

手册页第 5 部分：标准、环境和宏

版权所有 © 2011, Oracle 和/或其附属公司。保留所有权利。

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 RIGHTS

Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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.

目录

前言	7
简介	11
Intro(5)	12
标准、环境和宏	13
acl(5)	14
ad(5)	21
ascii(5)	22
attributes(5)	24
audit_binfile(5)	32
audit_flags(5)	34
audit_remote(5)	35
audit_syslog(5)	40
brands(5)	43
cancellation(5)	44
charmap(5)	50
condition(5)	54
crypt_bsdbf(5)	56
crypt_bsmd5(5)	57
crypt_sha256(5)	58
crypt_sha512(5)	59
crypt_sunmd5(5)	60
crypt_unix(5)	61
device_clean(5)	62
dhcp(5)	64
dhcp_modules(5)	66

environ(5)	67
eqnchar(5)	72
extendedFILE(5)	73
extensions(5)	75
filesystem(5)	76
fmri(5)	80
fnmatch(5)	98
formats(5)	102
fsattr(5)	107
grub(5)	118
gss_auth_rules(5)	119
hal(5)	120
iconv_1250(5)	121
iconv_1251(5)	127
iconv(5)	136
iconv_646(5)	140
iconv_852(5)	143
iconv_8859-1(5)	150
iconv_8859-2(5)	157
iconv_8859-5(5)	163
iconv_dhn(5)	171
iconv_koi8-r(5)	175
iconv_mac_cyr(5)	183
iconv_maz(5)	191
iconv_pc_cyr(5)	195
iconv_unicode(5)	201
ieee802.11(5)	206
ieee802.3(5)	207
ipfilter(5)	213
isalist(5)	216
kerberos(5)	218
krb5_auth_rules(5)	220
krb5envvar(5)	222
kssl(5)	225
labels(5)	226
largefile(5)	228

ldap(5)	232
lf64(5)	236
lfcompile(5)	243
lfcompile64(5)	246
locale(5)	248
locale_alias(5)	274
man(5)	298
mansun(5)	303
me(5)	307
mech_spnego(5)	312
mm(5)	314
ms(5)	321
mutex(5)	326
mwac(5)	328
nfssec(5)	329
nis(5)	332
nwam(5)	333
openssl(5)	335
pam_allow(5)	338
pam_authtok_check(5)	340
pam_authtok_get(5)	342
pam_authtok_store(5)	344
pam_deny(5)	345
pam_dhkeys(5)	347
pam_dial_auth(5)	349
pam_krb5(5)	350
pam_krb5_migrate(5)	360
pam_ldap(5)	363
pam_list(5)	368
pam_passwd_auth(5)	371
pam_pkcs11(5)	373
pam_rhosts_auth(5)	377
pam_roles(5)	378
pam_sample(5)	380
pam_smbfs_login(5)	382
pam_smb_passwd(5)	384

pam_tsol_account(5)	386
pam_unix_account(5)	388
pam_unix_auth(5)	390
pam_unix_cred(5)	392
pam_unix_session(5)	394
pam_zfs_key(5)	395
pkcs11_kernel(5)	398
pkcs11_kms(5)	400
pkcs11_softtoken(5)	403
pkcs11_tpm(5)	407
privileges(5)	410
prof(5)	422
rbac(5)	423
regex(5)	427
regexp(5)	436
resource_controls(5)	443
sgml(5)	452
smf(5)	456
smf_bootstrap(5)	463
smf_method(5)	465
smf_restarter(5)	470
smf_security(5)	471
smf_template(5)	473
solaris10(5)	482
solaris(5)	485
standards(5)	489
sticky(5)	495
tecla(5)	496
term(5)	518
threads(5)	522
trusted_extensions(5)	529
vgrindefs(5)	530
wbem(5)	533
zones(5)	536

前言

无论是初次使用 SunOS 操作系统的用户还是熟悉该操作系统的用户，均可通过联机手册页获取与系统及其功能有关的信息。手册页用于简要回答关于用途/用法的问题。手册页通常是参考手册的组成部分，并不用作教程。

概述

以下内容包含对手册页各部分及其所引用信息的简要说明：

- 第 1 部分按照字母顺序介绍了操作系统中提供的各种命令。
- 第 1M 部分按照字母顺序介绍了主要用于系统维护和管理的各种命令。
- 第 2 部分介绍了所有的系统调用。其中的大多数调用可能返回一个或多个错误。如果返回一个异常值，则表明有错误情况。
- 第 3 部分介绍了各种库中包含的函数，这些函数不属于第 2 部分介绍的那些直接调用 UNIX 系统原语 (primitive) 的函数。
- 第 4 部分简要介绍了各种文件的格式，并在适当之处给出了文件格式的 C 结构声明。
- 第 5 部分包含其他文档，如字符集表。
- 第 7 部分介绍了涉及特定硬件外围设备和设备驱动程序的各种特殊文件。还介绍了 STREAMS 软件驱动程序、模块和 STREAMS 通用的一组系统调用。
- 第 9E 部分介绍了 DDI (Device Driver Interface, 设备驱动程序接口) /DKI (Driver/Kernel Interface, 驱动程序/内核接口)、仅 DDI 和仅 DKI 入口点例程，开发者可以将这些例程包含在设备驱动程序中。
- 第 9F 部分介绍了可供设备驱动程序使用的内核函数。
- 第 9S 部分介绍了驱动程序用来在驱动程序和内核之间共享信息的数据结构。

下面是手册页的通用格式。每个手册的手册页部分通常遵循该顺序，但只包括需要的标题。例如，如果未报告任何已知问题，则不包括“已知问题”部分。有关每一部分的更多详细信息，请参见 [intro](#) 页；有关手册页的更多一般信息，请参见 [man\(1\)](#)。

名称 本部分提供了记录的命令或函数的名称，后跟其用途的简要说明。

用法概要	<p>本部分说明了命令或函数的语法。如果命令或文件不存在于标准路径中，则显示其全路径名。除非要求使用不同的参数顺序，否则选项和参数均按字母顺序排列，首先是单个字母的参数，接下来是带有参数的选项。</p> <p>本部分使用以下特殊字符：</p> <ul style="list-style-type: none">[] 方括号。括在方括号中的选项或参数是可选的。如果没有方括号，则必须指定参数。... 省略号。可以为该符号前面的参数提供多个值，或者可以多次指定该参数，例如“filename...”。 分隔符。一次只能指定一个由该字符分隔的参数。{ } 大括号。括在大括号内的选项和/或参数是相互依赖的，因此必须将大括号中的所有内容视为一个单元。
协议	本部分仅出现在第 3R 子部分，用于指示协议说明文件。
描述	本部分定义了服务的功能和行为。因此，它简明地介绍了命令执行哪些操作。它不讨论“选项”或引用“示例”。在“用法”下介绍了交互式命令、子命令、请求、宏和函数。
IOctl	本部分仅出现在第 7 部分的手册页中。只有为 <code>ioctl(2)</code> 系统调用提供了适当参数的设备类才被称为 <code>ioctl</code> ，并生成自己的标题。特定设备的 <code>ioctl</code> 调用按字母顺序显示在该特定设备的手册页中。 <code>ioctl</code> 调用用于特殊类别的设备。所有这些调用都以 <code>io</code> 结尾，例如 <code>mtio(7I)</code> 。
选项	本部分列出了各命令选项及每个选项用途的简明摘要。逐个列出各个选项，并以它们在“用法概要”部分显示的顺序排列。在选项下讨论各个选项可能的参数，还提供缺省值（如果适用）。
操作数	本部分列出了命令操作数，并介绍它们对命令操作的影响。
输出	本部分介绍了命令所生成的输出（标准输出、标准错误或输出文件）。
返回值	如果手册页记录返回值的函数，则本部分列出这些值并介绍返回这些值应满足的条件。如果函数只能返回常量值（例如 0 或 -1），则将在标记的段落中列出这些值。否则，会有单个段落介绍每个函数的返回值。声明为 <code>void</code> 的函数不返回值，因此不会在“返回值”中讨论这些函数。
错误	对于故障，大多数函数将指出它们出现故障的原因的错误代码置于全局变量 <code>errno</code> 中。本部分按字母顺序列出了函数可以生成的所有错误代码，并介绍了导致每个错误的条件。如果多个条件可以导致同一错误，则在错误代码下以单独的段落介绍每个条件。
用法	本部分列出了需要详细说明了特殊规则、功能和命令。此处列出的子部分用于说明内置功能：
命令	

