M)` command can be used to interact with snapshots.

**Special Property Groups** Some property groups are marked as non-persistent. These groups are not backed up in snapshots and their content is cleared during system boot. Such groups generally hold an active program state which does not need to survive system restart.

**Configuration Repository** The current state of each service instance, as well as the properties associated with services and service instances, is stored in a system repository managed by `svc.configd(1M)`.

The repository for service management facility data is managed by `svc.configd(1M)`.

**Service Bundles, Manifests, and Profiles** The information associated with a service or service instance that is stored in the configuration repository can be exported as XML-based files. Such XML files, known as service bundles, are portable and suitable for backup purposes. Service bundles are classified as one of the following types:

<code>manifests</code>	Files that contain the complete set of properties associated with a specific set of services or service instances.
<code>profiles</code>	Files that contain a set of service instances and values for the enabled property (type <code>boolean</code> in the general property group) on each instance.

Profiles can also contain configuration values for properties in services and instances. Template elements cannot be defined in a profile.

Profiles can use a relaxed set of elements from the DTD described in `service_bundle(4)`. To use these, the `DOCTYPE` entry should have the following definitions added:

```
<!ENTITY % profile "INCLUDE">
<!ENTITY % manifest "IGNORE">
```

Service bundles can be imported or exported from a repository using the [svccfg\(1M\)](#) command. See [service\\_bundle\(4\)](#) for a description of the service bundle file format with guidelines for authoring service bundles.

**Milestones** An smf milestone is a service that aggregates a multiple service dependencies. Usually, a milestone does nothing useful itself, but declares a specific state of system-readiness on which other services can depend. One example is the `name-services` milestone, which simply depends upon the currently enabled `name-services`.

**Legacy Startup Scripts** Startup programs in the `/etc/rc?.d` directories are executed as part of the corresponding run-level milestone:

```
/etc/rcS.d    milestone/single-user:default
/etc/rc2.d    milestone/multi-user:default
/etc/rc3.d    milestone/multi-user-server:default
```

Execution of each program is represented as a reduced-functionality service instance named by the program's path. These instances are held in a special `legacy-run` state.

These instances do not have an `enabled` property (type `boolean` in the general property group) and, generally, cannot be manipulated with the [svcadm\(1M\)](#) command. No error diagnosis or restart is done for these programs.

**See Also** [svcs\(1\)](#), [inetd\(1M\)](#), [snmp-notify\(1M\)](#), [smtp-notify\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [svc.configd\(1M\)](#), [svc.startd\(1M\)](#), [exec\(2\)](#), [fork\(2\)](#), [libscf\(3LIB\)](#), [scf\\_tmpl\\_validate\\_fmri\(3SCF\)](#), [strftime\(3C\)](#), [contract\(4\)](#), [service\\_bundle\(4\)](#), [smf\\_bootstrap\(5\)](#), [smf\\_method\(5\)](#), [smf\\_restarter\(5\)](#), [smf\\_security\(5\)](#)

<b>Name</b>	smf_bootstrap – service management facility boot, packaging, and compatibility behavior
<b>Description</b>	The service management facility establishes conventions for delivering service manifests, incorporating service manifest changes, describing service configuration stability, using service configuration overrides, and the use of service profiles.
Manifest Loading at Boot	<p>Manifests from the standard directory trees <code>/lib/svc/manifest</code> and <code>/var/svc/manifest</code> are processed during system boot and anytime an administrator or program runs:</p> <pre>\$ svcadm restart manifest-import</pre> <p>Manifests that have not been imported previously or have changed since the last time they were imported are processed. A hash is used to determine whether a manifest has changed.</p> <p>When a manifest in a standard location is imported for the first time, its properties, instances, and services are added to the repository as part of the <code>manifest</code> layer.</p> <p>Manifests in standard locations are automatically imported when they are updated. New services and instances are added, properties are upgraded if they are changed, and services, instances, and properties are deleted if they are removed.</p> <p>Manifests are processed in two different phases during boot.</p> <p>The service <code>svc:/system/early-manifest-import:default</code>, a pseudo service, is responsible for the first manifest processing. This service processes only manifests from the <code>/lib/svc/manifest</code> directory tree before <code>svc.startd(1M)</code> initializes any services thus enabling services delivered in <code>/lib/svc/manifest</code> to always start with their most updated definition. Since this is a pseudo service, <code>svcadm(1M)</code> commands are ignored though <code>svcs(1)</code> can be used to observe status and get log file information.</p> <p>The <code>svc:/system/manifest-import:default</code> service handles the second manifest processing and imports manifest files from both <code>/lib/svc/manifest</code> and <code>/var/svc/manifest</code> directory trees, in that respective order.</p> <p>Support for <code>/var/svc/manifest</code> is compatibility support for manifests delivered in that directory tree prior to the introduction of <code>system/early-manifest-import:default</code>. Services delivered in <code>/var/svc/manifest</code> can run into upgrade-related issues where a service might be started with an old repository configuration because its updated manifest is not yet imported. Similarly, a newly added service might not be available or a deleted service is still started during boot because its manifest file has not been processed. Developers are strongly encouraged to move a manifest to <code>/lib/svc/manifest</code> to avoid these issues.</p> <p>Only <code>*.xml</code> files are considered for the operations described in this section.</p>
Profile Application	Profiles are also applied by the <code>early-manifest-import</code> and <code>manifest-import</code> services.

The system-delivered profiles in `/etc/svc/profile/generic.xml` and `/etc/svc/profile/platform.xml` are imported into the `system-profile` layer.

Site-specific profiles in the `/etc/svc/profile/site` directory and legacy site files `/etc/svc/profile/site.xml` and `/var/svc/profile/site.xml` are imported into the `site-profile` layer.

Administrators can request that these profiles are reapplied by running:

```
$ svcadm restart manifest-import
```

The behavior of properties, instances, and services defined by profiles is identical to those defined by manifests.

Only `*.xml` files are considered for the operations described in this section.

**Manifest Handling During Packaging Operations** Service manifests within packages should be identified with the class `manifest`. Class action scripts that install and remove service manifests are included in the packaging subsystem. When `pkg install` is invoked, the service manifest is imported.

When `pkg uninstall` is invoked, instances in the manifest that are disabled are deleted. Instances in the manifest that are online or degraded are disabled first and then deleted. Any services in the manifest with no remaining instances are also deleted.

**Stability Declarations** Each service group and each property group delivered in a manifest should declare a stability level based on [attributes\(5\)](#) definitions. With knowledge of the stability level, an application developer can determine the likelihood that feature development based on the existence or components of a service or object is likely to remain functional across a release boundary.

In an [smf\(5\)](#) context, the stability value also identifies the expected scope of the changes to properties within the property group across a release boundary for the service, which can include patches for that service. The following two sections discuss this in more detail.

**Property Group Deletion** The [service\\_bundle\(4\)](#) document type definition includes a `delete` attribute, applicable to each property group in a service manifest. If set to `true`, the `delete` attribute instructs [svccfg\(1M\)](#) and other manifest import tools to delete this property group from the repository. If the `delete` attribute is absent or present but set to `false`, the property group in the repository is preserved.

Property groups declared as `Stable` or `Evolving` are not deleted. Property groups declared as `Unstable` can be deleted across any release boundary.

**See Also** [svcs\(1\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [svc.startd\(1M\)](#), [libscf\(3LIB\)](#), [service\\_bundle\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [smf\\_security\(5\)](#)

[pkg\(1\)](#)

**Notes** The present version of [smf\(5\)](#) does not support multiple repositories.

**Name** smf\_method – service management framework conventions for methods

**Description** The class of services managed by `svc.startd(1M)` in the service management framework, `smf(5)`, consists of applications that fit a simple `fork(2)-exec(2)` model. The `svc.startd(1M)` master daemon and other restarters support the `fork(2)-exec(2)` model, potentially with additional capabilities. The `svc.startd(1M)` daemon and other restarters require that the methods which activate, manipulate, or examine a service instance follow the conventions described in this manual page.

**Invocation form** The form of a method invocation is not dictated by convention. In some cases, a method invocation might consist of the direct invocation of the daemon or other binary executable that provides the service. For cases in which an executable script or other mediating executable is used, the convention recommends the form:

```
/path/to/method_executable abbr_method_name
```

The *abbr\_method\_name* used for the recommended form is a supported method such as `start` or `stop`. The set of methods supported by a restarter is given on the related restarter page. The `svc.startd(1M)` daemon supports `start`, `stop`, and `refresh` methods.

A restarter might define other kinds of methods beyond those referenced in this page. The conventions surrounding such extensions are defined by the restarter and might not be identical to those given here.

**Environment Variables** The restarter provides four environment variables to the method that determine the context in which the method is invoked.

**SMF\_FMRI**

The service fault management resource identifier (FMRI) of the instance for which the method is invoked.

**SMF\_METHOD**

The full name of the method being invoked, such as `start` or `stop`.

**SMF\_RESTARTER**

The service FMRI of the restarter that invokes the method

**SMF\_ZONENAME**

The name of the zone in which the method is running. This can also be obtained by using the `zonename(1)` command.

These variables should be removed from the environment prior to the invocation of any persistent process by the method. A convenience shell function, `smf_clear_env`, is given for service authors who use Bourne-compatible shell scripting to compose service methods in the include file described below.

The method context can cause other environment variables to be set as described below.

**Method Definition** A method is defined minimally by three properties in a propertygroup of type `method`.

These properties are:

<code>exec</code> ( <i>astring</i> )	Method executable string.
<code>timeout_seconds</code> ( <i>count</i> )	Number of seconds before method times out. See the <code>Timeouts</code> section for more detail.
<code>type</code> ( <i>astring</i> )	Method type. Currently always set to <code>method</code> .

A `Method Context` can be defined to further refine the execution environment of the method. See the `Method Context` section for more information.

**Method Tokens** When defined in the `exec` string of the method by the restarter `svc.startd`, a set of tokens are parsed and expanded with appropriate value. Other restarters might not support method tokens. The delegated restarter for inet services, `inetd(1M)`, does not support the following method expansions.

`%%`

`%`

`%r`

Name of the restarter, such as `svc.startd`

`%m`

The full name of the method being invoked, such as `start` or `stop`.

`%s`

Name of the service

`%i`

Name of the instance

`%f`

FMRI of the instance

`%{prop[: , ]}`

Value(s) of a property. The `prop` might be a property FMRI, a property group name and a property name separated by a `/`, or a property name in the `application` property group. These values can be followed by a `,` (comma) or `:` (colon). If present, the separators are used to separate multiple values. If absent, a space is used. The following shell metacharacters encountered in string values are quoted with a `\` (backslash):

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

An invalid expansion constitutes method failure.

Two explicit tokens can be used in the place of method commands.

`:kill [-signal]`

Sends the specified signal, which is SIGTERM by default, to all processes in the primary instance contract. Always returns SMF\_EXIT\_OK. This token should be used to replace common `kill` invocations.

`:true`

Always returns SMF\_EXIT\_OK. This token should be used for methods that are required by the restarter but which are unnecessary for the particular service implementation.

**Exiting and Exit Status** The required behavior of a start method is to delay exiting until the service instance is ready to answer requests or is otherwise functional.

The following exit status codes are defined in `<libscf.h>` and in the shell support file.

SMF_EXIT_OK	0	Method exited, performing its operation successfully.
SMF_EXIT_ERR_FATAL	95	Method failed fatally and is unrecoverable without administrative intervention.
SMF_EXIT_ERR_CONFIG	96	Unrecoverable configuration error. A common condition that returns this exit status is the absence of required configuration files for an enabled service instance.
SMF_EXIT_ERR_NOSMF	99	Method has been mistakenly invoked outside the <a href="#">smf(5)</a> facility. Services that depend on <a href="#">smf(5)</a> capabilities should exit with this status value.
SMF_EXIT_ERR_PERM	100	Method requires a form of permission such as file access, privilege, authorization, or other credential that is not available when invoked.
SMF_EXIT_ERR_OTHER	non-zero	Any non-zero exit status from a method is treated as an unknown error. A series of unknown errors can be diagnosed as a fault by the restarter or on behalf of the restarter.

In addition to the exit codes described above, a method may use the following exit codes in conjunction with `smf_method_exit()`, available via `smf_method_exit(3SCF)` and `smf_include.sh`:

SMF_EXIT_TEMP_DISABLE	101	Method exits successfully and requests a temporary disable.
SMF_EXIT_TEMP_TRANSIENT	105	Method exits successfully and requests that it be treated as if its service model was “transient”.

Use of a precise exit code allows the responsible restarters to categorize an error response as likely to be intermittent and worth pursuing restart or permanent and request administrative intervention.

**Timeouts** Each method can have an independent timeout, specified in seconds. The method timeout is specified by the `timeout_seconds` property.

A timeout is used as a last resort for the service's restarters to determine that a method has hung or is not making progress. If a timeout elapsed, many restarters place the service into the maintenance state. See [`svc.startd\(1M\)`](#). A significant margin of error is recommended when specifying a timeout in order to avoid premature failures when the method is making progress, but the system is temporarily responding very slowly due to memory, CPU, or I/O load.

60 seconds is a good starting value for methods expected to take only a second or two. 300 seconds (5 minutes) is appropriate for a method which commonly takes 30 seconds. Scale up as appropriate for methods which routinely take longer.

Shorter timeouts can be used if fast failure is desired in order to prompt administrative intervention. If administrative intervention is likely to only be to clear the service and start the method again, consider a longer timeout.

If `timeout_seconds` is set to 0, there is no timeout for the service. This setting is not preferred, but is available for services which absolutely require it. -1 is also accepted to specify no timeout, but is deprecated.

**Shell Programming Support** A set of environment variables that define the above exit status values is provided with convenience shell functions in the file `/lib/svc/share/smf_include.sh`. This file is a Bourne shell script suitable for inclusion via the source operator in any Bourne-compatible shell.

To assist in the composition of scripts that can serve as SMF methods as well as `/etc/init.d` scripts, the `smf_present()` shell function is provided. If the [`smf\(5\)`](#) facility is not available, `smf_present()` returns a non-zero exit status.

One possible structure for such a script follows:

```
if smf_present; then
    # Shell code to run application as managed service
    ....

    smf_clear_env
else
    # Shell code to run application as /etc/init.d script
    ....
fi
```

This example shows the use of both convenience functions that are provided.

**Method Context** The service management facility offers a common mechanism to set the context in which the `fork(2)-exec(2)` model services execute.

The desired method context should be provided by the service developer. All service instances should run with the lowest level of privileges possible to limit potential security compromises.

A method context can contain the following properties:

**use\_profile**

A boolean that specifies whether the profile should be used instead of the `user`, `group`, `privileges`, and `limit_privileges` properties.

**environment**

Environment variables to insert into the environment of the method, in the form of a number of `NAME=value` strings.

**profile**

The name of an RBAC (role-based access control) profile which, along with the method executable, identifies an entry in `exec_attr(4)`.

**user**

The user ID in numeric or text form.

**group**

The group ID in numeric or text form.

**supp\_groups**

An optional string that specifies the supplemental group memberships by ID, in numeric or text form.

**privileges**

An optional string specifying the privilege set as defined in `privileges(5)`. An Extended Policy can be specified here.

**limit\_privileges**

An optional string specifying the limit privilege set as defined in `privileges(5)`.

**working\_directory**

The home directory from which to launch the method. `:home` can be used as a token to indicate the home directory of the user whose `uid` is used to launch the method. If the property is unset, `:home` is used.