	修饰符 变量 表达式 输入语法
示例	本部分提供了用法的示例，或者如何使用命令或函数的示例。会尽可能显示包括命令行输入和计算机响应的完整示例。只要给出了示例，就会显示 example% 提示，如果用户必须为超级用户，则提示显示为 example# 。示例后面跟有说明、变量替换规则或返回值。大部分示例说明了“用法概要”、“说明”、“选项”和“用法”部分的概念。
环境变量	本部分列出了命令或函数影响的所有环境变量，其后附加了关于影响的简要说明。
退出状态	本部分列出了命令返回到调用程序或 shell 中的值以及导致返回这些值的条件。通常，返回零表示成功完成，返回非零值表示各种错误条件。
文件	本部分列出了手册页引用的所有文件名称、相关文件以及命令创建或所需的文件。每个文件名称后面都具有描述性摘要或说明。
属性	本部分通过定义属性类型及其相应的值列出了命令、实用程序和设备驱动程序的特征。有关更多信息，请参见 attributes(5) 。
另请参见	本部分列出了对其他手册页、内部文档和外部出版物的引用。
诊断	本部分列出了诊断消息以及导致错误的条件的简要说明。
警告	本部分列出了有关特殊条件的警告，这些条件可能会严重影响您的工作状况。此部分不是诊断列表。
附注	本部分列出了不属于页面任何部分的其他信息。它采用对用户旁白提示的形式，包含用户特别关注的要点。此处不包含关键信息。
已知问题	本部分介绍了已知问题，并尽可能给出解决方法。

参 考 文 档

简介

引用名	Intro – introduction to miscellany
描述	Among the topics presented in this section are: <div><div>Standards</div><div>The POSIX (IEEE) Standards and the X/Open Specifications are described on the standards page.</div></div> <div><div>Environments</div><div>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.</div></div> <div><div>Macros</div><div>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.</div></div> <div><div>Characters</div><div>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.</div></div>

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.

参 考 文 档

标准、环境和宏

引用名

acl – Access Control Lists

描述

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

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 descendents 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 `rxw` 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 ls(1) and stat(2) .
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 chmod(1) or setfacl(1) .
write_owner (o)	Permission to change the owner or the ability to execute chown(1) or chgrp(1) .
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:
-------------------------	-------------------------------------------------------------------------------------------------------------------------------

chmod	The chmod 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 chmod(1) for details.
compress	When a file is compressed any ACL associated with the original file is preserved with the compressed file.
cp	By default, cp ignores ACLs, unless the -p option is specified. When -p is specified the owner and group id, permission modes, modification and access times, ACLs, and extended attributes if applicable are preserved.
cpio	ACLs are preserved when the -P option is specified.
find	Find locates files with ACLs when the -acl flag is specified.
ls	By default ls does not display ACL information. When the -v option is specified, a file's ACL is displayed.
mv	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.
pack	When a file is packed, any ACL associated with the original file is preserved with the packed file.
rcp	rcp has been enhanced to support copying. A file's ACL is only preserved when the remote host supports ACLs.
tar	ACLs are preserved when the -p option is specified.
unpack	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 <code>acl_get()</code> and <code>acl_set()</code> functions. These functions support both POSIX draft ACLs and NFSv4 ACLs.
Retrieving a file's ACL	<pre>int acl_get(const char *path, int flag, acl_t **aclp); int facl_get(int fd, int flag, acl_t **aclp);</pre> <p>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 <code>ACL_NO_TRIVIAL</code> then only ACLs that are not trivial are retrieved. The ACL is returned in the <code>aclp</code> argument.</p>
Freeing ACL structure	<pre>void acl_free(acl_t *aclp);</pre> <p>The <code>acl_free()</code> function frees up memory allocated for the argument <code>aclp</code>.</p>

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.

示例

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

示例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);
}
...
```

示例 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);
}
...
```

示例 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);
...
```

示例 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);
...
```

另请参见

[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\(3SEC\)](#)

引用名	ad – Active Directory 用作命名系统信息库
描述	<p>Solaris 客户机可以从 Active Directory (AD) 服务器获取命名信息。</p> <p>Solaris 系统必须首先加入 AD 域，然后再将 <code>ad</code> 关键字添加到 <code>nsswitch.conf(4)</code> 文件的相应条目中。Solaris 系统通过使用 <code>kclient(1M)</code> 实用程序来加入 AD 域。AD 名称服务仅对 <code>passwd</code> 和 <code>group</code> 支持命名数据库。</p> <p>Windows 用户无法登录。<code>user_attr(4)</code> 数据库不包含 Windows 用户的对应条目，并且 <code>passwd(1)</code> 命令不支持与 AD 之间同步用户口令。</p> <p>Solaris AD 客户机使用自动搜索方法查找 AD 目录服务器，例如，域控制器和全局目录服务器。客户机还使用 LDAP v3 协议访问 AD 服务器中的命名信息。由于 AD 客户机与本机 AD 模式协同工作，因此无需修改 AD 服务器模式。Solaris AD 客户机使用 <code>idmap(1M)</code> 服务在 Windows 安全标识符 (security identifier, SID) 和 Solaris 用户标识符 (user identifier, UID) 及组标识符 (group identifier, GID) 之间映射。用户名和组名源自 AD 用户和组对象的 <code>sAMAccountName</code> 属性，这些名称使用对象所在的域进行标记。使用 <code>@</code> 字符分隔域名和用户名或组名。</p> <p>客户机使用 SASL/GSSAPI/KRB5 安全模型。<code>kclient</code> 实用程序用于将客户机加入到 AD。执行加入操作时，<code>kclient</code> 在客户机上配置 Kerberos v5。请参见 <code>kclient(1M)</code>。</p>
文件	<p><code>/etc/nsswitch.conf</code> 用于名称服务转换的配置文件。</p> <p><code>/etc/nsswitch.ad</code> 使用 <code>ad</code>、<code>dns</code> 和文件配置的名称服务转换的样例配置文件。</p> <p><code>/usr/lib/nss_ad.so.1</code> AD 的名称服务转换模块。</p>
另请参见	<code>passwd(1)</code> 、 <code>svcs(1)</code> 、 <code>idmap(1M)</code> 、 <code>idmapd(1M)</code> 、 <code>kclient(1M)</code> 、 <code>svcadm(1M)</code> 、 <code>svccfg(1M)</code> 、 <code>svc</code>

引用名

ascii – map of ASCII character set

用法概要

cat /usr/pub/ascii

描述

/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

文件

`/usr/pub/ascii` On-line chart of octal and hexadecimal values for the ASCII character set.