**corefile\_pattern**

An optional string that specifies the corefile pattern to use for the service, as per `coreadm(1M)`. Most restarters supply a default. Setting this property overrides local customizations to the global core pattern.

**project**

The project ID in numeric or text form. `:default` can be used as a token to indicate a project identified by `getdefaultproj(3PROJECT)` for the user whose `uid` is used to launch the method.

**resource\_pool**

The resource pool name on which to launch the method. `:default` can be used as a token to indicate the pool specified in the [project\(4\)](#) entry given in the project attribute above.

The method context can be set for the entire service instance by specifying a `method_context` property group for the service or instance. A method might override the instance method context by providing the method context properties on the method property group.

Invalid method context settings always lead to failure of the method, with the exception of invalid environment variables that issue warnings.

In addition to the context defined above, many [fork\(2\)-exec\(2\)](#) model restarters also use the following conventions when invoking executables as methods:

**Argument array**

The arguments in `argv[]` are set consistently with the result of `/bin/sh -c` of the exec string.

**File descriptors**

File descriptor `0` is `/dev/null`. File descriptors `1` and `2` are recommended to be a per-service log file.

**Files** `/lib/svc/share/smf_include.sh`  
Definitions of exit status values.

`/usr/include/libscf.h`  
Definitions of exit status codes.

**Examples** **EXAMPLE 1** Report a service-specific configuration error message in the service log.

A start method might want to use `smf_method_exit()` to report a service-specific configuration error message in the service log.

```
if [ ! -s "$my_config_file" ]; then
    smf_method_exit $SMF_EXIT_ERR_CONFIG \
        missing_or_empty_config_file \
        "$my_config_file is missing or empty"
fi
```

**EXAMPLE 2** Disable a service that should only run in the global zone when started in a non-global zone.

A service that should only run in the global zone might want to disable itself when started in a non-global zone.

```
if smf_is_nonglobalzone; then
    smf_method_exit $SMF_EXIT_TEMP_DISABLE global_zone_only \
        "$SMF_FMRI is not supported in a local zone" SUNW_OST_OSCMD
fi
```

**See Also** zonename(1), coreadm(1M), inetd(1M), svccfg(1M), svc.startd(1M), exec(2), fork(2), getdefaultproj(3PROJECT), smf\_method\_exit(3SCF), exec\_attr(4), project(4), service\_bundle(4), attributes(5), privileges(5), rbac(5), smf(5), smf\_bootstrap(5), zones(5)

**Notes** The present version of `smf(5)` does not support multiple repositories.

When a service is configured to be started as root but with privileges different from `limit_privileges`, the resulting process is privilege aware. This can be surprising to developers who expect `seteuid(<non-zero UID>)` to reduce privileges to basic or less.

**Name** smf\_restarter – service management facility conventions for restarters

**Description** All service instances in the service management facility must be managed by a restarter. This manual page describes configuration, functionality, and reporting characteristics that are common to all restarters in the framework. Characteristics specific to a particular restarter are described in the restarter's man page.

For each managed service, a restarter relies on retrieving properties on the service instance to determine configuration. The restarter manages a set of property groups to communicate the current disposition of a service with display tools such as [svcs\(1\)](#).

**Service Configuration** The common restarter configuration for all services is captured in the `general` property group. This group includes the following required and optional property settings.

<code>enabled</code>	This is a required property. If set, the restarter of an instance attempts to maintain availability of the service.
<code>restarter</code>	This is an optional property that allows the specification of an alternate restarter to manage the service instance. If the restarter property is empty or absent, the restarter defaults to <a href="#">svc.startd(1M)</a> .
<code>single_instance</code>	This is an optional property. When set, only one instance of the service is allowed to transition to an online or degraded status at any time.

**Service Reporting** All restarters report status using the `restarter` property group, which includes the following properties:

<code>next_state</code>	The current state and next state, if currently in transition, for instances stored in these properties. See <a href="#">smf(5)</a> for a description of the potential states.
<code>auxiliary_state</code>	An astring with no spaces that contains a precise term to describe the full restarter-specific state in combination with the restarter state property. The auxiliary state cannot always be set and is always cleared during transition out of any state. Each restarter must define the precise list of auxiliary states it uses.
<code>state_timestamp</code>	The time when the current state was reached.
<code>contract</code>	The primary process contract ID, if any, under which the service instance is executing.

**See Also** [svcs\(1\)](#), [svc.startd\(1M\)](#), [service\\_bundle\(4\)](#), [smf\(5\)](#), [smf\\_method\(5\)](#)

**Name** smf\_security – service management facility security behavior

**Description** The configuration subsystem for the service management facility, [smf\(5\)](#), requires privilege to modify the configuration of a service. Privileges are granted to a user by associating the authorizations described below to the user through [user\\_attr\(4\)](#) and [prof\\_attr\(4\)](#). See [rbac\(5\)](#).

The following authorization is used to manipulate services and service instances.

`solaris.smf.modify` Authorized to add, delete, or modify services, service instances, or their properties, and to read protected property values.

**Property Group Authorizations** The [smf\(5\)](#) configuration subsystem associates properties with each service and service instance. Related properties are grouped. Groups can represent an execution method, credential information, application data, or restarter state. The ability to create or modify property groups can cause [smf\(5\)](#) components to perform actions that can require operating system privilege. Accordingly, the framework requires appropriate authorization to manipulate property groups.

Each property group has a type corresponding to its purpose. The core property group types are `method`, `dependency`, `application`, and `framework`. Additional property group types can be introduced, provided they conform to the extended naming convention in [smf\(5\)](#). The following basic authorizations, however, apply only to the core property group types:

`solaris.smf.modify.method`  
Authorized to change values or create, delete, or modify a property group of type `method`.

`solaris.smf.modify.dependency`  
Authorized to change values or create, delete, or modify a property group of type `dependency`.

`solaris.smf.modify.application`  
Authorized to change values, read protected values, and create, delete, or modify a property group of type `application`.

`solaris.smf.modify.framework`  
Authorized to change values or create, delete, or modify a property group of type `framework`.

`solaris.smf.modify`  
Authorized to add, delete, or modify services, service instances, or their properties, and to read protected property values.

Property group-specific authorization can be specified by properties contained in the property group.

`modify_authorization`  
Authorizations allow the addition, deletion, or modification of properties within the property group, and the retrieval of property values from the property group if protected.

**value\_authorization**

Authorizations allow changing the values of any property of the property group except `modify_authorization`, and the retrieval of any property values except `modify_authorization` from the property group if protected.

**read\_authorization**

Authorizations allow the retrieval of property values within the property group. The presence of a string-valued property with this name identifies the containing property group as protected. This property has no effect on property groups of types other than application. See Protected Property Groups.

The above authorization properties are only used if they have type `string`. If an instance property group does not have one of the properties, but the instance's service has a property group of the same name with the property, its values are used.

**Protected Property Groups**

Normally, all property values in the repository can be read by any user without explicit authorization. Property groups of non-framework types can be used to store properties with values that require protection. They must not be revealed except upon proper authorization. A property group's status as protected is indicated by the presence of a string-valued `read_authorization` property. If this property is present, the values of all properties in the property group is retrievable only as described in Property Group Authorizations.

Administrative domains with policies that prohibit backup of data considered sensitive should exclude the SMF repository databases from their backups. In the face of such a policy, non-protected property values can be backed up by using the `svccfg(1M)` archive command to create an archive of the repository without protected property values.

**Service Action Authorization**

Certain actions on service instances can result in service interruption or deactivation. These actions require an authorization to ensure that any denial of service is a deliberate administrative action. Such actions include a request for execution of the refresh or restart methods, or placement of a service instance in the maintenance or other non-operational state. The following authorization allows such actions to be requested:

```
solaris.smf.manage    Authorized to request restart, refresh, or other state modification of
                        any service instance.
```

In addition, the `general/action_authorization` property can specify additional authorizations that permit service actions to be requested for that service instance. The `solaris.smf.manage` authorization is required to modify this property.

**Defined Rights Profiles**

Two rights profiles are included that offer grouped authorizations for manipulating typical `smf(5)` operations.

**Service Management**

A service manager can manipulate any service in the repository in any way. It corresponds to the `solaris.smf.manage` and `solaris.smf.modify` authorizations.

The service management profile is the minimum required to use the [pkg\(1\)](#) command to add or remove software packages that contain an inventory of services in its service manifest.

#### Service Operator

A service operator has the ability to enable or disable any service instance on the system, as well as request that its restart or refresh method be executed. It corresponds to the `solaris.smf.manage` and `solaris.smf.modify.framework` authorizations.

Sites can define additional rights profiles customized to their needs.

**Remote Repository Modification** Remote repository servers can deny modification attempts due to additional privilege checks. See NOTES.

**Examples** **EXAMPLE 1** Allow user to modify `system/cron` services without becoming root.

Adding the following line to `/etc/user_attr` allows the user “johndoe” to restart, enable, disable or other state modification of `system/cron` service without becoming root.

```
johndoe:::auths=solaris.smf.manage.cron
```

**EXAMPLE 2** Allow user to modify any property on any service and modify `system/cron` services without becoming root.

Adding the following line to `/etc/user_attr` allows the user “janedoe” to modify any property on any service, and restart, enable, disable or other state modification of `system/cron` service without becoming root

```
janedoe:::auths=solaris.smf.modify,solaris.smf.manage.cron
```

**See Also** [auths\(1\)](#), [profiles\(1\)](#), [svccfg\(1M\)](#), [prof\\_attr\(4\)](#), [user\\_attr\(4\)](#), [rbac\(5\)](#), [smf\(5\)](#)  
[pkg\(1\)](#)

**Notes** The present version of [smf\(5\)](#) does not support remote repositories.

When a service is configured to be started as root but with privileges different from `limit_privileges`, the resulting process is privilege aware. This can be surprising to developers who expect `seteuid(<non-zero UID>)` to reduce privileges to basic or less.

---

<b>Name</b>	smf_template – service management framework support for service metadata
<b>Description</b>	<p>Templates are defined by service developers to describe metadata about a service in general or individual configuration properties on a service, including human-consumable descriptions as well as definitions of valid configuration.</p> <p>Administrators are provided access to templates through SMF commands that describe configuration values and validate configuration against templates.</p> <p>Tool developers can use templates to provide more helpful user interfaces for service configuration.</p>
Template Data	Service metadata is defined in the template as part of the service manifest.
Consuming Template Data	<p>The <code>svcs -lv</code> and <code>svccfg describe</code> commands can be used to access metadata about properties in a human-readable format.</p> <p><code>svccfg(1M)</code>'s <code>validate</code> subcommand can be used to validate a service instance or manifest against template data. A set of <code>libscf(3LIB)</code> interfaces is available to access template data.</p>
Template Definition	<p>The sole interface to define templates is the service manifest.</p> <p>Service authors should provide template metadata including <code>common_names</code>, <code>descriptions</code>, <code>choices</code> and <code>constraints</code> for service-specific property groups and properties which they introduce. At a minimum, service authors must provide descriptions for property groups and properties in the C locale. Service authors should not provide template metadata for framework-delivered property groups such as methods and dependencies.</p> <p>See the <code>EXAMPLES</code> section for an example of authoring a template definition for a service.</p>
Template Composition	<p>All template interfaces search for template data about a property group first on the instance, then on the service, then on the service's restarter, and finally globally.</p> <p>A property group template is defined by its author to apply to a specific instance, to a service and all of its instances, to a restarter's delegates, or globally. A typical service author defines the template on an instance or on a service. A template defined on an instance is applied to that instance only, and can override a template for that property group defined on the service. A template defined on the service is applied to all instances of that service.</p> <p>Restarter authors can define templates in their manifest that apply to any service which uses their restarter, which is also known as a <i>delegate</i>. SMF framework authors have defined templates for property groups with well-known meanings to the entire SMF framework in the manifest for <code>svc:/system/svc/global</code>.</p> <p>Templates defined globally or by the restarter and re-defined by the service or instance are flagged as a validation error. Service authors can avoid these errors by creating templates only for property groups specific to their service and not consumed by the SMF framework.</p> <p>Property group templates can also be wildcarded by name or type. Only the most specific template definition applicable to a property group is honored.</p>

## Template Details

### Service and Instance Templates

The `template` element defines the start of a template block. All further definitions below can be included in a template block. A `template` element can be contained in either a `service` or `instance` element. If it is contained in the `service` element, it applies to the service and all instances of that service. If it is contained in the `instance` element, it applies to only that instance of the service.

Whenever possible, we recommend defining the template data for the entire service.

```
<service ... >
  <template>
  </template>
</service>
```

### Service and Instance Common Names

The entire service or instance can define a common name to describe the purpose of the service/instance.

```
<template>
  <common_name>
    <loctext xml:lang='C'>console login</loctext>
  </common_name>
</template>
```

`common_name` is a free-form string, but is intended to be used as a label in a GUI or CLI.

Use the following guidelines when defining a common name:

- Be brief. A word or two is usually appropriate. Limit a name to under 40 characters.
- Be clear. The service, property group, or property name might not be helpful for humans, but `common_name` should help clarify the purpose of the entity.
- No punctuation. `common_name` is not a sentence or a paragraph. It should not contain clauses or phrases. Punctuation should only be present to meet trademark requirements.
- Capital letters must be used only for acronyms or proper names. For locales other than English, use appropriate capitalization for a sentence fragment.

### Service and Instance Descriptions

The `description` element contains a longer description of the property group, suitable for a status line or a tool-tip:

```
<template>
  <description>
    <loctext xml:lang='C'>Provide the text login prompt on console.
    </loctext>
  </description>
</template>
```

### *description* Guidelines

- Use proper grammar. *description* is a sentence meant to be read by humans.
- Be brief. A few sentences are usually most appropriate.

### Documentation

Documentation for this service can be defined explicitly, so that when the service is experiencing issues, or a consumer of the service wants more information on it, they can find it easily.

### Property Groups

The `pg_pattern` element contains the definitions for a property group:

```
<template>
  <pg_pattern name="pgname" type="pgtype" target="this" required="true">
    </pg_pattern>
</template>
```

*name* is the property group's name, and *type* is the property group's type.

*target* specifies what the target of this definition is. "this" would refer to the defining service or instance. "instance" can only be used in a service's template block, and means the definition applies to all instances of this service. "delegate" can only be used in a restarter's template block, and applies to all instances that are delegated to that restarter. "all", only usable by the master restarter, would refer to all services on the system. The default value of *target* is "this".

*required* indicates whether this property group is required or not. The default value of *required* is false. If *required* is true, both *name* and *type* must be specified.

*name* and/or *type* can be omitted. If either of these attributes is omitted it is treated as a wildcard. For instance, if the name attribute is omitted from the `pg_pattern` definition, the `pg_pattern` is applied to all property groups that have the specified type.

### Property Group Names

The `common_name` element contains the localized, human-readable name for the property group:

```
<pg_pattern ...>
  <common_name>
    <loctext xml:lang='C'>start method</loctext>
  </common_name>
</pg_pattern>
```

*common\_name* is a free-form string, but is intended to be used as a label in a GUI or CLI.

See the guidelines for *common\_name* under “Service Instance and Common Names,” above.

## Property Group Description

The *description* element contains a longer description of the property group, suitable for a status line or a tool-tip:

```
<pg_pattern ...>
  <description>
    <loctext xml:lang='C'>A required method which starts the service.
  </loctext>
  </description>
</pg_pattern>
```

See the guidelines for specifying a *description* under “Service and Instance Descriptions,” above.