引用名

attributes, architecture, availability, CSI, stability, MT-Level, standard – 接口属性

描述

手册页的**属性**部分包含一个定义属性类型及其相应值的表。下面提供了属性表的一个示例。并非所有属性类型都适用于所有接口类型。

属性类型	属性值
体系结构	SPARC
可用性	system/kernel
CSI	Enabled（已启用）
接口稳定性	Committed（已确定）
MT 级别	Safe（安全）
标准	请参见 standards(5) 。

体系结构

体系结构定义处理器或特定硬件。请参见 [uname\(1\)](#) 的 -p 选项。在某些情况下，体系结构可能表示所需适配器或外围设备。

可用性

这表示包含本手册页介绍的命令或组件的软件包。要使用此命令，必须安装指示的软件包。有关如何添加软件包的信息，请参见 [pkg\(1\)](#)。

代码集独立性 (Code Set Independence, CSI)

不依赖于任何代码集的属性的 OS 实用程序和库具有代码集独立性 (Code Set Independence, CSI)。它们具有启用 CSI 的属性。这区别于许多仅使用扩展 Unix 代码集 (Extended Unix Codeset, EUC) 等编码方法的命令和实用程序。扩展 Unix 代码集编码方法允许同时支持最多四个代码集并且通常用于表示亚洲字符集。

不过，出于实际原因，这种独立性不是绝对的。某些假设对当前 CSI 实现依然适用：

- 文件代码是 ASCII 的超集。
- 要支持多字节字符和以 null 结尾的 UNIX 文件名，任何多字节字符不能包含 NULL 和 /（斜杠）字符。
- 仅支持“无状态”文件代码编码。无状态编码可避免移位、锁定移位、指定、调用等，但单一移位未排除在外。
- 进程代码（wchar_t 值）取决于实现，并且可随时间、实现或语言环境而发生更改。
- 并非每个对象的名称都可由任意字符组成。下列对象的名称必须由 ASCII 字符组成：
 - 用户名、组名和口令
 - 系统名称
 - 打印机和特殊设备的名称
 - 终端名称 (/dev/tty*)
 - 进程 ID 号

- 消息队列、信号量和共享内存标签。
- 下列各项可由 ISO Latin-1 或 EUC 字符组成：
 - 文件名
 - 目录名称
 - 命令名称
 - Shell 变量和环境变量名称
 - 文件系统挂载点
 - NIS 键名和域名
- NFS 共享文件的名称应由 ASCII 字符组成。尽管文件和目录的名称及内容可由非 ASCII 代码集中的字符组成，但如果仅使用 ASCII 代码集，则允许在任何计算机上挂载 NFS，而不管是否已本地化。所有启用了 CSI 的命令和实用程序均可处理在 2.6 中发布的单字节和多字节语言环境。为使应用程序获得完整的国际化服务支持，必须应用动态绑定。只有 C 和 POSIX 语言环境才支持静态绑定程序。

接口稳定性

Oracle 通常使开发人员提前就能够接触到新技术，这使得开发人员能够尽早对这些技术进行评估。遗憾的是，新技术容易发生更改，并且标准化新技术往往会导致接口与以前的版本不兼容。

为了进行合理的风险评估，开发人员必须了解接口在将来发行版中发生更改的可能性。为了帮助开发人员进行上述评估，某些手册页中提供了命令、入口点和文件格式的接口稳定性信息。

由于 Oracle 将尽力确保在将来的次要发行版中继续使用接口，因此这些接口将更加稳定，可供几乎所有应用程序安全地使用。仅依赖于已确定接口的应用程序应在将来的次要发行版中（而不一定是在早期主要发行版中）继续可靠地正常运行。

欠稳定的接口可用于进行实验和设计原型，但使用时必须了解这些接口可能会发生不兼容的更改，甚至可能在将来的次要发行版中被删除或替换为其他接口。

Oracle 未记录的“接口”（例如，大多数内核数据结构以及系统头文件中的某些符号）可能是实现工件。此类内部接口不仅会发生不兼容的更改或被删除，我们还可能不会在发行说明中提及此类更改。

发行版级别

产品具有指定的发行版级别和名称，这些有助于进行兼容性介绍。每个发行版级别还可能包括适合较低级别的更改。

发行版	版本	含义
主	x.0	可能包含增加的主要功能；遵循可能不兼容的不同标准修订；可能会更改、删除或替换已确定接口（虽然这些情况不太可能发生）。产品初始发行版通常为 1.0。

发行版	版本	含义
次	x.y	与 x.0 或早期发行版 (y!=0) 相比，此发行版可能包括：增加的功能、对已确定接口所做的兼容更改、或者可能对未确定或 Volatile（可变）接口所做的不兼容更改。
微	x.y.z	应为与上一个发行版 (z!=0) 兼容的接口，但可能修复了更多错误、改进了性能并且支持其他硬件。可能对 Volatile（可变）接口进行了不兼容的更改。

在接口稳定性上下文中，更新发行版（有时称为修补发行版）应视为等效于微发行版。

分类

下表概述了稳定性级别分类与发行版级别的关系。第一列中列出了稳定性级别。第二列中列出了不兼容的更改所对应的发行版级别，第三列中列出了其他注释。有关各分类的完整讨论，请参见下面的相应小节。

稳定性	发行版	注释
已确定	主发行版 (x.0)	极少发生不兼容情况。
未确定	次发行版 (x.y)	经常发生不兼容情况。
Volatile（可变）	微发行版 (x.y.z)	经常发生不兼容情况。

除非另行说明，否则本手册页中介绍的接口稳定性级别分类适用于源代码接口和二进制接口。所有稳定性级别分类都是公共的，但**专用**分类除外。除非明确说明，否则不会指定公共接口（即本手册页中记录的接口）的确切稳定性级别。未记录接口的稳定性级别缺省为**专用**。

除了 Solaris 产品中包含的文档以外，其他现有文档不应解释为暗指 Solaris 产品所提供接口的任何稳定性级别。Solaris 手册页是稳定性级别信息的唯一来源。

Committed（已确定）

已确定接口的用途在于使第三方能够根据这些接口开发和发布应用程序，并确信这些应用程序能够在引入接口的产品发行版的所有后续发行版（属于同一主要发行版）中正常运行。即使在主要发行版中，不兼容的更改也应当极少出现，并且应具有正当理由。

作为行业标准定义和控制的接口通常视为已确定接口。在这种情况下，属性表中的“标准”条目或其他文档位置通常会说明监管机构和/或公共文档版本。

虽然不兼容的更改很少发生，但是如果相关缺陷极其严重（如本文档的“例外情况”部分中所述），在任何发行版中都可能发生不兼容的更改；或者在次要发行版中，可能会通过“功能终止”过程来执行不兼容的更改。如果必须停止对已确定接口的支持，Oracle 将会尝试提供通知并将稳定性级别标记为“过时”。

Uncommitted (未确定)

不对这些接口在不同次要发行版中的源代码或二进制兼容性进行任何承诺。甚至在次要发行版中可能会发生接口删除等重大的不兼容更改。未确定的接口通常不适用于与发行版无关的产品。

对接口进行不兼容更改旨在对接口进行实质性的改进（包括考虑到易用性等因素）。一般情况下，未确定的接口不太可能进行不兼容的更改，如果发生此类更改，这些更改将影响甚微，并且通常具有减轻风险计划。

未确定的接口通常属于下列子类别之一：

1. 实验性或过渡性接口。这些接口通常旨在使外部开发者可以及早接触到一些新兴的、不断发展变化的技术，或者提供一个临时的解决问题的办法，有待将来再寻求更通用的解决方案。
2. 其规范由外部机构控制的接口，但 Oracle 希望在提供与外部规范同步的下一个次要发行版之前尽力与以前的发行版保持兼容。
3. 相较于稳定性而言，其目标用户更重视创新（或者易用性）的接口。此属性通常与较高层组件的管理接口相关联。

对于未确定的接口，Oracle 不会对不同次要发行版之间的源代码或二进制兼容性做任何声明。根据这些接口开发的应用程序可能无法在将来的次要发行版中运行。

Volatile (可变)

Volatile (可变) 接口可能出于任何原因而随时发生更改。

通过 Volatile (可变) 接口稳定性级别，Oracle 产品可以快速跟上不断发展变化的规范。在许多情况下，与为接口提供额外的稳定性相比，Volatile (可变) 接口能够更好地满足使用者的期望，因此人们更喜欢使用 Volatile (可变) 接口。

此分类级别最常应用于由 Oracle 以外的机构控制的接口，但与重视接口兼容性的标准机构或免费/开源软件 (Free or Open Source Software, FOSS) 社区控制的规范不同的是，无法声明极少对接口规范进行的不兼容更改。此外，此接口还适用于由 FOSS 控制的软件，对于此类软件，大家认为在最短时间内了解社区动态比向我们的客户提供稳定性更为重要。

通常还可以在可靠组织或广泛认可的组织定义接口的过程中，将 Volatile (可变) 分类级别应用于接口。这些级别通常称为标准草案。“IETF Internet 草案”就是一个广为人知的正在开发的规范的示例。

此外，实验性接口也可以是 Volatile (可变) 接口。

我们不对任何两个发行版（包括修补程序）之间的 Volatile (可变) 接口的源代码或二进制兼容性做出任何声明。包含这些接口的应用程序可能无法在将来的任何发行版中正常运行。

Not-an-Interface (不是接口)

有时会出现以下情况：推断存在的某个实体可能是一个接口，但实际上却并非接口。常见示例包括：仅供人员使用的 CLI 的输出以及 GUI 的确切布局。

此分类是一个适合用于阐明确定可能存在此类混淆的术语。如果无法对实体应用此术语，也并不意味着该实体就是某种形式的接口。它仅表明未确定可能存在此类混淆。

Private (专用)

专用接口是由组件（或产品）提供的专用于该组件的接口。专用接口仍可对其他组件可见或可由其他组件访问。由于使用其他组件的专用接口存在巨大的稳定性风险，因此明确不支持这种使用方式。并非由 Oracle Corporation 提供的组件不应使用专用接口。

大多数专用接口都未予以记录。我们很少记录专用接口。记录专用接口的原因包括（但不限于）：接口用途可能会在将来重新分类为某个公共稳定性级别分类，或者该接口会无规律地显现。

Obsolete (已过时)

“已过时”是可与上述分类级别一起显示的修饰符。“已过时”修饰符表示接口“已弃用”并且/或者建议不要继续用于一般用途。通过应用“已过时”修饰符，现有接口可从某个其他状态（例如，已确定或未确定）发生降级，以便建议客户先从该接口进行迁移，然后再删除（或以不兼容方式更改）该接口。

当前发行版支持已过时接口，但计划在将来的（次要）发行版中将其删除。停止接口支持之前，Oracle 将会首先尝试提供通知，然后再停止支持相应接口。使用已过时接口会生成警告消息。

例外情况

在极少情况下，为了维护 Oracle 和客户的最佳利益，需要违反接口稳定性承诺。下表包含接口提供者违反接口稳定性确定的常见已知原因，但不排除存在其他原因。

1. 存在安全漏洞（接口固有的漏洞）。
2. 发生数据损坏（接口固有的漏洞）。
3. 违反标准的情况（由一致性测试的解释或改进中的更改所揭示）。
4. 并非由 Oracle 控制的接口规范已发生不兼容的更改，并且大多数接口使用者都希望提供更新的接口。
5. 对于客户来说，不进行不兼容更改是无法接受的。例如，在放弃 DOS 8.3 命名限制后，如果不对 pcfs 进行不兼容的更改，就会是一个这样的示例。

例外情况允许的不兼容更改将始终尽可能在“最主要的”发行版中提供。但是，漏洞带来的后果或合同相关规定要求往往会强制在修补程序中提供。

与早期接口分类机制的兼容性

在 Solaris 10 以及先前的发行版中，采用不同的接口分类机制。下表概述了新旧分类机制之间的映射。

旧机制	新建	注释
标准	已确定	应显示标准属性类型的属性表中的条目。
稳定	已确定	名称更改。
发展中	未确定	实际承诺相符。
不稳定	未确定	名称更改。
外部	Volatile（可变）	名称更改，同时扩展了允许的使用情况。
已过时	(已过时)	以前为分类，现在为修饰符。

免费/开源软件的重要性日益提高，促使将名称从“稳定/不稳定”更改为“已确定/未确定”。“稳定”一词与该术语在 FOSS 社区中的常见用途相冲突。

“正在改进”的定义比较模糊，导致很难理解此术语。在迁移到新分类机制的过程中，以前的许多“正在改进”的接口都已升级为“已确定”。不过，在遇到术语“正在改进”时，应推断为“未确定”。

MT 级别

库分为若干类别，这些类别定义了其支持多个线程的能力。包含属于多个或不同级别的函数的手册页在**注释**或**用法**部分中对此方面进行了介绍。

Safe（安全）

“安全”是可从多线程应用程序调用的代码的属性。调入安全接口或安全代码段的作用是：即使由多个线程调用，结果仍然有效。人们常常忽视的一点是：此安全接口或安全代码段的结果可产生影响所有线程的全局后果。例如，从一个线程打开或关闭文件的操作对进程中的所有线程都可见。多线程应用程序负责安全地使用这些接口，这与此接口是否安全有所不同。例如，关闭应用程序中其他线程仍在使用的文件的多线程应用程序未安全地使用 `close(2)` 接口。

Unsafe（非安全）

非安全库包含不受保护的全局和静态数据。除非应用程序安排每次仅在库中执行一个线程，否则使用此库会不安全。非安全库可能包含安全函数；不过，库包含的大多数函数在调用时都是不安全的。某些非安全函数具有多线程安全的可重入函数。可重入函数由附加到函数名称的 `_r` 后缀指定。