## Properties

The *prop\_pattern* element contains the definitions for a specific property:

```
<pg_pattern ...>
  <prop_pattern name="proptype" type="proptype" required="true">
  </prop_pattern>
</pg_pattern>
```

*name* is the property's name, and *type* is the property's type.

*required* indicates whether this property is required. The default value of *required* is *false*.

*name* is always required. *type* is optional only if *required* is *false*.

## Property Names

The *common\_name* element contains the localized, human-readable name for the property:

*common\_name* is a free-form string field, but is intended to be used as a label in a GUI or CLI.

```
<prop_pattern ...>
  <common_name>
    <loctext xml:lang='C'>retry interval</loctext>
  </common_name>
</prop_pattern>
```

See the guidelines for *common\_name* under “Service Instance and Common Names,” above.

## Property units

The *units* element contains the localized, human-readable units for a numerical property:

```
<prop_pattern ...>
  <units>
    <loctext xml:lang='C'>seconds</loctext>
  </units>
</prop_pattern>
```

### *units* Guidelines

- Be brief. Strive to use only a single word or label. The plural form is usually the most appropriate.
- No punctuation. *units* is not a sentence or a paragraph. It should not contain clauses or phrases. Punctuation should be present only to meet trademark requirements.

### Property description

The *description* element contains a longer description of the property, suitable for a status line or a tool-tip:

```
<prop_pattern ...>
  <description> <loctext xml:lang='C'>
    The number of seconds to wait before retry.
  </loctext> </description>
</prop_pattern>
```

See the guidelines for specifying a *description* under “Service and Instance Descriptions,” above.

### Property visibility

The *visibility* element specifies whether simplified views in higher level software might want to display this property.

```
<prop_pattern ...>
  <visibility value="hidden | readonly | readwrite"/>
</prop_pattern>
```

Some properties are internal implementation details and should not be presented as a configuration setting. Others might merely be read-only. This property is used to specify these restrictions. A value of *hidden* indicates that the property shouldn't be displayed, *readonly* means that the property isn't intended to be modified, and *readwrite* indicates the property is modifiable.

This is not a security mechanism, it is solely intended to help prevent the user from shooting himself in the foot, and to remove unnecessary clutter from CLI output or a GUI display. Hidden properties is visible in full-disclosure modes of many commands and UIs.

### Property format

The *cardinality* and *internal\_separators* elements constrain the structure of a property:

```
<prop_pattern ...>
  <cardinality min="1" max="1"/>
  <internal_separators>,</internal_separators>
</prop_pattern>
```

*cardinality* indicates the acceptable number of property values. *min* is the minimum number, and *max* is the maximum number. Both are optional. If neither is specified, `<cardinality/>` is the same as the default, zero or more values.

*internal\_separators* specify the separator characters used for those property values into which multiple real values are packed.

### Value constraints

The *constraints* element specifies what values are acceptable for a property:

```
<prop_pattern ...>
<constraints>
  <value name="blue" />
  <range min="1" max="7"/>
  <include_values type="values"/>
</constraints>
</prop_pattern>
```

The *value* element includes a possible property value. *range* includes an integer range.

*value* and *range* can be used in any combination, as restricting their use would prohibit many valid descriptions. If no value constraints are specified, the property can take on any value.

*include\_values* includes all values specified by the values block (see Value Descriptions section).

### Value choices

The choices block indicates which values a UI should offer the user:

```
<prop_pattern ...>
<choices>
  <range min="1" max="3"/>
  <value name="vt100" />
  <value name="xterm" />
  <include_values type="constraints"/>
  <include_values type="values"/>
</choices>
</prop_pattern>
```

*range* and *value* include ranges and individual values as they do for constraints.

*include\_values* includes all values specified by either the constraints block or the values block (see next section).

### Value Descriptions

Like property names, the values a property can take on can also have inscrutable representations. The *values* element contains localized, human-readable descriptions for specific property values:

```

<prop_pattern>
<values>
  <value name="blue">
    <common_name>
      <loctext xml:lang='C'>blue</loctext>
    </common_name>
    <description>
      <loctext xml:lang='C'>
        The color between green and indigo.
      </loctext>
    </description>
  </value>
</values>
</prop_pattern>

```

*common\_name* is a free-form string field, but is intended to be used as a label in a GUI or CLI.

See the guidelines for *common\_name* under “Service Instance and Common Names,” above.

**Examples** Assuming a basic service which wants to define basic templates data looks like this:

```

<?xml version="1.0"?
<!DOCTYPE service_bundle SYSTEM "/usr/share/lib/xml/dtd/service_bundle.dtd.1">
<service_bundle type='manifest' name='FOOfoo:foo'>
<service name='system/foo' type='service' version='1'>
  <dependency>
    name='multi-user'
    type='service'
    grouping='require_all'
    restart_on='none'
    <service_fmri value='svc:/milestone/multi-user' />
  </dependency>
  <exec_method
    type='method'
    name='start'
    exec='/opt/foo/food'
    timeout_seconds='60'>
  </exec_method>
  <exec_method
    type='method'
    name='stop'
    exec=':kill'
    timeout_seconds='60'>
  </exec_method>
  <property_group name='config' type='application'>
    <propval name='local_only' type='boolean' value='false' />
    <propval name='config_file' type='astring'
      value='/opt/foo/foo.conf' />
  </property_group name='modules' type='astring'>

```

```
        <astring_list>
            <value_node value='bar' />
            <value_node value='baz' />
        </astring_list>
    </property>
</property_group>

    <instance name='default' enabled='false' />
</service>
</service_bundle>
```

That service could define some basic templates data to help an administrator using this service inside of the `<service>` tags. The most helpful things are to document the purpose of the service itself and the service-specific configuration.

```
<template>
    <common_name> <loctext xml:lang='C'>
        all-purpose demonstration
    </loctext> </common_name>
    <documentation>
        <manpage title='food' section='1M'
            manpath='/opt/foo/man' />
    </documentation>

    <pg_pattern name='config' type='application' target='this'
        required='true'>
        <description> <loctext xml:lang='C'>
            Basic configuration for foo.
        </loctext> </description>
        <prop_pattern name='local_only' type='boolean'
            required='false'>
            <description> <loctext xml:lang='C'>
                Only listen to local connection requests.
            </loctext> </description>
        </prop_pattern>
        <prop_pattern name='config_file' type='astring'
            required='true'>
            <cardinality min='1' max='1' />
            <description> <loctext xml:lang='C'>
                Configuration file for foo.
            </loctext> </description>
        </prop_pattern>
        <prop_pattern name='modules' type='astring'
            required='false'>
            <description> <loctext xml:lang='C'>
                Plugin modules for foo.
            </loctext> /description>
            <values>
                <value name='bar'>
```

```

        <description> <loctext xml:lang='C'>
            Allow foo to access the bar.
        </loctext> </description>
    </value>
    <value name='baz'>
        <description> <loctext xml:lang='C'>
            Allow foo to access baz functions.
        </loctext> </description>
    </value>
    <value name='qux'>
        <description> <loctext xml:lang='C'>
            Allow foo to access qux functions.
        </loctext> </description>
    </value>
</values>
<choices>
    <include_values type='values' />
</choices>
<prop_pattern>
</prop_pattern>
</template>

```

**Files** /usr/share/lib/xml/dtd/service\_bundle.dtd.1

**See Also** [svcs\(1\)](#), [svccfg\(1M\)](#), [libscf\(3LIB\)](#), [service\\_bundle\(4\)](#), [smf\(5\)](#)

/usr/apache2/2.2/man

<http://httpd.apache.org>

**Name** solaris10 – Solaris 10 branded zone

**Description** The `solaris10` brand uses the branded zones framework described in [brands\(5\)](#) to enable Solaris 10 binary applications to run unmodified on a machine with the latest Solaris Operating System kernel.

Oracle Solaris 10 Zones are `solaris10` branded zones that host x86 and SPARC Solaris 10 9/10 (or later released Oracle Solaris 10 update) user environments running on the Oracle Solaris 11 kernel.

Note that it is possible to use an earlier Oracle Solaris 10 release if you first install the kernel patch 142909-17 (SPARC) or 142910-17 (x86/x64), or later version, on the original system.

The `solaris10` brand includes the tools necessary to install a Solaris 10 system image into a non-global zone. It also supports the tools necessary to migrate a Solaris 10 native zone to a `solaris10` branded zone. The brand supports the execution of 32-bit and 64-bit Solaris 10 applications on either SPARC or x86 machines running the latest Solaris operating system.

**Configuration and Administration** The `solaris10` brand supports the whole root non-global zone model. All of the required Solaris 10 software and any additional packages are installed into the private file systems of the zone.

The zone must reside on its own [zfs\(1M\)](#) dataset and only ZFS is supported. The ZFS dataset created automatically when the zone is installed or attached. If a ZFS dataset cannot be created, the zone is not installed or attached.

The [zonecfg\(1M\)](#) utility is used to configure a `solaris10` branded zone. The `SYSsolaris10` template can be used when creating the zone or the configuration can be set up manually. Once a branded zone has been installed, that zone's brand cannot be changed or removed. The [zoneadm\(1M\)](#) utility is used to report the zone's brand type and administer the zone. The [zlogin\(1\)](#) utility is used to log in to the zone.

The support for delegated ZFS dataset configurations is currently experimental and has not yet been tested. Support for running these zones in a para-virtualized xVM domain is experimental and there are known problems with 64-bit x86 applications within the zone. The `/dev/sound` device cannot be configured into the branded zone. In addition, [mdb\(1\)](#) and [dtrace\(1M\)](#) are not fully functional when used in the global zone to examine processes executing within a `solaris10` branded zone.

The `solaris10` brand installer supports installing the zone from an image of an installed Solaris 10 system. This can be a full [flash\\_archive\(4\)](#), [cpio\(1\)](#), or [pax\(1\)](#) xustar archive. The `cpio` archive can be compressed with [gzip\(1\)](#) or [bzip2\(1\)](#). The image can also be a level 0 [ufsdump\(1M\)](#), or a path to the top-level of a Solaris 10 system's root directory tree. The zone cannot be installed from standard Solaris 10 distribution media.

To migrate a native zone from a Solaris 10 system to the latest Solaris Operating System kernel, the `attach` subcommand supports installing the zone from an archive of an installed Solaris 10 native zone. As with the installer, this can be a [cpio\(1\)](#) or [pax\(1\)](#) xustar archive of

the `zonepath`. The `cpio` archive can be compressed with `gzip(1)` or `bzip2(1)`. The image can also be a path to the top-level of a Solaris 10 zone's `zonepath` directory tree. In addition to migrating from a Solaris 10 native zone, the same migration options can be used when migrating a `solaris10` branded zone from one host to another. When migrating from Solaris 10, it is possible that the zone is configured as a `sparseroot` zone. In this case, the zone should be readied on the host before the archive is made. This ensures that the inherited directories are included in the archive.

**Sub-commands** The following arguments of `zoneadm(1M)` brand-specific subcommand are supported:

`attach [-a archive | -d path] [-c sysidcfg]`

Attach the specified Solaris 10 native zone image into the branded zone. If neither `-a` or `-d` is specified, the zone's `zonepath` is assumed to already be properly installed with the zone's files.

`-a archive` The path to a `cpio(1)`, `pax(1)` `xustar`, or `zfs` archive of either an installed Solaris 10 native zone or a `solaris10` branded zone's `zonepath`. `cpio` and `zfs` archives can be compressed using `gzip` or `bzip2`.

This option may not be used if the zone's state is 'unavailable'.

Warning: This option may be removed in a future release of Solaris. Use of `'zoneadm install -a archive'` is suggested.

`-c sysidcfg` Specifying a `sysidcfg` file causes a `sys-unconfig` to occur on the zone after `attach`. The `sysidcfg` file is then applied to the zone.

`-d path` The path to the `zonepath` directory of either an installed Solaris 10 native zone or a `solaris10` branded zone's `zonepath`.

This option may not be used if the zone's state is 'unavailable'.

Warning: This option may be removed in a future release of Solaris. Use of `'zoneadm install -a archive'` is suggested.

`clone [-c sysidcfg]`

Install a zone by copying an existing installed zone. This subcommand is an alternative way to install the zone.

`-c sysidcfg` Provides a `sysidcfg` file to apply after unconfiguration of the cloned zone.

`install [-a archive] [-d path] [-p] [-s] [-u] [-v] \`  
`[-c sysidcfg]`

Install the specified Solaris 10 system image into the zone. Either the `-u` or `-p` option is required *and* either the `-a` or `-d` option is required.

- a *archive*  
The path to a `flash_archive(4)`, `cpio(1)`, `pax(1)` xustar, zfs archive, or a level 0 `ufsdump(1M)` of an installed Solaris 10 system, an installed Solaris 10 native zone, or a solaris10 branded zone. The `cpio` and `zfs` archives can be compressed using `gzip` or `bzip2`.
- c *sysidcfg*  
Provides a *sysidcfg* file to apply after installation.
- d *path*  
The path to the root directory of an installed Solaris 10 system.
- p  
Preserve the system configuration after installing the zone.
- s  
Install silently.
- u  
Run `sys-unconfig` on the zone after installing it.
- v  
Verbose output from the install process.

**Application Support** The `solaris10` zone only supports user-level Solaris 10 applications. You cannot use Solaris 10 device drivers or Solaris 10 kernel modules from inside a `solaris10` zone. However, depending on the kernel module, you might be able to use the latest Solaris kernel module version with the Solaris 10 user-level application.

#### **Examples** EXAMPLE 1 Creating a ZFS Flash Archive for Install

The following example shows how to create an archive for a physical to virtual (P2V) migration. This is performed in the global zone of a system that is running Solaris 10. The Solaris 10 system must not have any non-global zones configured, installed, or running. The Solaris 10 system can use ZFS or UFS as its root file system.

```
# flarcreate -n s10box -c /net/somehost/p2v/s10box.flar
```

#### EXAMPLE 2 Installing a solaris10 Branded Zone Using a Flash Archive

The following example installs a zone using the archive from Example 1. It assumes the zone has already been configured with `zonecfg(1M)` and has the brand property set to `solaris10`.

```
# zoneadm -z s10p2v install -a /net/somehost/p2v/s10box.flar -p
```

#### EXAMPLE 3 Creating a ZFS Archive for Attach

The following example shows how to create an archive for a virtual to virtual (V2V) migration. It assumes that the `zonepath` for the `solaris10` branded zone is `/zones/v2vzone`.

First, determine the name of `zonepath` dataset.

**EXAMPLE 3** Creating a ZFS Archive for Attach *(Continued)*

```
# dataset=$(zfs list -H -o name /zones/v2vzone)
```

Next, create a snapshot of the zone's datasets.

```
# zfs snapshot -r $dataset@v2v
```

Finally, generate a ZFS self-contained recursive stream that is compressed with bzip2.

```
# zfs send -rc $dataset@v2v | bzip2 > /net/somehost/v2v/v2v.zfs.bz2
```

**EXAMPLE 4** Attaching a Zone Using a ZFS Archive

The following example attaches a zone using a ZFS archive. It assumes that the zone has already been configured using [zonecfg\(1M\)](#) and that the brand property is set to `solaris10`.

```
# zoneadm -z v2vzone attach -a /net/somehost/v2v/v2v.zfs.bz2
```

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/zones/brand/brand-solaris10
Interface Stability	Committed

**See Also** [cpio\(1\)](#), [mdb\(1\)](#), [pax\(1\)](#), [zlogin\(1\)](#), [dtrace\(1M\)](#), [ufsdump\(1M\)](#), [zfs\(1M\)](#), [zoneadm\(1M\)](#), [zonecfg\(1M\)](#), [flash\\_archive\(4\)](#), [attributes\(5\)](#), [brands\(5\)](#), [zones\(5\)](#)

**Name** solaris, ipkg – solaris branded zone

**Description** The `solaris` brand uses the branded zones framework described in [brands\(5\)](#) to run zones installed with the same software as is installed in the global zone. The system software must always be in sync with the global zone when using a `solaris` brand. The system software packages within the zone are managed using the image packaging system. See [pkg\(5\)](#).