MT-Safe（MT 安全）

多线程安全库是为多线程访问而充分准备的库。它通过锁定保护其全局和静态数据，并且可提供合理数目的并发性。可以安全使用的库并不能视为多线程安全。例如，使用监视器监视整个库可使库保持安全，但它不支持并发性，因此不能视为多线程安全。多线程安全库必须允许合理数目的并发性。（此定义的目的是精确定义安全库的含义。安全库的定义不会指定该库是否支持并发性。多线程安全定义明确指明该库是安全的，并且支持一定程度的并发性。这阐明了安全定义，它可以表示从单线程到任何并发度的多线程的所有内容。）

Async-Signal-Safe (异步信号安全)

“异步信号安全”表示可从信号处理程序安全地调用的特定库函数。执行异步信号安全函数的线程在被信号中断时，自身不会发生死锁。信号只会为获取锁定的多线程安全函数带来问题。

异步信号安全函数也具有多线程安全性。在异步信号安全函数中获取锁定时，将禁用信号。这些信号用于防止调用可能获取相同锁定的信号处理程序。

MT-Safe with Exceptions (MT 安全，但存在异常)

有关异常的说明，请参见这些页面的**注释**或**用法**部分。

Safe with Exceptions (安全，但存在异常)

有关异常的说明，请参见这些页面的“附注”或“用法”部分。

Fork-Safe (Fork 安全)

fork(2) 函数仅在子进程中复制调用线程。**fork1(2)** 函数的存在目的是为了与以前版本兼容，它与 **fork()** 同义。当调用 **fork()** 时，如果未在执行派生的其他线程保持锁定，该锁定仍将保持在子进程中，但是由于未复制所属线程，因此没有锁定所有者。调用尝试获取锁定的函数的子进程自身将发生死锁。

当调用 **fork()** 时，Fork 安全库安排仅让执行派生的线程保留该库的所有内部锁定。这通常是使用 **pthread_atfork(3C)** 实现的，该函数在初始化库时调用。

在极少情况下，如果进程需要在执行派生时复制其所有线程，**forkall(2)** 函数会提供此功能。调用 **forkall()** 时，不会执行 **pthread_atfork()** 操作。调用 **forkall()** 存在相应的危险。当某个线程调用 **forkall()** 时，如果进程中的某些其他线程正在执行 I/O 操作，这些线程将继续在父进程和子进程中执行相同的 I/O 操作，这可能会导致数据损坏。出于此原因以及其他竞争情况原因，不建议使用 **forkall()**。

在 Solaris 10 之前的所有 Solaris 发行版中，**fork()** 的行为取决于应用程序是否与 **-lpthread** 相链接（有关 POSIX 线程，请参见 **standards(5)**）。如果与 **-lpthread** 链接，**fork()** 的行为与 **fork1()** 相似，否则与 **forkall()** 相似。为了避免产生有关 **fork()** 行为的任何混淆，应用程序可以根据需要明确地调用 **fork1()** 或 **forkall()**。

Cancel-Safety (取消安全)

如果多线程应用程序使用 **pthread_cancel(3C)** 取消（即中止）线程，目标线程在中止时可能会保留某项资源，例如锁定或分配的内存。如果线程未安装有适当的取消清除处理程序来释放相应资源（请参见 **pthread_cancel(3C)**），该应用程序即为“取消不安全”，也就是说，从线程取消方面来说，该应用程序不安全。由于取消的线程未释放锁定，这种非安全性可能导致死锁或资源泄漏；例如，不会在取消线程时释放内存。使用 **pthread_cancel(3C)** 的所有应用程序都应确保它们在“取消安全”环境中运行。此外，如果库具有取消点并且获取锁定等资源或动态分配内存，也会导致与这些库关联的应用程序的取消不安全性。这为多线程程序中的库引入了另一个安全级别：取消安全。取消安全包含两个子类别：延迟取消安全以及异步取消安全。如果应用程序对于取消类型为 **PTHREAD_CANCEL_DEFERRED** 的线程为取消安全时，该应用程序被视为延迟取消安全。如果应用程序对于取消类型为

`PTHREAD_CANCEL_ASYNCHRONOUS` 的线程为取消安全时，该应用程序被视为异步取消安全。由于具有延迟取消类型的线程只能在正确定义的取消点取消，而具有异步取消类型的线程可在任意位置取消，因此延迟取消安全比异步取消安全更容易实现。缺省情况下，创建的所有线程都具有延迟取消类型，因此可能永远不需要担心异步取消安全。大多数应用程序和库都应当始终为异步取消不安全。根据定义，异步取消安全的应用程序同时也是延迟取消安全的。

标准

许多接口都作为行业标准进行定义和控制。在这种情况下，本部分中将说明监管机构和/或公共文档版本。

程序员在生成可移植应用程序时，应该遵照此应用程序应符合的标准或规范中提供的接口说明，而不能遵照基于公共标准的接口的手册页说明。当标准或规范允许备用实现选项时，本手册页通常仅介绍由 Oracle 实现的备用选项。此外，本手册页还会介绍由 Oracle 提供的标准接口的基本定义的所有兼容扩展。

对于文中引用的监管机构或文档，并不意味着我们将其认可为“标准”条目。监管机构可以是非常正式的组织（例如 ISO 或 ANSI）、较不正式但广泛接受的组织（例如 IETF）或非正式的独立贡献者（例如 FOSS（Free or Open Source Software，免费/开源软件）贡献者）。

另请参见

[uname\(1\)](#)、[Intro\(3\)](#)、[standards\(5\)](#)

[pkg\(1\)](#)

引用名	audit_binfile – generation of Solaris audit logs
用法概要	<code>/usr/lib/security/audit_binfile.so</code>
描述	<p>The <code>audit_binfile</code> plugin module for Solaris audit, <code>/usr/lib/security/audit_binfile.so</code>, 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).</p> <p>The <code>audit_binfile</code> plugin is loaded by <code>auditd</code> if the plugin is configured as an active via <code>auditconfig</code>. Use the <code>auditconfig -setplugin</code> option to change all the plugin related configuration parameters.</p>
Object Attributes	<p>The following attributes specify the configuration of <code>audit_binfile</code> plugin:</p> <p><code>p_dir</code></p> <p><code>dir1[,dir2],... [,dirn]</code></p> <p>A list of directories, where the audit files will be created. Any valid writable directory can be specified.</p> <p><code>p_minfree</code></p> <p>A percentage, which indicates the amount of free space required on the target <code>p_dir</code>. 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> <p><code>p_fsize</code></p> <p>The <code>p_fsize</code> 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 <code>audit -ncommand</code> 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 <code>p_fsize</code> 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.</p>
示例	<p>The following directives cause <code>audit_binfile.so</code> to be loaded, specify the directories for writing audit logs, and specify the percentage of required free space per directory.</p> <pre>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"</pre>
属性	See attributes(5) for a description of the following attributes:

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

另请参见

[auditconfig\(1M\)](#), [auditd\(1M\)](#), [audit_warn\(1M\)](#), [syslog.conf\(4\)](#), [attributes\(5\)](#)
[《Oracle Solaris 管理：安全服务》](#)

引用名

audit_flags – audit preselection flags

描述

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.

示例

示例 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
```

示例 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
```

另请参见

[profiles\(1\)](#), [auditconfig\(1M\)](#), [auditd\(1M\)](#), [usermod\(1M\)](#), [audit_class\(4\)](#),
[audit_event\(4\)](#), [prof_attr\(4\)](#), [user_attr\(4\)](#)

《Oracle Solaris 管理：安全服务》

引用名	audit_remote – send Solaris audit logs to a remote server
用法概要	/usr/lib/security/audit_remote.so
描述	<p>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).</p> <p>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.</p>
Object Attributes	<p>The following attributes specify the configuration of audit_remote plugin:</p> <p>p_hosts</p> <pre>host1[:[port1][:mech1]][,host2[:[port2][:mech2]],... \ hostn[:[portn][:mechn]]]</pre> <p>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.</p> <p>The <i>host</i> part of a p_hosts entry can be in any form acceptable to getipnodebyname(3SOCKET).</p> <p>The <i>port</i> 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).</p> <p>The <i>mech</i> 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 <code>kerberos_v5</code>.</p> <p>p_retries</p> <p>The number of retries for connecting to and sending data to a server.</p> <p>The default value is 3.</p> <p>p_timeout</p> <p>The number of seconds in which a connection/sending data timeouts.</p> <p>The default value is 5 seconds.</p> <p>qsize</p> <p>The maximum number of outstanding audit records to keep.</p> <p>The default is the value of the kernel queue control high water mark. See auditconfig(1M).</p> <p>If set to 0, the default value is the value of the kernel queue control high water mark. See auditconfig(1M).</p>

GSS SESSION

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

示例

示例 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"
```

示例 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"
```

属性 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.

另请参见 [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\(4\)](#), [attributes\(5\)](#), [kerberos\(5\)](#), [tcp\(7P\)](#)

附注 `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.

引用名	audit_syslog – realtime conversion of Solaris audit data to syslog messages														
用法概要	/usr/lib/security/audit_syslog.so														
描述	<p>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.</p> <p>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.</p> <p>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.</p> <p>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:</p> <pre>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</pre> <p>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.</p> <table><tr><td>ok or failed</td><td>Comes from the return or exit token.</td></tr><tr><td>session <#></td><td><#> is the session ID from the subject token.</td></tr><tr><td>by <name></td><td><name> is the audit ID from the subject token.</td></tr><tr><td>as <name>:<group></td><td><name> is the effective user ID and <group> is the effective group ID from the subject token.</td></tr><tr><td>in <zone name></td><td>The zone name. This field is generated only if the zonename audit policy is set.</td></tr><tr><td>from <terminal></td><td><terminal> is the text machine address from the subject token.</td></tr><tr><td>obj <path></td><td><path> 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 (...).</td></tr></table>	ok or failed	Comes from the return or exit token.	session <#>	<#> is the session ID from the subject token.	by <name>	<name> is the audit ID from the subject token.	as <name>:<group>	<name> is the effective user ID and <group> is the effective group ID from the subject token.	in <zone name>	The zone name. This field is generated only if the zonename audit policy is set.	from <terminal>	<terminal> is the text machine address from the subject token.	obj <path>	<path> 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 (...).
ok or failed	Comes from the return or exit token.														
session <#>	<#> is the session ID from the subject token.														
by <name>	<name> is the audit ID from the subject token.														
as <name>:<group>	<name> is the effective user ID and <group> is the effective group ID from the subject token.														
in <zone name>	The zone name. This field is generated only if the zonename audit policy is set.														
from <terminal>	<terminal> is the text machine address from the subject token.														
obj <path>	<path> 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 (...).														

`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.

示例

示例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"
```

示例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"
```

属性

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.

另请参见

[auditconfig\(1M\)](#), [auditd\(1M\)](#), [audit_class\(4\)](#), [syslog.conf\(4\)](#), [user_attr\(4\)](#), [attributes\(5\)](#), [audit_flags\(5\)](#)

《Oracle Solaris 管理：安全服务》

附注

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.

引用名	brands – 非全局区域的备用操作环境
描述	<p>标记区域 (Branded Zone, BrandZ) 框架扩展了 zones(5) 中介绍的 Solaris Zones 基础结构，在其中包含标记创建功能，这些标记提供包含非本机操作环境的非全局区域。</p> <p>术语“标记”可以指各种操作环境。所有标记管理都通过扩展当前区域结构来执行。</p> <p>每个区域都配有一个关联标记。标记类型用来确定安装和引导区域时执行的脚本。此外，区域的标记还可用来在应用程序启动时确定正确的应用程序类型。缺省标记由全局区域中安装的分发软件决定。</p> <p>一个标记区域只支持一个非本地二进制命令标记，这意味着一个标记区域只提供一种操作环境。一旦为区域分配了标记，便不能更改或删除该标记。</p> <p>BrandZ 通过以下方式扩展区域工具：</p> <ul style="list-style-type: none">■ 标记是区域的一个属性，在创建区域时设置。■ <code>zonecfg</code> 工具（参见 zonecfg(1M)）用于设置区域的标记类型并配置区域。■ <code>zoneadm</code> 工具（参见 zoneadm(1M)）用于报告区域的标记类型并管理区域。
设备支持	每个区域支持的设备都记录在该标记的手册页和其他文档中。区域基础结构检测任何尝试添加不支持设备的操作并向管理员发出警告。如果管理员不顾警告而选择添加不支持的设备，该设备可能无法如预期那样工作。该配置将会是未经过测试的，并且不受支持。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	system/zones
接口稳定性	Committed（已确定）

另请参见 [mdb\(1\)](#)、[zlogin\(1\)](#)、[zonename\(1\)](#)、[dtrace\(1M\)](#)、[in.rlogind\(1M\)](#)、[sshd\(1M\)](#)、[zoneadm\(1M\)](#)、[z](#)

引用名

cancellation – overview of concepts related to POSIX thread cancellation

描述

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)`, `aio_suspend(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)`, `writew(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)	Disabled

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).

示例

示例 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.)