**Configuration and Administration** The `solaris` brand supports the whole root non-global zone model. All of the required system software and any additional packages are installed into the private file systems of the zone. The zone must reside on its own [zfs\(1M\)](#) dataset and only ZFS is supported. The ZFS dataset is created automatically when the zone is installed or attached. If a ZFS dataset cannot be created, the zone is not installed or attached.

**Sub-commands** The following `solaris` brand-specific subcommand options are supported by [zoneadm\(1M\)](#).

`attach [-a archive | -d path | -z ZBE] [-u | -U] [-c config_profile.xml | dir]`

Attach the specified `solaris` branded zone image into the zone. If neither `-a` or `-d` is specified, the zone's `zonpath` is assumed to already be properly installed with the zone's files. `zoneadm` checks package levels on the machine to which the zone is to be attached. If the packages that the zone depends on from the global zone are different (have different revision numbers) from the dependent packages on the source machine, `zoneadm` reports these conflicts and does not perform the attach.

If the destination system has only newer dependent packages (higher revision numbers) than those on the source system, you can use the `-u` or `-U` option to update the dependent packages to match the revision of the packages that exist on the new system.

When attaching a zone, multiple zone boot environments (ZBEs) can exist and the `attach` subcommand must determine which one to attach. The selection criteria is as follows, with the first match being used.

- If the `-z` option is used to specify a ZBE, it is selected.
- If there is only one ZBE, it is selected.
- If there is only one ZBE associated with this global zone boot environment, it is selected.
- If there is only one active ZBE associated with this global zone boot environment, it is selected.
- If there is only one active ZBE, it is selected.

If the selected ZBE is associated with another global zone boot environment the ZBE is cloned and the clone of the selected ZBE is attached. See [beadm\(1M\)](#) for more information regarding boot environments.

`-a archive`

Warning: This option may be removed in a future release of Solaris. Use of `'zoneadm install -a archive'` is suggested.

The path to a `cpio(1)` or `pax(1)` xustar or `zfs(1M)` archive of an installed Oracle Solaris branded zone.

`cpio` and `pax` archives must be of the `zonelibrary` or the `zonelibrary` and must not contain absolute paths (paths must not start with `/`). ZFS archives can be of the `zonelibrary` dataset, `zone rpool` dataset or a single zone boot environment.

If a ZFS archive contains more than one ZBE, the attach can fail with a message indicating that a specific ZBE must be attached using the `-z` option. In such a case, the ZBEs from the archive are remain extracted and the `-a` and `-d` options are invalid for further use with this zone until the `-z` option is used to attach one of the extracted ZBEs. If, instead of attaching one of the extracted ZBEs, it is desired to delete the extracted ZBEs, use `zoneadm -z <zone> mark incomplete` followed by `zoneadm -z <zone> uninstall`.

`cpio` and ZFS archives can be compressed using `gzip` or `bzip2`.

This option may not be used if the zone's state is 'unavailable'.

`-c config_profile.xml | dir`

Provides a profile or a directory of profiles to apply after installation from the repository.

All profiles must have an `.xml` extension.

`-d path`

Warning: This option may be removed in a future release of Solaris. Use of `'zoneadm install -a archive'` is suggested.

The path to the `zonelibrary` directory of a solaris branded zone's `zonelibrary`.

This option may not be used if the zone's state is 'unavailable'.

`-u`

Update the minimal number of packages within the zone to allow the zone's packages to be compatible with the packages installed in the global zone.

`-U`

Update all packages within the zone to their latest versions which are compatible with the packages installed in the global zone.

`-z ZBE`

Attach the specified existing zone boot environment. If the specified zone boot environment is associated with a different global zone, the specified ZBE is cloned and a clone of the ZBE is attached.

`clone [-c config_profile.xml | dir]`

`-c config_profile.xml | dir`

Provides a profile or a directory of profiles to apply after installation from the repository.

All profiles must have an `.xml` extension.

```
install [-m manifest.xml] [-c config_profile.xml | dir]
install [-a archive | -d path] [-U] [-p] [-s] [-u] [-v] [-c config_profile.xml | dir]
```

The solaris brand installer supports installing the zone from either the software repository or from an image of an installed system running the same release. This can be a [cpio\(1\)](#), [pax\(1\)](#) xustar, or ZFS archive. The cpio or ZFS archive can be compressed with gzip or bzip2. The image can also be a path to the top-level of a system's root tree, or a pre-existing zone path.

If neither the -a nor -d options are specified, the zone is installed from the repository. To install additional packages in a zone the default zone manifest, `/usr/share/auto_install/manifest/zone_default.xml`, can be copied and edited to include the needed packages. This modified manifest should be specified to install with the -m option.

To install the zone from a system or zone image, either the -a or -d options is required. If required, the software in the image's active ZBE will be updated with the minimal changes required to make it compatible with the global zone's packages. If the -U option is specified, all software in the image's active ZBE will be updated to the latest version compatible with the active ZBE. If either the -a or -d options is used, either the -u or -p option is also required.

**-a *archive*** The path to a [cpio\(1\)](#) or [pax\(1\)](#) xustar or ZFS archive of an installed global zone or non-global zone.

If a ZFS archive contains multiple boot environments, the active boot environment are installed. If install is unable to determine which boot environment is the active boot environment, install provides a list of boot environments extracted and suggest an attach command that uses the -z option to attach a specific boot environment.

cpio and ZFS archives can be compressed using gzip or bzip2.

**-c *config\_profile.xml* | *dir*** Provides a profile or a directory of profiles to apply after installation from the repository.

All profiles must have an .xml extension.

**-d *path*** The path to the zonepath directory of a solaris branded zone's zonepath or Solaris 11 global zone root directory.

**-m *manifest.xml*** Manifest file to be specified to the automated installer.

**-p** Preserve the system configuration after installing the zone from an archive or a path.

**-s** Install silently

---

-u	Unconfigure the system after installing it.
-U	Update all packages within the zone to their latest versions which are compatible with the packages installed in the global zone. The -U option may only be used if either of the -a or -d options is used.
-v	Verbose output from the install process.
-z <i>ZBE</i>	Attach the specified existing zone boot environment. If the specified zone boot environment is associated with a different global zone, the specified ZBE is cloned and a clone of the ZBE is attached.

**Examples** EXAMPLE 1 Creating a ZFS archive for P2V

The following example shows how to create an archive for a physical to virtual (P2V) migration. This is performed in the global zone of a system that has no non-global zones, configured, installed, or running. It assumes the root pool is named `rpool`.

First, create a snapshot of the entire root pool.

```
# zfs snapshot -r rpool@p2v
```

Next, destroy the snapshots associated with swap and dump devices, as there is no need for them on the target system.

```
# zfs destroy rpool/swap@p2v
# zfs destroy rpool/dump@p2v
```

Finally, generate a ZFS replication stream archive that is compressed with `gzip`. In this example, it is stored on a remote NFS server.

```
# zfs send -R rpool@p2v | gzip > /net/somehost/p2v/p2v.zfs.gz
```

EXAMPLE 2 Create a ZFS Archive for V2V

The following example shows how to create an archive for a virtual to virtual (V2V) migration. It assumes that the `zonename` for the zone is `/zones/v2vzone`.

First, determine the name of `zonename` dataset.

```
# dataset=$(zfs list -H -o name /zones/v2vzone)
```

Next, create a snapshot of the zone's datasets.

```
# zfs snapshot -r $dataset@v2v
```

Finally, generate a ZFS self-contained recursive stream that is compressed with `bzip2`.

```
# zfs send -rc $dataset@v2v | bzip2 > /net/somehost/v2v/v2v.zfs.bz2
```

**EXAMPLE 3** Installing a Zone Using a ZFS Archive

The following example installs a zone using a ZFS archive.

```
# zoneadm -z v2vzone install -p -a /net/somehost/v2v/v2v.zfs.bz2
```

A similar command can be used for installing the P2V archive created in Example 1.

**Attributes** See [attributes\(5\)](#) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/zones
Interface Stability	Uncommitted

**See Also** [cpio\(1\)](#), [pax\(1\)](#), [beadm\(1M\)](#), [sysconfig\(1M\)](#), [zfs\(1M\)](#), [zlogin\(1\)](#), [zonename\(1\)](#), [zoneadm\(1M\)](#), [zonectg\(1M\)](#), [attributes\(5\)](#), [brands\(5\)](#), [privileges\(5\)](#), [zones\(5\)](#)

[pkg\(5\)](#), available in the IPS consolidation

**Name** standards, ANSI, C, C++, ISO, POSIX, POSIX.1, POSIX.2, SUS, SUSv2, SUSv3, SVID, SVID3, XNS, XNS4, XNS5, XPG, XPG3, XPG4, XPG4v2 – standards and specifications supported by Oracle Solaris

**Description** Oracle Solaris 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
POSIX.1-2001	POSIX.1-1990, POSIX.1b-1993, POSIX.1c-1996, POSIX.2-1992, and POSIX.2a-1992 updates	Solaris 10
POSIX.1-2004	POSIX.1-2001 update	Oracle Solaris 11

Oracle Solaris 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); Single UNIX Specification, Version 2 (SUSv2); and Single UNIX Specification, Version 3 (SUSv3). 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

X/Open CAE		
Specification	Description	Release
SUSv2	superset of SUS extended to support POSIX.1b-1993, POSIX.1c-1996, and ISO/IEC 9899 (C Standard) Amendment 1	Solaris 7
XNS5	superset and LP64-clean derivative of XNS4.	Solaris 7
SUSv3	same as POSIX.1–2001	Solaris 10

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 or SUSv3, which have LP64-clean interfaces that are portable across ILP32 and LP64 environments. Solaris releases 7 through Oracle Solaris 11 support both the ILP32 and LP64 environments.

Solaris releases 7 through 10 have been branded to conform to The Open Group's UNIX 98 Product Standard. Solaris 10 through Oracle Solaris 11 have been branded to conform to The Open Group's UNIX 03 Product Standard.

Solaris releases 2.0 through Oracle Solaris 11 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 Oracle Solaris Studio 12.3 C Compiler is installed, Oracle Solaris 11 supports the ANSI X3.159-1989 Programming Language - C and ISO/IEC 9899:1990 Programming Language - C (C) interfaces.

When Oracle Solaris Studio 12.3 C Compiler is installed, Oracle Solaris 11 supports ISO/IEC 9899:1990 Amendment 1:1995: C Integrity.

When Oracle Solaris Studio 12.3 C Compiler is installed, Oracle Solaris 11 supports ISO/IEC 9899:1999 Programming Languages – C.

When Oracle Solaris Studio 12.3 C++ Compiler is installed, Oracle Solaris 11 supports ISO/IEC 14882:1998 Programming Languages - C++. Unsupported features of that standard are described in the compiler README file.

**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`. If the behavior required by POSIX.1–2001 or SUSv3 conflicts with historical Solaris utility behavior, a new version that is standard-conforming has been provided in `/usr/xpg4/bin` or in `/usr/xpg6/bin`. If the behavior required by POSIX.1–2001 or SUSv3 conflicts with POSIX.2, POSIX.2a, SUS, or SUSv2, a new version that is SUSv3 standard-conforming has been provided in `/usr/xpg6/bin`.

An application that wants to use standard-conforming utilities must set the `PATH` ([sh\(1\)](#) or [ksh\(1\)](#)) or `path` ([csh\(1\)](#)) environment variable to specify the directories listed below in the order specified to get the appropriate utilities:

#### SVID3, XPG3

1. `/usr/bin`
2. directory containing binaries for your compiler
3. other directories containing binaries needed by the application

#### POSIX.2, POSIX.2a, SUS, SUSv2, XPG4

1. `/usr/xpg4/bin`
2. `/usr/bin`
3. directory containing binaries for your compiler
4. other directories containing binaries needed by the application

#### POSIX.1–2001, SUSv3

1. `/usr/xpg6/bin`
2. `/usr/xpg4/bin`
3. `/usr/bin`
4. directory containing binaries for your compiler
5. other directories containing binaries needed by the application

When an application uses `execvp()` or `execvp()` (see [exec\(2\)](#)) to execute a shell file, or uses [system\(3C\)](#), the shell used to interpret the shell file depends on the standard to which the caller conforms:

Standard	Shell Used
1989 ANSI C, 1990 ISO C, 1999 ISO C, POSIX.1 (1990–2001), SUS, SUSv2, SUSv3, XPG4	<code>/usr/xpg4/bin/sh</code>
POSIX.1 (1988), SVID3, XPG3, no standard specified	<code>/usr/bin/sh</code>

**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 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 on the compile command line or within the application source files.

#### 1989 ANSI C, 1990 ISO C, 1999 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`, predefined by the compiler based on compiler defaults and command-line options, is set to a value of 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.)

C++ bindings are not defined for POSIX or X/Open CAE, so specifying feature test macros such as `_POSIX_SOURCE`, `_POSIX_C_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>
POSIX.1-2001	<code>_POSIX_C_SOURCE=200112L</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 `_XOPEN_SOURCE`. If `_XOPEN_SOURCE` is defined with a value, the value must be less than 500.

XPG4	The application must define <code>_XOPEN_SOURCE</code> and set <code>_XOPEN_VERSION=4</code> . If <code>_XOPEN_SOURCE</code> is defined with a value, the value must be less than 500.
SUS (XPG4v2)	The application must define <code>_XOPEN_SOURCE</code> and set <code>_XOPEN_SOURCE_EXTENDED=1</code> . If <code>_XOPEN_SOURCE</code> is defined with a value, the value must be less than 500.
SUSv2	The application must define <code>_XOPEN_SOURCE=500</code> .
SUSv3	The application must define <code>_XOPEN_SOURCE=600</code> .

Compilation The Oracle Solaris Studio 12.3 C Compiler provides the ISO/IEC 99899:1999 (1999 ISO C Language) standard-conforming compilation system and the `c99` utility.

When `ld` is used directly to link applications, `/usr/lib/values-xpg4.o` must be specified on any link/load command line, unless the application is POSIX.1-2001- or SUSv3-conforming, in which case `/usr/lib/values-xpg6.o` must be specified on any link/load compile line. When `cc` or `CC` is used to link applications, the compiler automatically adds the appropriate file. The preferred way to build applications, however, is described in the table 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 Oracle Solaris Studio 12.3 C Compiler and the Oracle Solaris Studio 12.3 C++ Compiler define the macro `__PRAGMA_REDEFINE_EXTNAME` to indicate that they support 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.