示例1 Cancellation example (续)

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";
}
```

示例1 Cancellation example (续)

```
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));
}
```

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

ATTRIBUTETYPE	ATTRIBUTEVALUE
MT-Level	MT-Safe

另请参见 [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\(5\)](#)

引用名

charmap – character set description file

描述

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>		<ETB>	<IS1>	<RS>	<SYN>
<CR>	<DLE>	<ETX>	<IS2>	<SI>	<US>
<DC1>		<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.

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

Format

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

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

or

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

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

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

A character set mapping definition line must exist for all symbolic names and must define the coded character value that corresponds to the character glyph indicated in the table, or the coded character value that corresponds with the control character symbolic name. If the control characters commonly associated with the symbolic names are supported by the implementation, the symbolic name and the corresponding encoding value must be included in the file. Additional unique symbolic names may be included. A coded character value can be represented by more than one symbolic name.

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

"%cd%d", <escape_char>, <decimal byte value>

"%cX%X", <escape_char>, <hexadecimal byte value>

"%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 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.

另请参见

[locale\(1\)](#), [localedef\(1\)](#), [nl_langinfo\(3C\)](#), [extensions\(5\)](#), [locale\(5\)](#)

引用名

condition – concepts related to condition variables

描述

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.

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.

Initialize

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 [mmap\(2\)](#)) and are initialized for this purpose.

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.

A condition variable must not be simultaneously initialized by multiple threads or re-initialized while in use by other threads.

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, *type*, passed to their initialization function.

Condition Wait

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.

- 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.

属性

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

另请参见

[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\(5\)](#)

附注

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.

引用名

crypt_bsdbf – password hashing module using Blowfish cryptographic algorithm

用法概要

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

描述

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
```

属性

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

另请参见

[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\(5\)](#)

引用名 crypt_bsmd5 – password hashing module using MD5 message hash algorithm

用法概要 /usr/lib/security/\$ISA/crypt_bsmd5.so

描述 The crypt_bsmd5 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_bsmd5 is 255 characters.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

另请参见 [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\(5\)](#)

引用名

crypt_sha256 – password hashing module using SHA–256 message hash algorithm

用法概要

/usr/lib/security/\$ISA/crypt_sha256.so

描述

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.

属性

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT-Level	Safe

另请参见

`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)`

引用名	crypt_sha512 – password hashing module using SHA–512 message hash algorithm						
用法概要	/usr/lib/security/\$ISA/crypt_sha512.so						
描述	<p>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.</p> <p>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.</p> <p>The maximum password length for crypt_sha512 is 255 characters.</p> <p>The following options can be passed to the module by means of crypt.conf(4):</p> <p>rounds=<positive_number></p> <p>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.</p> <p>The number of additional rounds is stored in the salt string returned by crypt_gensalt(3C). For example:</p> <p>\$6,rounds=6000\$nlxmTtpz\$</p> <p>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.</p>						
属性	See attributes(5) for descriptions of the following attributes:						
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Interface Stability</td><td>Committed</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Committed	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Committed						
MT-Level	Safe						
另请参见	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)						

引用名

crypt_sunmd5 – password hashing module using MD5 message hash algorithm

用法概要

/usr/lib/security/\$ISA/crypt_sunmd5.so

描述

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:

`$md5,rounds=1000$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.

属性

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

另请参见

`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)`

引用名	crypt_unix – traditional UNIX crypt algorithm
描述	<p>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.</p> <p>The algorithm identifier for policy.conf(4) is <code>__unix__</code>. There is no entry in crypt.conf(4) for this algorithm.</p> <p>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 "\$".</p> <p>This algorithm is based on a one-way encryption algorithm with variations intended (among other things) to frustrate use of hardware implementations of a key search. Only the first eight characters of the key passed to <code>crypt()</code> are used with this algorithm; the rest are silently ignored. The salt is a two-character string chosen from the set [a-zA-Z0-9./]. This string is used to perturb the hashing algorithm in one of 4096 different ways.</p> <p>The maximum password length for crypt_unix is 8 characters.</p>
用法	The return value of the crypt_unix algorithm might not be portable among standard-conforming systems. See standards(5) .
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

另请参见	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)
------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

引用名 device_clean – device clean programs

描述 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`.

- audio_clean audio devices
- st_clean tape devices
- sr_clean CD-ROM devices

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

- disk_clean CD-ROM and other removable media devices. This program mounts the device during the execution of `allocate`, if required.
- audio_clean_wrapper wrapper to make audio_clean work with CDE
- wdwrapper wrapper to make other cleaning programs work with CDE
- wdwmsg 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:

- device-name* The name of the device that is to be cleaned. Use `list_devices` to obtain the list of allocatable devices.
- i Invoke boot-time initialization.
- f Force cleanup by the administrator.
- s Invoke standard cleanup by the user.
- I Same as -i, 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 “/”.

退出状态

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.

文件

/etc/security/lib/* device clean programs

属性

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.

另请参见

[allocate\(1\)](#), [deallocate\(1\)](#), [list_devices\(1\)](#), [attributes\(5\)](#)

《Oracle Solaris 管理：安全服务》

引用名	dhcp – Dynamic Host Configuration Protocol					
描述	<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, dhcpcagent(1M).</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 dhcpinfo(1) 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 ifconfig(1M) utility. Check the status of the daemon using the netstat(1M) and ifconfig(1M) commands.</p>					
Oracle Solaris DHCP Server	<p>The Oracle Solaris DHCP server is implemented as a background daemon, in.dhcpd(1M). This daemon can deliver network configuration information to either BOOTP or DHCP clients. The Oracle Solaris DHCP service can be managed using the dhcpcmgr(1M) GUI or the command line utilities dhcpcconfig(1M), dhtadm(1M), and pntadm(1M).</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><tr><td>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 dhcptab for the DHCP service. The dhcptab(4) can be viewed and modified using the dhtadm(1M) command or dhcpcmgr(1M) graphical utility. See dhcptab(4) for more information about the syntax of dhcptab records. See dhcp_inittab(4) for more information about the DHCP options and symbols.</td></tr><tr><td>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 dhcptab for the DHCP service. The dhcptab(4) can be viewed and modified using the dhtadm(1M) command or dhcpcmgr(1M) graphical utility. See dhcptab(4) for more information about the syntax of dhcptab records. See dhcp_inittab(4) 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 dhcptab for the DHCP service. The dhcptab(4) can be viewed and modified using the dhtadm(1M) command or dhcpcmgr(1M) graphical utility. See dhcptab(4) for more information about the syntax of dhcptab records. See dhcp_inittab(4) 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 `dhcpgmr` graphical utility. See [dhcp_network\(4\)](#) for more information about network tables.

另请参见

[dhcpinfo\(1\)](#), [dhcpgagent\(1M\)](#), [dhcpconfig\(1M\)](#), [dhcpgmr\(1M\)](#), [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.

引用名	dhcp_modules – data storage modules for the DHCP service
描述	<p>This man page describes the characteristics of data storage modules (public modules) for use by the Solaris Dynamic Host Configuration Protocol (DHCP) service.</p> <p>Public modules are the part of the DHCP service architecture that encapsulate the details of storing DHCP service data in a data storage service. Examples of data storage services are Oracle and ufs file systems.</p> <p>Public modules are dynamic objects which can be shipped separately from the Solaris DHCP service. Once installed, a public module is visible to the DHCP service, and can be selected for use by the service through the DHCP service management interfaces (dhcpgmgr(1M), dhcpconfig(1M), dhtadm(1M), and pntadm(1M)).</p> <p>Public modules may be provided by Sun Microsystems, Inc or by third parties.</p> <p>The Solaris DHCP service management architecture provides a mechanism for plugging in public module-specific administration functionality into the dhcpgmgr(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.</p> <p>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.</p>
SUNWfiles	<p>This module stores its data in ASCII files. Although the format is ASCII, hand-editing is discouraged. It is useful for DHCP service environments that support several hundred to a couple thousand of clients and lease times are a few hours or more.</p> <p>This module's data may be shared between DHCP servers through the use of NFS.</p>
SUNWbinfiles	<p>This module stores its data in binary files. It is useful for DHCP service environments with many networks and many thousands of clients. This module provides an order of magnitude increase in performance and capacity over SUNWfiles.</p> <p>This module's data cannot be shared between DHCP servers.</p>
另请参见	<p>crontab(1), dhcpconfig(1M), dhcpgmgr(1M), dhtadm(1M), pntadm(1M), dhcpsvc.conf(4), dhcp(5)</p> <p>《Oracle Solaris DHCP Service Developer's Guide》</p>

引用名

environ – 用户环境

描述

当进程开始执行时，`exec` 函数系列中的一个函数将激活一个字符串数组，该字符串数组被称为环境；参见 [exec\(2\)](#)。根据约定，这些字符串的形式为 `variable=value`。例如，`PATH=/sbin:/usr/sbin`。这些环境变量提供了使程序环境相关信息可用于程序的方法。

可以通过 [sh\(1\)](#) 中的 `export` 命令和 `name=value` 参数或某一个 `exec` 函数将名称放入环境中。某些 shell 变量（例如，`MAIL`、`PS1`、`PS2` 和 `IFS`）经常被 `.profile` 文件导出，与这些变量发生冲突是不可取的；参见 [profile\(4\)](#)。

以下环境变量可由应用程序使用，并且应在目标运行时环境中设置。

HOME

用户登录目录的名称，由 [login\(1\)](#) 通过口令文件设置；参见 [passwd\(4\)](#)。

LANG

用于指定国际化信息的字符串，通过国际化信息，用户可以使用不同的国家约定。[setlocale\(3C\)](#) 函数检查 `LANG` 环境变量（当通过 `""` 将它作为 `locale` 参数调用时）。如果特定类别的相应环境变量未设置或为 `Null`，`LANG` 将作为缺省语言环境使用。如果 `LC_ALL` 设置为有效的非空值，将使用其内容覆盖 `LANG` 和其他 `LC_*` 变量。例如，在以 `setlocale(LC_CTYPE, "")` 形式调用时，`setlocale()` 将首先查询 `LC_CTYPE` 环境变量来查看它是否已设置并且不为 `Null`。如果 `LC_CTYPE` 未设置或者为 `Null`，则 `setlocale()` 将检查 `LANG` 环境变量来查看它是否已设置并且不为 `Null`。如果 `LANG` 和 `LC_CTYPE` 均未设置或为 `NULL`，将使用缺省的 "C" 语言环境来设置 `LC_CTYPE` 类别。

大部分命令都会在执行任何其他处理之前调用 `setlocale(LC_ALL, "")`。这样，通过设置相应环境变量，可以将命令与不同的国家约定一起使用。

以下环境变量对应于 [setlocale\(3C\)](#) 的每个类别：

LC_ALL

如果设置为有效的非空字符串值，将覆盖 `LANG` 以及所有其他 `LC_*` 变量的值。

LC_COLLATE

此类别指定所使用字符排序规则序列。此类别对应的信息存储在 [localedef\(1\)](#) 命令创建的数据库中。此环境变量影响 [strcoll\(3C\)](#) 和 [strxfrm\(3C\)](#)。

LC_CTYPE

此类别指定字符分类、字符转换以及多字节字符的宽度。当 `LC_CTYPE` 设置为有效值时，调用实用程序可以显示并处理包含下列字符的文本和文件名：该语言环境的有效字符；扩展 Unix 编码 (Extended Unix Code, EUC) 字符，其中任何单个字符的宽度可以为 1、2 或 3 个字节；以及 1、2 或 3 个列宽的 EUC 字符。缺省 "C" 语言环境对应于 7 位 ASCII 字符集；只有 ISO 8859-1 字符有效。此类别对应的信息存储在 [localedef\(1\)](#) 命令创建的数据库中。此环境变量由 [ctype\(3C\)](#)、[mblen\(3C\)](#) 以及许多其他命令（例如，[cat\(1\)](#)、[ed\(1\)](#)、[ls\(1\)](#) 和 [vi\(1\)](#)）使用。

LC_MESSAGES

此类别指定所使用消息数据库的语言。例如，应用程序的一个消息数据库可能包含法语消息，而另一个数据库可能包含德语消息。消息数据库由 `mkmsgs(1)` 命令创建。此环境变量由 `exstr(1)`、`gettext(1)`、`srchtxt(1)`、`gettext(3C)` 以及 `gettext(3C)` 使用。

LC_MONETARY

此类别指定用于特定语言环境的货币符号和分隔符。此类别对应的信息存储在 `localedef(1)` 命令创建的数据库中。此环境变量由 `localeconv(3C)` 使用。

LC_NUMERIC

此类别指定小数点分隔符和千分位分隔符。此类别对应的信息存储在 `localedef()` 命令创建的数据库中。在缺省的 C 语言环境中，使用 "." 作为小数点分隔符，不使用千分位分隔符。此环境变量由 `localeconv(3C)`、`printf(3C)` 以及 `strtod(3C)` 使用。

LC_TIME

此类别指定日期和时间格式。此类别对应的信息存储在 `localedef()` 指定的数据库中。在缺省的 C 语言环境中，使用美国的日期和时间格式。此环境变量由许多命令和函数使用。例如：`at(1)`、`calendar(1)`、`date(1)`、`strftime(3C)` 以及 `getdate(3C)`。

MSGVERB

控制在向 `stderr` 显示消息时 `fmtmsg` 选择哪些标准格式消息组件；参见 `fmtmsg(1)` 和 `fmtmsg(3C)`。

NETPATH

冒号分隔的网络标识符列表。网络标识符是系统的网络选择组件使用的字符串，用于提供应用程序特定的缺省网络搜索路径。网络标识符必须由非 Null 字符构成且长度至少为 1。不指定最大长度。网络标识符通常由系统管理员选择。网络标识符还是任意 `/etc/netconfig` 文件条目中的第一个字段。因此，`NETPATH` 提供了指向 `/etc/netconfig` 文件的链接以及有关该网络条目中包含的网络的信息。`/etc/netconfig` 由系统管理员维护。`getnetpath(3NSL)` 中介绍的库例程可以访问 `NETPATH` 环境变量。

NLSPATH

包含一个模板序列，当 `catopen(3C)` 和 `gettext(3C)` 尝试定位消息目录时将使用这些模板。每个模板都由一个可选前缀、一个或多个替换字段、一个文件名以及一个可选后缀组成。例如：

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

定义 `catopen()` 应搜索目录 `/system/nlslib` 中的所有消息目录，其中目录名由传递到 `catopen()` 的 `name` 参数（即 %N）以及 `.cat` 后缀构成。

替换字段包含一个 % 符号，后面跟一个单字母关键字。当前定义了下列关键字：

%N

传递到 `catopen()` 的 `name` 参数的值。

%L

LANG 或 LC_MESSAGES 的值。

%l

LANG 或 LC_MESSAGES 中的语言元素。

%t

LANG 或 LC_MESSAGES 中的地区元素。

%c

LANG 或 LC_MESSAGES 中的代码集元素。

%%

单个 % 字符。

如果指定的值当前未定义，将替换为空字符串。分隔符 "_" 和 "." 不包含在 %t 和 %c 替换中。

NLSPATH 中定义的模板由冒号 (:) 分隔。前导冒号或两个相邻的冒号 (::) 相当于指定 %N。例如：

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

指示 catopen() 应