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 the Oracle Solaris Studio 12.3 C Compiler, 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
1989 ANSI C and 1990 ISO C	<code>c89</code>	none
1999 ISO C	<code>c99</code>	none
SVID3	<code>cc -Xt -xc99=none</code>	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
POSIX.1-2001	c99	_POSIX_C_SOURCE=200112L
POSIX.1c-1996	c89	_POSIX_C_SOURCE=199506L
CAE XPG3	cc -Xa -xc99=none	_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
SUSv3	c99	_XOPEN_SOURCE=600

For platforms supporting the LP64 (64-bit) programming environment, 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
```

Similar SUSv3–conforming LP64 applications should be built with command lines of the form:

```
c99 $(getconf POSIX_V6_LP64_OFF64_CFLAGS) -D_XOPEN_SOURCE=600 \  
    $(getconf POSIX_V6_LP64_OFF64_LDFLAGS) foo.c -o foo \  
    $(getconf POSIX_V6_LP64_OFF64_LIBS) -lxnet
```

### SUSv3

```
c99 _XOPEN_SOURCE=600
```

**See Also** [csh\(1\)](#), [ksh\(1\)](#), [sh\(1\)](#), [exec\(2\)](#), [sysconf\(3C\)](#), [system\(3C\)](#), [environ\(5\)](#), [lf64\(5\)](#)

**Name** sticky – mark files for special treatment

**Description** The *sticky bit* (file mode bit `01000`, see [chmod\(2\)](#)) 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 can only be removed or renamed by a user who has write permission on the directory, and either owns the file, owns the directory, has write permission on the file, or is a privileged user. Setting the sticky bit is useful for directories such as `/tmp`, which must be publicly writable but should deny users permission to arbitrarily delete or rename the files of others.

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.

Any user may create a sticky directory. See [chmod](#) for details about modifying file modes.

**See Also** [chmod\(1\)](#), [chmod\(2\)](#), [chown\(2\)](#), [mkdir\(2\)](#), [rename\(2\)](#), [unLink\(2\)](#)

**Bugs** The [mkdir\(2\)](#) function will not create a directory with the sticky bit set.

**Name** suri – storage URIs

**Description** Storage URIs uniquely identify storage objects across different nodes.

The following Storage URIs are supported:

Dev URI `dev:<local-path-under-/dev>`  
`dev:///<path-with-dev>`  
`dev:<absolute-path-with-dev>`

Examples:

```
dev:dsk/c0t0d0s0
dev:///dev/dsk/c0t0d0s0
dev:/dev/dsk/c0t0d0s0
```

The dev URI specifies an object in the form of an existing device. All forms of the dev URI are functionally equivalent. The “dev:dsk” format with no authority section and an implied “/dev/” prefix is usually preferred. The path must be a device in the /dev name space.

Components of the device path may only include characters from POSIX.1-2008 Portable Filename Character Set, 0-9a-zA-Z.\_-.

Logical Unit URI `lu:luname.naa.<ID>`  
`lu:initiator.naa.<ID>,target.naa.<ID>,luname.naa.<ID>`

Examples:

```
lu:luname.naa.5000c5000288fa25
lu:initiator.naa.2100001d38089fb0,target.naa.2100001d38089fb0,luname.naa.5000c5000288fa25
```

Logical unit URI specifies a logical unit attached via fibre channel or serial-attached SCSI. In the luname-only URI form, the ID describes a logical unit name. In the initiator,target,luname form, an initiator specifies an initiator port and a target specifies a target port, and together they specify a path to the logical unit whose name must match the URI's luname part as in the luname-only form.

IDs are represented by using the SCSI-3 VPD page (0x83) global unique identifier (GUID) formatted as hexadecimal numbers (64/128 bits). IDs are case insensitive. Only GUIDs based on NAA identifiers are supported. See NOTES for information on NAA identifier format.

In general it is recommended using multipathing in conjunction with luname-only URIs. If multipathing is disabled and a luname-only URI is used, a random path to the specified logical unit will be chosen. To avoid this random path assignment an initiator,target,luname URI form can be used to select a specific path to a logical unit. If multipathing is enabled and an initiator,target,luname URI form is used then the multipathing framework controls which paths are used to access the logical unit and the URI is only used to identify the unit, not the access path.

If an initiator or/and a target IDs are not matched but the logical unit is found based on a logical unit name, the unit will be accessed via an alternate path.

iSCSI URI `iscsi:///lunname.naa.<ID>`  
`iscsi://<host>[:<port>]/lunname.naa.<ID>`

Examples:

```
iscsi:///lunname.naa.600144f03d70c80000004ea57da10001
iscsi://[::1]/lunname.naa.600144f03d70c80000004ea57da10001
iscsi://127.0.0.1/lunname.naa.600144f03d70c80000004ea57da10001
iscsi://127.0.0.1:3620/lunname.naa.600144f03d70c80000004ea57da10001
iscsi://hostname:3620/lunname.naa.600144f03d70c80000004ea57da10001
```

The iSCSI URI specifies an object based on a logical unit number accessible via iSCSI protocol. Optional `hostname[:<port>]` authority section provides information for a consumer to set up a send-targets discovery address. IPv6 addresses must be enclosed in square brackets.

Format of an ID is explained in the “Logical Unit URI” section. See NOTES for information on NAA identifier format.

**Getting URIs** To avoid errors by generating URIs manually, the `suriadm(1M)` subcommand `lookup-uri` is recommended to be used to generate URIs based on existing system device paths. Such device paths must be identified via other means. For Fibre Channel devices, `fcadm(1M)` may be used. For Serial Attached SCSI devices, `sasinfo(1M)` may be used. For iSCSI devices, `iscsiadm(1M)` may be used.

**Attributes** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/library/storage/suri
Interface Stability	Committed

**See Also** `fcadm(1M)`, `iscsiadm(1M)`, `sasinfo(1M)`, `suriadm(1M)`, `attributes(5)`

Small Computer System Interface-3 (SCSI-3)

**Notes** NAA stands for “Name Address Authority” and it is a field within an identifier that specifies the format and length of that identifier. NAA identifier format is defined in SCSI Primary Commands - 3 specification, section “7.6.3.6 NAA identifier format”. WWN (World Wide Name) identifiers use an NAA identifier format.

**Name** tecla, teclarc – User interface provided by the tecla library.

**Description** This man page describes the command-line editing features that are available to users of programs that read keyboard input via the tecla library. Users of the `tcsh` shell will find the default key bindings very familiar. Users of the `bash` shell will also find it quite familiar, but with a few minor differences, most notably in how forward and backward searches through the list of historical commands are performed. There are two major editing modes, one with emacs-like key bindings and another with vi-like key bindings. By default emacs mode is enabled, but `vi(1)` mode can alternatively be selected via the user's configuration file. This file can also be used to change the bindings of individual keys to suit the user's preferences. By default, tab completion is provided. If the application hasn't reconfigured this to complete other types of symbols, then tab completion completes file names.

**Key Sequence Notation** In the rest of this man page, and also in all tecla configuration files, key sequences are expressed as follows.

- `^A` or `C-a` This is a 'CONTROL-A', entered by pressing the CONTROL key at the same time as the 'A' key.
- `\E` or `M-` In key sequences, both of these notations can be entered either by pressing the ESCAPE key, then the following key, or by pressing the META key at the same time as the following key. Thus the key sequence `M-p` can be typed in two ways, by pressing the ESCAPE key, followed by pressing 'P', or by pressing the META key at the same time as 'P'.
- `up` This refers to the up-arrow key.
- `down` This refers to the down-arrow key.
- `left` This refers to the left-arrow key.
- `right` This refers to the right-arrow key.
- `a` This is just a normal 'A' key.

**The Tecla Configuration File** By default, tecla looks for a file called `.teclarc` in your home directory (ie. `~/ .teclarc`). If it finds this file, it reads it, interpreting each line as defining a new key binding or an editing configuration option. Since the emacs key-bindings are installed by default, if you want to use the non-default vi editing mode, the most important item to go in this file is the following line:

```
edit-mode vi
```

This will re-configure the default bindings for vi-mode. The complete set of arguments that this command accepts are:

- `vi` Install key bindings like those of the vi editor.
- `emacs` Install key bindings like those of the emacs editor. This is the default.
- `none` Use just the native line editing facilities provided by the terminal driver.

To prevent the terminal bell from being rung, such as when an unrecognized control-sequence is typed, place the following line in the configuration file:

```
nobeeep
```

An example of a key binding line in the configuration file is the following.

```
bind M-[2~ insert-mode
```

On many keyboards, the above key sequence is generated when one presses the insert key, so with this key binding, one can toggle between the emacs-mode insert and overwrite modes by hitting one key. One could also do it by typing out the above sequence of characters one by one. As explained above, the M- part of this sequence can be typed either by pressing the ESCAPE key before the following key, or by pressing the META key at the same time as the following key. Thus if you had set the above key binding, and the insert key on your keyboard didn't generate the above key sequence, you could still type it in either of the following 2 ways.

1. Hit the ESCAPE key momentarily, then press '[', then '2', then finally '~'.
2. Press the META key at the same time as pressing the '[' key, then press '2', then '~'.

If you set a key binding for a key sequence that is already bound to a function, the new binding overrides the old one. If in the new binding you omit the name of the new function to bind to the key sequence, the original binding becomes undefined.

Starting with versions of `libtecla` later than 1.3.3 it is now possible to bind key sequences that begin with a printable character. Previously key sequences were required to start with a CONTROL or META character.

Note that the special keywords "up", "down", "left", and "right" refer to the arrow keys, and are thus not treated as key sequences. So, for example, to rebind the up and down arrow keys to use the history search mechanism instead of the simple history recall method, you could place the following in your configuration file:

```
bind up history-search-backwards
bind down history-search-backwards
```

To unbind an existing binding, you can do this with the bind command by omitting to name any action to rebind the key sequence to. For example, by not specifying an action function, the following command unbinds the default beginning-of-line action from the ^A key sequence:

```
bind ^A
```

If you create a `~/ .teclarc` configuration file, but it appears to have no effect on the program, check the documentation of the program to see if the author chose a different name for this file.

**Filename and Tilde Completion** With the default key bindings, pressing the TAB key (aka. `^I`) results in tecla attempting to complete the incomplete file name that precedes the cursor. Tecla searches backwards from the cursor, looking for the start of the file name, stopping when it hits either a space or the start of the line. If more than one file has the specified prefix, then tecla completes the file name up to the point at which the ambiguous matches start to differ, then lists the possible matches.

In addition to literally written file names, tecla can complete files that start with `~/` and `~user/` expressions and that contain `$envvar` expressions. In particular, if you hit TAB within an incomplete `~user,` expression, tecla will attempt to complete the username, listing any ambiguous matches.

The completion binding is implemented using the `cpl_complete_word()` function, which is also available separately to users of this library. See the [cpl\\_complete\\_word\(3TECLA\)](#) man page for more details.

**Filename Expansion** With the default key bindings, pressing `^X*` causes tecla to expand the file name that precedes the cursor, replacing `~/` and `~user/` expressions with the corresponding home directories, and replacing `$envvar` expressions with the value of the specified environment variable, then if there are any wildcards, replacing the so far expanded file name with a space-separated list of the files which match the wild cards.

The expansion binding is implemented using the `ef_expand_file()` function. See the [ef\\_expand\\_file\(3TECLA\)](#) man page for more details.

**Recalling Previously Typed Lines** Every time that a new line is entered by the user, it is appended to a list of historical input lines maintained within the `GetLine` resource object. You can traverse up and down this list using the up and down arrow keys. Alternatively, you can do the same with the `^P`, and `^N` keys, and in `vi` command mode you can alternatively use the `k` and `j` characters. Thus pressing up-arrow once, replaces the current input line with the previously entered line. Pressing up-arrow again, replaces this with the line that was entered before it, etc.. Having gone back one or more lines into the history list, one can return to newer lines by pressing down-arrow one or more times. If you do this sufficient times, you will return to the original line that you were entering when you first hit up-arrow.

Note that in `vi` mode, all of the history recall functions switch the library into command mode.

In emacs mode the `M-p` and `M-n` keys work just like the `^P` and `^N` keys, except that they skip all but those historical lines which share the prefix that precedes the cursor. In `vi` command mode the upper case `'K'` and `'J'` characters do the same thing, except that the string that they search for includes the character under the cursor as well as what precedes it.

Thus for example, suppose that you were in emacs mode, and you had just entered the following list of commands in the order shown:

```
ls ~/tecla/
cd ~/tecla
ls -l getline.c
emacs ~/tecla/getline.c
```

If you next typed:

```
ls
```

and then hit M-p, then rather than returning the previously typed emacs line, which doesn't start with "ls", tecla would recall the "ls -l getline.c" line. Pressing M-p again would recall the "ls ~/tecla/" line.

Note that if the string that you are searching for, contains any of the special characters, \*, ?, or '[', then it is interpreted as a pattern to be matched. Thus, continuing with the above example, after typing in the list of commands shown, if you then typed:

```
*tecla*
```

and hit M-p, then the "emacs ~/tecla/getline.c" line would be recalled first, since it contains the word tecla somewhere in the line. Similarly, hitting M-p again, would recall the "ls ~/tecla/" line, and hitting it once more would recall the "ls ~/tecla/" line. The pattern syntax is the same as that described for file name expansion, in the [ef\\_expand\\_file\(3TECLA\)](#).

**History Files** Authors of programs that use the tecla library have the option of saving historical command-lines in a file before exiting, and subsequently reading them back in from this file when the program is next started. There is no standard name for this file, since it makes sense for each application to use its own history file, so that commands from different applications don't get mixed up.

**International Character Sets** Since `libtecla` version 1.4.0, tecla has been 8-bit clean. This means that all 8-bit characters that are printable in the user's current locale are now displayed verbatim and included in the returned input line. Assuming that the calling program correctly contains a call like the following,

```
setlocale(LC_CTYPE, "");
```

then the current locale is determined by the first of the environment variables `LC_CTYPE`, `LC_ALL`, and `LANG`, that is found to contain a valid locale name. If none of these variables are defined, or the program neglects to call `setlocale`, then the default C locale is used, which is US 7-bit ASCII. On most unix-like platforms, you can get a list of valid locales by typing the command:

```
locale -a
```

at the shell prompt.

### Meta Keys and Locales

Beware that in most locales other than the default C locale, META characters become printable, and they are then no longer considered to match M-c style key bindings. This allows international characters to be entered with the compose key without unexpectedly triggering META key bindings. You can still invoke META bindings, since there are actually two ways to do this. For example the binding M-c can also be invoked by pressing the ESCAPE key momentarily, then pressing the c key, and this will work regardless of locale. Moreover, many modern terminal emulators, such as gnome's gnome-terminal's and KDE's konsole terminals, already generate escape pairs like this when you use the META key, rather than a real meta character, and other emulators usually have a way to request this behavior, so you can continue to use the META key on most systems.

For example, although xterm terminal emulators generate real 8-bit meta characters by default when you use the META key, they can be configured to output the equivalent escape pair by setting their `EightBitInput` X resource to `False`. You can either do this by placing a line like the following in your `~/.Xdefaults` file,

```
XTerm*EightBitInput: False
```

or by starting an xterm with an `-xrm '*EightBitInput: False'` command-line argument. In recent versions of xterm you can toggle this feature on and off with the 'Meta Sends Escape' option in the menu that is displayed when you press the left mouse button and the CONTROL key within an xterm window. In CDE, dtterms can be similarly coerced to generate escape pairs in place of meta characters, by setting the `Dtterm*KshMode` resource to `True`.

### Entering International Characters

If you don't have a keyboard that generates all of the international characters that you need, there is usually a compose key that will allow you to enter special characters, or a way to create one. For example, under X windows on unix-like systems, if your keyboard doesn't have a compose key, you can designate a redundant key to serve this purpose with the `xmodmap` command. For example, on many PC keyboards there is a microsoft-windows key, which is otherwise useless under Linux. On a laptop, for example, the `xev` program might report that pressing this key generates keycode 115. To turn this key into a COMPOSE key, do the following:

```
xmodmap -e 'keycode 115 = Multi_key'
```

Type this key followed by a " character to enter an 'I' with a umlaut over it.

#### The Available Key Binding Functions

The following is a list of the editing functions provided by the tecla library. The names in the leftmost column of the list can be used in configuration files to specify which function a given key or combination of keys should invoke. They are also used in the next two sections to list the default key bindings in emacs and vi modes.

user-interrupt	Send a SIGINT signal to the parent process.
suspend	Suspend the parent process.

---

stop-output	Pause terminal output.
start-output	Resume paused terminal output.
literal-next	Arrange for the next character to be treated as a normal character. This allows control characters to be entered.
cursor-right	Move the cursor one character right.
cursor-left	Move the cursor one character left.
insert-mode	Toggle between insert mode and overwrite mode.
beginning-of-line	Move the cursor to the beginning of the line.
end-of-line	Move the cursor to the end of the line.
delete-line	Delete the contents of the current line.
kill-line	Delete everything that follows the cursor.
backward-kill-line	Delete all characters between the cursor and the start of the line.
forward-word	Move to the end of the word which follows the cursor.
forward-to-word	Move the cursor to the start of the word that follows the cursor.
backward-word	Move to the start of the word which precedes the cursor.
goto-column	Move the cursor to the 1-relative column in the line specified by any preceding digit-argument sequences (see Entering Repeat Counts below).
find-parenthesis	If the cursor is currently over a parenthesis character, move it to the matching parenthesis character. If not over a parenthesis character move right to the next close parenthesis.
forward-delete-char	Delete the character under the cursor.
backward-delete-char	Delete the character which precedes the cursor.
list-or-eof	This is intended for binding to ^D. When invoked when the cursor is within the line it displays all possible completions then redisplay the line unchanged. When invoked on an empty line, it signals end-of-input (EOF) to the caller of <code>gl_get_line()</code> .
del-char-or-list-or-eof	This is intended for binding to ^D. When invoked when the cursor is within the line it invokes forward-delete-char. When invoked at the end of the line it displays all possible

	completions then redisplay the line unchanged. When invoked on an empty line, it signals end-of-input (EOF) to the caller of <code>gl_get_line()</code> .
forward-delete-word	Delete the word which follows the cursor.
backward-delete-word	Delete the word which precedes the cursor.
upcase-word	Convert all of the characters of the word which follows the cursor, to upper case.
downcase-word	Convert all of the characters of the word which follows the cursor, to lower case.
capitalize-word	Capitalize the word which follows the cursor.
change-case	If the next character is upper case, toggle it to lower case and vice versa.
redisplay	Redisplay the line.
clear-screen	Clear the terminal, then redisplay the current line.
transpose-chars	Swap the character under the cursor with the character just before the cursor.
set-mark	Set a mark at the position of the cursor.
exchange-point-and-mark	Move the cursor to the last mark that was set, and move the mark to where the cursor used to be.
kill-region	Delete the characters that lie between the last mark that was set, and the cursor.
copy-region-as-kill	Copy the text between the mark and the cursor to the cut buffer, without deleting the original text.
yank	Insert the text that was last deleted, just before the current position of the cursor.
append-yank	Paste the current contents of the cut buffer, after the cursor.
up-history	Recall the next oldest line that was entered. Note that in <code>vi</code> mode you are left in command mode.
down-history	Recall the next most recent line that was entered. If no history recall session is currently active, the next line from a previous recall session is recalled. Note that in <code>vi</code> mode you are left in command mode.
history-search-backward	Recall the next oldest line whose prefix matches the string which currently precedes the cursor (in <code>vi</code> command-mode

---

	the character under the cursor is also included in the search string). Note that in <code>vi</code> mode you are left in command mode.
<code>history-search-forward</code>	Recall the next newest line who's prefix matches the string which currently precedes the cursor (in <code>vi</code> command-mode the character under the cursor is also included in the search string). Note that in <code>vi</code> mode you are left in command mode.
<code>history-re-search-backward</code>	Recall the next oldest line who's prefix matches that established by the last invocation of either <code>history-search-forward</code> or <code>history-search-backward</code> .
<code>history-re-search-forward</code>	Recall the next newest line who's prefix matches that established by the last invocation of either <code>history-search-forward</code> or <code>history-search-backward</code> .
<code>complete-word</code>	Attempt to complete the incomplete word which precedes the cursor. Unless the host program has customized word completion, file name completion is attempted. In <code>vi</code> command mode the character under the cursor is also included in the word being completed, and you are left in <code>vi</code> insert mode.
<code>expand-filename</code>	Within the command line, expand wild cards, tilde expressions and dollar expressions in the file name which immediately precedes the cursor. In <code>vi</code> command mode the character under the cursor is also included in the file name being expanded, and you are left in <code>vi</code> insert mode.
<code>list-glob</code>	List any file names which match the wild-card, tilde and dollar expressions in the file name which immediately precedes the cursor, then redraw the input line unchanged.
<code>list-history</code>	Display the contents of the history list for the current history group. If a repeat count of <code>&gt; 1</code> is specified, only that many of the most recent lines are displayed. See the Entering Repeat Counts section.
<code>read-from-file</code>	Temporarily switch to reading input from the file who's name precedes the cursor.
<code>read-init-files</code>	Re-read <code>teclarc</code> configuration files.
<code>beginning-of-history</code>	Move to the oldest line in the history list. Note that in <code>vi</code> mode you are left in command mode.
<code>end-of-history</code>	Move to the newest line in the history list (ie. the current line). Note that in <code>vi</code> mode this leaves you in command mode.

digit-argument	Enter a repeat count for the next key binding function. For details, see the Entering Repeat Counts section.
newline	Terminate and return the current contents of the line, after appending a newline character. The newline character is normally '\n', but will be the first character of the key sequence that invoked the newline action, if this happens to be a printable character. If the action was invoked by the '\n' newline character or the '\r' carriage return character, the line is appended to the history buffer.
repeat-history	Return the line that is being edited, then arrange for the next most recent entry in the history buffer to be recalled when tecla is next called. Repeatedly invoking this action causes successive historical input lines to be re-executed. Note that this action is equivalent to the 'Operate' action in ksh.
ring-bell	Ring the terminal bell, unless the bell has been silenced via the nobeep configuration option (see The Tecla Configuration File section).
forward-copy-char	Copy the next character into the cut buffer (NB. use repeat counts to copy more than one).
backward-copy-char	Copy the previous character into the cut buffer.
forward-copy-word	Copy the next word into the cut buffer.
backward-copy-word	Copy the previous word into the cut buffer.
forward-find-char	Move the cursor to the next occurrence of the next character that you type.
backward-find-char	Move the cursor to the last occurrence of the next character that you type.
forward-to-char	Move the cursor to the character just before the next occurrence of the next character that the user types.
backward-to-char	Move the cursor to the character just after the last occurrence before the cursor of the next character that the user types.
repeat-find-char	Repeat the last backward-find-char, forward-find-char, backward-to-char or forward-to-char.
invert-refind-char	Repeat the last backward-find-char, forward-find-char, backward-to-char, or forward-to-char in the opposite direction.

---

delete-to-column	Delete the characters from the cursor up to the column that is specified by the repeat count.
delete-to-parenthesis	Delete the characters from the cursor up to and including the matching parenthesis, or next close parenthesis.
forward-delete-find	Delete the characters from the cursor up to and including the following occurrence of the next character typed.
backward-delete-find	Delete the characters from the cursor up to and including the preceding occurrence of the next character typed.
forward-delete-to	Delete the characters from the cursor up to, but not including, the following occurrence of the next character typed.
backward-delete-to	Delete the characters from the cursor up to, but not including, the preceding occurrence of the next character typed.
delete-refind	Repeat the last *-delete-find or *-delete-to action.
delete-invert-refind	Repeat the last *-delete-find or *-delete-to action, in the opposite direction.
copy-to-column	Copy the characters from the cursor up to the column that is specified by the repeat count, into the cut buffer.
copy-to-parenthesis	Copy the characters from the cursor up to and including the matching parenthesis, or next close parenthesis, into the cut buffer.
forward-copy-find	Copy the characters from the cursor up to and including the following occurrence of the next character typed, into the cut buffer.
backward-copy-find	Copy the characters from the cursor up to and including the preceding occurrence of the next character typed, into the cut buffer.
forward-copy-to	Copy the characters from the cursor up to, but not including, the following occurrence of the next character typed, into the cut buffer.
backward-copy-to	Copy the characters from the cursor up to, but not including, the preceding occurrence of the next character typed, into the cut buffer.
copy-refind	Repeat the last *-copy-find or *-copy-to action.

copy-invert-refind	Repeat the last *-copy-find or *-copy-to action, in the opposite direction.
vi-mode	Switch to vi mode from emacs mode.
emacs-mode	Switch to emacs mode from vi mode.
vi-insert	From vi command mode, switch to insert mode.
vi-overwrite	From vi command mode, switch to overwrite mode.
vi-insert-at-bol	From vi command mode, move the cursor to the start of the line and switch to insert mode.
vi-append-at-eol	From vi command mode, move the cursor to the end of the line and switch to append mode.
vi-append	From vi command mode, move the cursor one position right, and switch to insert mode.
vi-replace-char	From vi command mode, replace the character under the cursor with the next character entered.
vi-forward-change-char	From vi command mode, delete the next character then enter insert mode.
vi-backward-change-char	From vi command mode, delete the preceding character then enter insert mode.
vi-forward-change-word	From vi command mode, delete the next word then enter insert mode.
vi-backward-change-word	From vi command mode, delete the preceding word then enter insert mode.
vi-change-rest-of-line	From vi command mode, delete from the cursor to the end of the line, then enter insert mode.
vi-change-line	From vi command mode, delete the current line, then enter insert mode.
vi-change-to-bol	From vi command mode, delete all characters between the cursor and the beginning of the line, then enter insert mode.
vi-change-to-column	From vi command mode, delete the characters from the cursor up to the column that is specified by the repeat count, then enter insert mode.
vi-change-to-parenthesis	Delete the characters from the cursor up to and including the matching parenthesis, or next close parenthesis, then enter vi insert mode.

vi-forward-change-find	From vi command mode, delete the characters from the cursor up to and including the following occurrence of the next character typed, then enter insert mode.
vi-backward-change-find	From vi command mode, delete the characters from the cursor up to and including the preceding occurrence of the next character typed, then enter insert mode.
vi-forward-change-to	From vi command mode, delete the characters from the cursor up to, but not including, the following occurrence of the next character typed, then enter insert mode.
vi-backward-change-to	From vi command mode, delete the characters from the cursor up to, but not including, the preceding occurrence of the next character typed, then enter insert mode.
vi-change-refind	Repeat the last vi- <i>*</i> -change-find or vi- <i>*</i> -change-to action.
vi-change-invert-refind	Repeat the last vi- <i>*</i> -change-find or vi- <i>*</i> -change-to action, in the opposite direction.
vi-undo	In vi mode, undo the last editing operation.
vi-repeat-change	In vi command mode, repeat the last command that modified the line.

#### Default Key Bindings In emacs Mode

The following default key bindings, which can be overridden by the tecla configuration file, are designed to mimic most of the bindings of the unix tcsh shell, when it is in emacs editing mode.

This is the default editing mode of the tecla library.

Under UNIX the terminal driver sets a number of special keys for certain functions. The tecla library attempts to use the same key bindings to maintain consistency. The key sequences shown for the following 6 bindings are thus just examples of what they will probably be set to. If you have used the stty command to change these keys, then the default bindings should match.

^C	user-interrupt
^\literal>	abort
^Z	suspend
^Q	start-output
^S	stop-output
^V	literal-next

The cursor keys are referred to by name, as follows. This is necessary because different types of terminals generate different key sequences when their cursor keys are pressed.

right    cursor-right  
left     cursor-left  
up       up-history  
down    down-history

The remaining bindings don't depend on the terminal settings.

^F                    cursor-right  
^B                    cursor-left  
M-i                   insert-mode  
^A                    beginning-of-line  
^E                    end-of-line  
^U                    delete-line  
^K                    kill-line  
M-f                   forward-word  
M-b                   backward-word  
^D                    del-char-or-list-or-eof  
^H                    backward-delete-char  
^?                    backward-delete-char  
M-d                   forward-delete-word  
M-^H                  backward-delete-word  
M-^?                  backward-delete-word  
M-u                   upcase-word  
M-l                   downcase-word  
M-c                   capitalize-word  
^R                    redisplay  
^L                    clear-screen  
^T                    transpose-chars  
^@                    set-mark

---

<code>^X^X</code>	exchange-point-and-mark
<code>^W</code>	kill-region
<code>M-w</code>	copy-region-as-kill
<code>^Y</code>	yank
<code>^P</code>	up-history
<code>^N</code>	down-history
<code>M-p</code>	history-search-backward
<code>M-n</code>	history-search-forward
<code>^I</code>	complete-word
<code>^X*</code>	expand-filename
<code>^X^F</code>	read-from-file
<code>^X^R</code>	read-init-files
<code>^Xg</code>	list-glob
<code>^Xh</code>	list-history
<code>M-&lt;</code>	beginning-of-history
<code>M-&gt;</code>	end-of-history
<code>\n</code>	newline
<code>\r</code>	newline
<code>M-o</code>	repeat-history
<code>M-^V</code>	vi-mode
<code>M-0, M-1, ... M-9</code>	digit-argument (see below)

Note that `^I` is what the TAB key generates, and that `^@` can be generated not only by pressing the CONTROL key and the @ key simultaneously, but also by pressing the CONTROL key and the space bar at the same time.

#### Default Key Bindings in vi Mode

The following default key bindings are designed to mimic the vi style of editing as closely as possible. This means that very few editing functions are provided in the initial character input mode, editing functions instead being provided by the vi command mode. The vi command mode is entered whenever the ESCAPE character is pressed, or whenever a key sequence that starts with a meta character is entered. In addition to mimicing vi, libtecla provides bindings for tab completion, wild-card expansion of file names, and historical line recall.

To learn how to tell the tecla library to use vi mode instead of the default emacs editing mode, see the earlier section entitled The Tecla Configuration File.

Under UNIX the terminal driver sets a number of special keys for certain functions. The tecla library attempts to use the same key bindings to maintain consistency, binding them both in input mode and in command mode. The key sequences shown for the following 6 bindings are thus just examples of what they will probably be set to. If you have used the stty command to change these keys, then the default bindings should match.

<code>^C</code>	user-interrupt
<code>^\</code>	abort
<code>^Z</code>	suspend
<code>^Q</code>	start-output
<code>^S</code>	stop-output
<code>^V</code>	literal-next
<code>M-^C</code>	user-interrupt
<code>M-^\literal&gt;</code>	abort
<code>M-^Z</code>	suspend
<code>M-^Q</code>	start-output
<code>M-^S</code>	stop-output

Note that above, most of the bindings are defined twice, once as a raw control code like `^C` and then a second time as a META character like `M-^C`. The former is the binding for vi input mode, whereas the latter is the binding for vi command mode. Once in command mode all key sequences that the user types that they don't explicitly start with an ESCAPE or a META key, have their first key secretly converted to a META character before the key sequence is looked up in the key binding table. Thus, once in command mode, when you type the letter `i`, for example, the tecla library actually looks up the binding for `M-i`.

The cursor keys are referred to by name, as follows. This is necessary because different types of terminals generate different key sequences when their cursor keys are pressed.

<code>right</code>	cursor-right
<code>left</code>	cursor-left
<code>up</code>	up-history
<code>down</code>	down-history

The cursor keys normally generate a key sequence that start with an ESCAPE character, so beware that using the arrow keys will put you into command mode (if you aren't already in command mode).

The following are the terminal-independent key bindings for vi input mode.

<code>^D</code>	list-or-eof
<code>^G</code>	list-glob
<code>^H</code>	backward-delete-char
<code>^I</code>	complete-word
<code>\r</code>	newline
<code>\n</code>	newline
<code>^L</code>	clear-screen
<code>^N</code>	down-history
<code>^P</code>	up-history
<code>^R</code>	redisplay
<code>^U</code>	backward-kill-line
<code>^W</code>	backward-delete-word
<code>^X*</code>	expand-filename
<code>^X^F</code>	read-from-file
<code>^X^R</code>	read-init-files
<code>^?</code>	backward-delete-char

The following are the key bindings that are defined in vi command mode, this being specified by them all starting with a META character. As mentioned above, once in command mode the initial meta character is optional. For example, you might enter command mode by typing ESCAPE, and then press 'H' twice to move the cursor two positions to the left. Both 'H' characters get quietly converted to M-h before being compared to the key binding table, the first one because ESCAPE followed by a character is always converted to the equivalent META character, and the second because command mode was already active.

<code>M-\</code>	cursor-right (META-space)
<code>M-\$</code>	end-of-line
<code>M-*</code>	expand-filename
<code>M-+</code>	down-history

M - -	up-history
M - <	beginning-of-history
M - >	end-of-history
M - ^	beginning-of-line
M -	repeat-find-char
M - ,	invert-refind-char
M -	goto-column
M - ~	change-case
M - .	vi-repeat-change
M - %	find-parenthesis
M - a	vi-append
M - A	vi-append-at-eol
M - b	backward-word
M - B	backward-word
M - C	vi-change-rest-of-line
M - cb	vi-backward-change-word
M - cB	vi-backward-change-word
M - cc	vi-change-line
M - ce	vi-forward-change-word
M - cE	vi-forward-change-word
M - cw	vi-forward-change-word
M - cW	vi-forward-change-word
M - cF	vi-backward-change-find
M - cf	vi-forward-change-find
M - cT	vi-backward-change-to
M - ct	vi-forward-change-to
M - c ;	vi-change-refind
M - c ,	vi-change-invert-refind
M - ch	vi-backward-change-char

---

M-c^H	vi-backward-change-char
M-c^?	vi-backward-change-char
M-cl	vi-forward-change-char
M-c\	vi-forward-change-char (META-c-space)
M-c^	vi-change-to-bol
M-c0	vi-change-to-bol
M-c\$	vi-change-rest-of-line
M-c	vi-change-to-column
M-c%	vi-change-to-parenthesis
M-dh	backward-delete-char
M-d^H	backward-delete-char
M-d^?	backward-delete-char
M-dl	forward-delete-char
M-d	forward-delete-char (META-d-space)
M-dd	delete-line
M-db	backward-delete-word
M-dB	backward-delete-word
M-de	forward-delete-word
M-dE	forward-delete-word
M-dw	forward-delete-word
M-dW	forward-delete-word
M-dF	backward-delete-find
M-df	forward-delete-find
M-dT	backward-delete-to
M-dt	forward-delete-to
M-d;	delete-refind
M-d,	delete-invert-refind
M-d^	backward-kill-line
M-d0	backward-kill-line

M-d\$	kill-line
M-D	kill-line
M-d	delete-to-column
M-d%	delete-to-parenthesis
M-e	forward-word
M-E	forward-word
M-f	forward-find-char
M-F	backward-find-char
M-.	up-history
M-h	cursor-left
M-H	beginning-of-history
M-i	vi-insert
M-I	vi-insert-at-bol
M-j	down-history
M-J	history-search-forward
M-k	up-history
M-K	history-search-backward
M-l	cursor-right
M-L	end-of-history
M-n	history-re-search-forward
M-N	history-re-search-backward
M-p	append-yank
M-P	yank
M-r	vi-replace-char
M-R	vi-overwrite
M-s	vi-forward-change-char
M-S	vi-change-line
M-t	forward-to-char
M-T	backward-to-char

---

M-u	vi-undo
M-w	forward-to-word
M-W	forward-to-word
M-x	forward-delete-char
M-X	backward-delete-char
M-yh	backward-copy-char
M-y^H	backward-copy-char
M-y^?	backward-copy-char
M-y\l	forward-copy-char
M-y\ M-ye	forward-copy-char (META-y-space)
M-yE	forward-copy-word
M-yw	forward-copy-word
M-yW	forward-copy-word
M-yb	backward-copy-word
M-yB	backward-copy-word
M-yf	forward-copy-find
M-yF	backward-copy-find
M-yt	forward-copy-to
M-yT	backward-copy-to
M-y;	copy-refind
M-y,	copy-invert-refind
M-y^	copy-to-bol
M-y0	copy-to-bol
M-y\$	copy-rest-of-line
M-yy	copy-line
M-Y	copy-line
M-y	copy-to-column
M-y%	copy-to-parenthesis

M-^E	emacs-mode
M-^H	cursor-left
M-^?	cursor-left
M-^L	clear-screen
M-^N	down-history
M-^P	up-history
M-^R	redisplay
M-^D	list-or-eof
M-^I	complete-word
M-\ r	newline
M-\ n	newline
M-^X^R	read-init-files
M-^Xh	list-history
M-0, M-1, . . . M-9	digit-argument (see below)

Note that ^I is what the TAB key generates.

Entering Repeat Counts Many of the key binding functions described previously, take an optional count, typed in before the target key sequence. This is interpreted as a repeat count by most bindings. A notable exception is the goto-column binding, which interprets the count as a column number.

By default you can specify this count argument by pressing the META key while typing in the numeric count. This relies on the digit-argument action being bound to 'META-0', 'META-1' etc. Once any one of these bindings has been activated, you can optionally take your finger off the META key to type in the rest of the number, since every numeric digit thereafter is treated as part of the number, unless it is preceded by the literal-next binding. As soon as a non-digit, or literal digit key is pressed the repeat count is terminated and either causes the just typed character to be added to the line that many times, or causes the next key binding function to be given that argument.

For example, in emacs mode, typing:

M-12a

causes the letter 'a' to be added to the line 12 times, whereas

M-4M-c

Capitalizes the next 4 words.

In vi command mode the meta modifier is automatically added to all characters typed in, so to enter a count in vi command-mode, just involves typing in the number, just as it does in the vi editor itself. So for example, in vi command mode, typing:

```
4w2x
```

moves the cursor four words to the right, then deletes two characters.

You can also bind digit-argument to other key sequences. If these end in a numeric digit, that digit gets appended to the current repeat count. If it doesn't end in a numeric digit, a new repeat count is started with a value of zero, and can be completed by typing in the number, after letting go of the key which triggered the digit-argument action.

**Files**

<code>/usr/lib/libtecla.so</code>	The tecla library
<code>/usr/include/libtecla.h</code>	The tecla header file
<code>~/.teclarc</code>	The personal tecla customization file

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	library/libtecla
Interface Stability	Committed

**See Also** [vi\(1\)](#), [cpl\\_complete\\_word\(3TECLA\)](#), [ef\\_expand\\_file\(3TECLA\)](#), [gl\\_get\\_line\(3TECLA\)](#), [gl\\_io\\_mode\(3TECLA\)](#), [libtecla\(3LIB\)](#), [pca\\_lookup\\_file\(3TECLA\)](#), [attributes\(5\)](#)

**Name** term – conventional names for terminals

**Description** Terminal names are maintained as part of the shell environment in the environment variable TERM. See [sh\(1\)](#), [profile\(4\)](#), and [environ\(5\)](#). These names are used by certain commands (for example, `tabs`, `tput`, and `vi`) and certain functions (for example, see [curses\(3CURSES\)](#)).

Files under `/usr/share/lib/terminfo` are used to name terminals and describe their capabilities. These files are in the format described in [terminfo\(4\)](#). Entries in `terminfo` source files consist of a number of comma-separated fields. To print a description of a terminal *term*, use the command `infocmp -I term`. See [infocmp\(1M\)](#). White space after each comma is ignored. The first line of each terminal description in the `terminfo` database gives the names by which `terminfo` 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 `TERMINFO` in `$HOME/.profile`; see [profile\(4\)](#)), 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.

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, `att4425`. 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 `a` through `z` and `0` through `9`, 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 `att4425-w`. The following suffixes should be used where possible:

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>
<code>-rv</code>	Reverse video	<code>att4415-rv</code>

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
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\(5\)](#)

**Name** threads, pthreads – POSIX pthreads and Solaris threads concepts

### Synopsis

```
POSIX cc -mt [ flag... ] file... [ -lrt library... ]
```

```
#include <pthread.h>
```

```
Solaris cc -mt [ flag... ] file... [ library... ]
```

```
#include <sched.h>
```

```
#include <thread.h>
```

**Description** POSIX and Solaris threads each have their own implementation within `libc(3LIB)`. Both implementations are interoperable, their functionality similar, and can be used within the same application. Only POSIX threads are guaranteed to be fully portable to other POSIX-compliant environments. POSIX and Solaris threads require different source, include files and linking libraries. See SYNOPSIS.

**Similarities** Most of the POSIX and Solaris threading functions have counterparts with each other. POSIX function names, with the exception of the semaphore names, have a “pthread” prefix. Function names for similar POSIX and Solaris functions have similar endings. Typically, similar POSIX and Solaris functions have the same number and use of arguments.

**Differences** POSIX pthreads and Solaris threads differ in the following ways:

- POSIX threads are more portable.
- POSIX threads establish characteristics for each thread according to configurable attribute objects.
- POSIX pthreads implement thread cancellation.
- POSIX pthreads enforce scheduling algorithms.
- POSIX pthreads allow for clean-up handlers for `fork(2)` calls.
- Solaris threads can be suspended and continued.
- Solaris threads implement daemon threads, for whose demise the process does not wait.

**Function Comparison** The following table compares the POSIX pthreads and Solaris threads functions. When a comparable interface is not available either in POSIX pthreads or Solaris threads, a hyphen (–) appears in the column.

Functions Related to Creation	POSIX	Solaris
<code>pthread_create()</code>		<code>thr_create()</code>
<code>pthread_attr_init()</code>		–

---

pthread_attr_setdetachstate()	–
pthread_attr_getdetachstate()	–
pthread_attr_setinheritsched()	–
pthread_attr_getinheritsched()	–
pthread_attr_setschedparam()	–
pthread_attr_getschedparam()	–
pthread_attr_setschedpolicy()	–
pthread_attr_getschedpolicy()	–
pthread_attr_setscope()	–
pthread_attr_getscope()	–
pthread_attr_setstackaddr()	–
pthread_attr_getstackaddr()	–
pthread_attr_setstacksize()	–
pthread_attr_getstacksize()	–
pthread_attr_getguardsize()	–
pthread_attr_setguardsize()	–
pthread_attr_destroy()	–
–	thr_min_stack()

---



---

**Functions Related to Exit**

	POSIX	Solaris
pthread_exit()		thr_exit()
pthread_join()		thr_join()
pthread_detach()		–

---

**Functions Related to Thread Specific Data**

	POSIX	Solaris
pthread_key_create()		thr_keycreate()
pthread_setspecific()		thr_setspecific()
pthread_getspecific()		thr_getspecific()
pthread_key_delete()		–

---

Functions Related to Signals	POSIX	Solaris
	pthread_sigmask()	thr_sigsetmask()
	pthread_kill()	thr_kill()
Functions Related to IDs	POSIX	Solaris
	pthread_self()	thr_self()
	pthread_equal()	–
	–	thr_main()
Functions Related to Scheduling	POSIX	Solaris
	–	thr_yield()
	–	thr_suspend()
	–	thr_continue()
	pthread_setconcurrency()	thr_setconcurrency()
	pthread_getconcurrency()	thr_getconcurrency()
	pthread_setschedparam()	thr_setprio()
	pthread_setschedprio()	thr_setprio()
	pthread_getschedparam()	thr_getprio()
Functions Related to Cancellation	POSIX	Solaris
	pthread_cancel()	–
	pthread_setcancelstate()	–
	pthread_setcanceltype()	–
	pthread_testcancel()	–
	pthread_cleanup_pop()	–
	pthread_cleanup_push()	–
Functions Related to Mutexes	POSIX	Solaris
	pthread_mutex_init()	mutex_init()

---

<code>pthread_mutexattr_init()</code>	–
<code>pthread_mutexattr_setpshared()</code>	–
<code>pthread_mutexattr_getpshared()</code>	–
<code>pthread_mutexattr_setprotocol()</code>	–
<code>pthread_mutexattr_getprotocol()</code>	–
<code>pthread_mutexattr_setprioceiling()</code>	–
<code>pthread_mutexattr_getprioceiling()</code>	–
<code>pthread_mutexattr_settype()</code>	–
<code>pthread_mutexattr_gettype()</code>	–
<code>pthread_mutexattr_setrobust()</code>	–
<code>pthread_mutexattr_getrobust()</code>	–
<code>pthread_mutexattr_destroy()</code>	–
<code>pthread_mutex_setprioceiling()</code>	–
<code>pthread_mutex_getprioceiling()</code>	–
<code>pthread_mutex_lock()</code>	<code>mutex_lock()</code>
<code>pthread_mutex_trylock()</code>	<code>mutex_trylock()</code>
<code>pthread_mutex_unlock()</code>	<code>mutex_unlock()</code>
<code>pthread_mutex_destroy()</code>	<code>mutex_destroy()</code>

---

**Functions Related to  
Condition Variables**


---

	POSIX	Solaris
<code>pthread_cond_init()</code>		<code>cond_init()</code>
<code>pthread_condattr_init()</code>		–
<code>pthread_condattr_setpshared()</code>		–
<code>pthread_condattr_getpshared()</code>		–
<code>pthread_condattr_destroy()</code>		–
<code>pthread_cond_wait()</code>		<code>cond_wait()</code>
<code>pthread_cond_timedwait()</code>		<code>cond_timedwait()</code>
<code>pthread_cond_signal()</code>		<code>cond_signal()</code>
<code>pthread_cond_broadcast()</code>		<code>cond_broadcast()</code>

---

	POSIX	Solaris
	<code>pthread_cond_destroy()</code>	<code>cond_destroy()</code>
<b>Functions Related to Reader/Writer Locking</b>		
	<code>pthread_rwlock_init()</code>	<code>rwlock_init()</code>
	<code>pthread_rwlock_rdlock()</code>	<code>rw_rdlock()</code>
	<code>pthread_rwlock_tryrdlock()</code>	<code>rw_tryrdlock()</code>
	<code>pthread_rwlock_wrlock()</code>	<code>rw_wrlock()</code>
	<code>pthread_rwlock_trywrlock()</code>	<code>rw_trywrlock()</code>
	<code>pthread_rwlock_unlock()</code>	<code>rw_unlock()</code>
	<code>pthread_rwlock_destroy()</code>	<code>rwlock_destroy()</code>
	<code>pthread_rwlockattr_init()</code>	–
	<code>pthread_rwlockattr_destroy()</code>	–
	<code>pthread_rwlockattr_getpshared()</code>	–
	<code>pthread_rwlockattr_setpshared()</code>	–
<b>Functions Related to Semaphores</b>		
	<code>sem_init()</code>	<code>sema_init()</code>
	<code>sem_open()</code>	–
	<code>sem_close()</code>	–
	<code>sem_wait()</code>	<code>sema_wait()</code>
	<code>sem_trywait()</code>	<code>sema_trywait()</code>
	<code>sem_post()</code>	<code>sema_post()</code>
	<code>sem_getvalue()</code>	–
	<code>sem_unlink()</code>	–
	<code>sem_destroy()</code>	<code>sema_destroy()</code>
<b>Functions Related to <code>fork()</code> Clean Up</b>		
	<code>pthread_atfork()</code>	–

Functions Related to Limits	POSIX	Solaris
	pthread_once()	–
Functions Related to Debugging	POSIX	Solaris
	–	thr_stksegment()

## Locking

**Synchronization** Multithreaded behavior is asynchronous, and therefore, optimized for concurrent and parallel processing. As threads, always from within the same process and sometimes from multiple processes, share global data with each other, they are not guaranteed exclusive access to the shared data at any point in time. Securing mutually exclusive access to shared data requires synchronization among the threads. Both POSIX and Solaris implement four synchronization mechanisms: mutexes, condition variables, reader/writer locking (*optimized frequent-read occasional-write mutex*), and semaphores.

Synchronizing multiple threads diminishes their concurrency. The coarser the grain of synchronization, that is, the larger the block of code that is locked, the lesser the concurrency.

**MT fork()** If a threads program calls `fork(2)`, it implicitly calls `fork1(2)`, which replicates only the calling thread. Should there be any outstanding mutexes throughout the process, the application should call `pthread_atfork(3C)` to wait for and acquire those mutexes prior to calling `fork()`.

## Scheduling

**POSIX Threads** Solaris supports the following three POSIX scheduling policies:

- `SCHED_OTHER` Traditional Timesharing scheduling policy. It is based on the timesharing (TS) scheduling class.
- `SCHED_FIFO` First-In-First-Out scheduling policy. Threads scheduled to this policy, if not preempted by a higher priority, will proceed until completion. Such threads are in real-time (RT) scheduling class. The calling process must have the `{PRIV_PROC_PRIORCTL}` privilege asserted in its effective set.
- `SCHED_RR` Round-Robin scheduling policy. Threads scheduled to this policy, if not preempted by a higher priority, will execute for a time period determined by the system. Such threads are in real-time (RT) scheduling class and the calling process must have the `{PRIV_PROC_PRIORCTL}` privilege asserted in its effective set.

In addition to the POSIX-specified scheduling policies above, Solaris also supports these scheduling policies:

- SCHED\_IA Threads are scheduled according to the Inter-Active Class (IA) policy as described in [prioctl\(2\)](#).
- SCHED\_FSS Threads are scheduled according to the Fair-Share Class (FSS) policy as described in [prioctl\(2\)](#).
- SCHED\_FX Threads are scheduled according to the Fixed-Priority Class (FX) policy as described in [prioctl\(2\)](#).

SolarisThreads Only scheduling policy supported is SCHED\_OTHER, which is timesharing, based on the TS scheduling class.

**Errors** In a multithreaded application, EINTR can be returned from blocking system calls when another thread calls [forkall\(2\)](#).

### Usage

-mt compiler option The -mt compiler option compiles and links for multithreaded code. It compiles source files with `-D_REENTRANT` and augments the set of support libraries properly.

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
MT-Level	MT-Safe, Fork 1-Safe

**See Also** [crle\(1\)](#), [fork\(2\)](#), [prioctl\(2\)](#), [libpthread\(3LIB\)](#), [librt\(3LIB\)](#), [libthread\(3LIB\)](#), [pthread\\_atfork\(3C\)](#), [pthread\\_create\(3C\)](#), [attributes\(5\)](#), [privileges\(5\)](#), [standards\(5\)](#)

*Linker and Libraries Guide*

**Name** trusted\_extensions – Solaris Trusted Extensions

**Description** Solaris Trusted Extensions software is a specific configuration of the Solaris Operating System (Solaris OS). Solaris Trusted Extensions (Trusted Extensions) provides labels for local objects and processes, for the desktop and windowing system, for zones and file systems, and for network communications. These labels are used to implement a Multilevel Security (MLS) policy that restricts the flow of information based on label relationships. In contrast to Discretionary Access Control (DAC) based on ownership, the MLS policy enforced by Trusted Extensions is an example of Mandatory Access Control (MAC).

By default, Trusted Extensions software is disabled. It is enabled and disabled (but not configured) by the `labeld(1M)` service, identified by the FMRI:

```
svc:/system/labeld:default
```

Refer to the Administrator's Guide listed below for the required configuration of Trusted Extensions software necessary before use. The system must be rebooted after enabling or disabling `labeld` to activate or deactivate Trusted Extensions software.

**See Also** `labeld(1M)`, `label_encodings(4)`, `labels(5)`

*Trusted Extensions Configuration and Administration*

*Trusted Extensions User's Guide*

**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** `vgrindefs` uses regular expressions similar to those of `ex(1)` and `lex(1)`. The characters ‘^’, ‘\$’, ‘.’, and ‘\’ are reserved characters and must be ‘quoted’ with a preceding ‘\’ if they are to be included as normal characters. The metasympols and their meanings are:

- \$ The end of a line
- ^ The beginning of a line
- \d A delimiter (space, tab, newline, start of line)
- \a Matches any string of symbols (like ‘.’ in `lex`)
- \p Matches any identifier. In a procedure definition (the ‘pb’ capability) the string that matches this symbol is used as the procedure name.
- () Grouping
- | Alternation
- ? Last item is optional
- \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.

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.

**Keyword List** 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.

**Examples** **EXAMPLE 1** A sample program.

The following entry, which describes the C language, is typical of a language entry.

```
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:
```

Note that the first field is just the language name (and any variants of it). Thus the C language could be specified to `vgrind(1)` as ‘c’ or ‘C’.

**Files** `/usr/lib/vgrindefs` file containing vgrind descriptions

**See Also** [ex\(1\)](#), [lex\(1\)](#), [troff\(1\)](#), [vgrind\(1\)](#)

<b>Name</b>	wbem – Web-Based Enterprise Management				
<b>Description</b>	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:				
Common Information Model (CIM)	<p>CIM is 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:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="vertical-align: top;">CIM Specification</td> <td>Consists of the language and methodology that describes management data.</td> </tr> <tr> <td style="vertical-align: top;">CIM Schema</td> <td>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.</td> </tr> </table> <p>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> <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>	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 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.				
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"> <li>▪ Solaris WBEM Services</li> <li>▪ Solaris WBEM Software Development Kit (SDK)</li> </ul>				
Solaris WBEM Services	<p>These services consist of a set of value-added Services 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 consists of:</p> <ul style="list-style-type: none"> <li>▪ CIM Object Manager, CIM Repository, and MOF Compiler</li> </ul>				

- 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 SDK** The Solaris WBEM SDK is a set of application programming interfaces (APIs) that contain the components necessary to write management applications. These applications communicate with WBEM-enabled management devices by using XML and HTTP communication standards.

Solaris WBEM applications request information or services from the Common Information Model (CIM) Object Manager through the WBEM APIs. These APIs represent CIM objects as Java classes. The APIs are used to describe managed objects and to retrieve information about managed objects in a system environment. The advantage of modeling managed resources by using CIM is that those objects can be shared across any system that is CIM-compliant.

For more information on the Solaris WBEM SDK, see the *Solaris WBEM Developer's Guide*. The Solaris WBEM API documentation is available in Javadoc format with the Solaris OS installation 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, SUNWwbco, SUNWwbco, 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\(5\)](#)

**Name** zones – Solaris application containers

**Description** The zones facility in Solaris provides an isolated environment for running applications. Processes running in a zone are prevented from monitoring or interfering with other activity in the system. Access to other processes, network interfaces, file systems, devices, and inter-process communication facilities are restricted to prevent interaction between processes in different zones.

The privileges available within a zone are restricted to prevent operations with system-wide impact. See [privileges\(5\)](#).

You can configure and administer zones with the [zoneadm\(1M\)](#) and [zonecfg\(1M\)](#) utilities. You can specify the configuration details a zone, install file system contents including software packages into the zone, and manage the runtime state of the zone. You can use the [zlogin\(1\)](#) to run commands within an active zone. You can do this without logging in through a network-based login server such as [in.rlogind\(1M\)](#) or [sshd\(1M\)](#).

The autobooting of zones is enabled and disabled by the zones service, identified by the FMRI:

```
svc:/system/zones:default
```

See [zoneadm\(1M\)](#). Note that a zone has an `autoboot` property, which can be set to `true` (always autoboot). However, if the zones service is disabled, autoboot will not occur, regardless of the setting of the `autoboot` property for a given zone. See [zonecfg\(1M\)](#).

An alphanumeric name and numeric ID identify each active zone. Alphanumeric names are configured using the [zonecfg\(1M\)](#) utility. Numeric IDs are automatically assigned when the zone is booted. The [zonename\(1\)](#) utility reports the current zone name, and the [zoneadm\(1M\)](#) utility can be used to report the names and IDs of configured zones.

A zone can be in one of several states:

**CONFIGURED**

Indicates that the configuration for the zone has been completely specified and committed to stable storage.

**INCOMPLETE**

Indicates that the zone is in the midst of being installed or uninstalled, or was interrupted in the midst of such a transition.

**INSTALLED**

Indicates that the zone's configuration has been instantiated on the system: packages have been installed under the zone's root path.

**READY**

Indicates that the “virtual platform” for the zone has been established. For instance, file systems have been mounted, devices have been configured, but no processes associated with the zone have been started.

**RUNNING**

Indicates that user processes associated with the zone application environment are running.

**SHUTTING\_DOWN  
DOWN**

Indicates that the zone is being halted. The zone can become stuck in one of these states if it is unable to tear down the application environment state (such as mounted file systems) or if some portion of the virtual platform cannot be destroyed. Such cases require operator intervention.

**UNAVAILABLE**

Indicates that the zone has been installed but cannot be booted. A zone enters the unavailable state when the zone's storage is unavailable while `svc:/system/zones:default` is onlining or when the zone tries to boot; when archive-based installations fail after successful archive extraction; and when the zone's software is incompatible with the global zone's software, such as after an improper forced attach.

**Process Access Restrictions** Processes running inside a zone (aside from the global zone) have restricted access to other processes. Only processes in the same zone are visible through `/proc` (see [proc\(4\)](#)) or through system call interfaces that take process IDs such as `kill(2)` and `prioctl(2)`. Attempts to access processes that exist in other zones (including the global zone) fail with the same error code that would be issued if the specified process did not exist.

**Privilege Restrictions** Processes running within a non-global zone are restricted to a subset of privileges, in order to prevent one zone from being able to perform operations that might affect other zones. The set of privileges limits the capabilities of privileged users (such as the super-user or root user) within the zone. The list of privileges available within a zone can be displayed using the `ppriv(1)` utility. For more information about privileges, see [privileges\(5\)](#).

**Device Restrictions** The set of devices available within a zone is restricted, to prevent a process in one zone from interfering with processes in other zones. For example, a process in a zone should not be able to modify kernel memory using `/dev/kmem`, or modify the contents of the root disk. Thus, by default, only a few pseudo devices considered safe for use within a zone are available. Additional devices can be made available within specific zones using the `zonecfg(1M)` utility.

The device and privilege restrictions have a number of effects on the utilities that can run in a non-global zone. For example, the `eeprom(1M)`, `prtdiag(1M)`, and `prtconf(1M)` utilities do not work in a zone since they rely on devices that are not normally available.

**Brands** A zone can be assigned a brand when it is initially created. A branded zone is one whose software does not match that software found in the global zone. The software can include Solaris software configured or laid out differently, or it can include non-Solaris software. The particular collection of software is called a “brand” (see [brands\(5\)](#)). Once installed, a zone's brand can not be changed unless the zone is first uninstalled.

**File Systems** Each zone has its own section of the file system hierarchy, rooted at a directory known as the zone root. Processes inside the zone can access only files within that part of the hierarchy, that is, files that are located beneath the zone root. This prevents processes in one zone from corrupting or examining file system data associated with another zone. The `chroot(1M)` utility can be used within a zone, but can only restrict the process to a root path accessible within the zone.

In order to preserve file system space, sections of the file system can be mounted into one or more zones using the read-only option of the `lofs(7FS)` file system. This allows the same file system data to be shared in multiple zones, while preserving the security guarantees supplied by zones.

NFS and autofs mounts established within a zone are local to that zone; they cannot be accessed from other zones, including the global zone. The mounts are removed when the zone is halted or rebooted.

ZFS datasets that are delegated to a zone are manageable within the zone. Within a delegated dataset, child datasets can be created. Datasets that are created within a delegated dataset are themselves delegated. Delegated datasets other than the top level delegated dataset can be destroyed. Most, but not all, properties can be set on delegated datasets. See `zfs(1M)` for details.

Each zone has a top-level delegated dataset, which in turn contains the `ROOT` and potentially other datasets such as `.../export` and `.../export/home`. Datasets that exist under the `ROOT` dataset make up the zone's boot environment(s). Boot environment datasets should only be created or destroyed using the `zoneadm(1M)` or `beadm(1M)` commands.

**Networking** A zone has its own port number space for TCP, UDP, and SCTP applications and typically one or more separate IP addresses (but some configurations of Trusted Extensions share IP address(es) between zones).

For the IP layer (IP routing, ARP, IPsec, IP Filter, and so on) a zone can either share the configuration and state with the global zone (a shared-IP zone), or have its distinct IP layer configuration and state (an exclusive-IP zone).

If a zone is to be connected to the same datalink, that is, be on the same IP subnet or subnets as the global zone, then it is appropriate for the zone to use the shared IP instance.

If a zone needs to be isolated at the IP layer on the network, for instance being connected to different VLANs or different LANs than the global zone and other non-global zones, then for isolation reasons the zone should have its exclusive IP.

A shared-IP zone is prevented from doing certain things towards the network (such as changing its IP address or sending spoofed IP or Ethernet packets), but an exclusive-IP zone has more or less the same capabilities towards the network as a separate host that is connected to the same network interface. In particular, the superuser in such a zone can change its IP address and spoof ARP packets.

The shared-IP zones are assigned one or more network interface names and IP addresses in [zonecfg\(1M\)](#). The network interface name(s) must also be configured in the global zone.

The exclusive-IP zones are assigned one or more network interface names in [zonecfg\(1M\)](#). The network interface names must be exclusively assigned to that zone, that is, it (or they) can not be assigned to some other running zone, nor can they be used by the global zone.

The full IP-level functionality in the form of DHCP client, IPsec and IP Filter, is available in exclusive-IP zones and not in shared-IP zones.

**Host Identifiers** A zone is capable of emulating a 32-bit host identifier, which can be configured via [zonecfg\(1M\)](#), for the purpose of system consolidation. If a zone emulates a host identifier, then commands such as [hostid\(1\)](#) and [sysdef\(1M\)](#) as well as C interfaces such as [sysinfo\(2\)](#) and [gethostid\(3C\)](#) that are executed within the context of the zone will display or return the zone's emulated host identifier rather than the host machine's identifier.

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

**See Also** [hostid\(1\)](#), [zlogin\(1\)](#), [zonename\(1\)](#), [beadm\(1M\)](#), [in.rlogind\(1M\)](#), [sshd\(1M\)](#), [sysdef\(1M\)](#), [zfs\(1M\)](#), [zoneadm\(1M\)](#), [zonecfg\(1M\)](#), [kill\(2\)](#), [priocntl\(2\)](#), [sysinfo\(2\)](#), [gethostid\(3C\)](#), [getzoneid\(3C\)](#), [ucred\\_get\(3C\)](#), [proc\(4\)](#), [attributes\(5\)](#), [brands\(5\)](#), [privileges\(5\)](#), [crgetzoneid\(9F\)](#)

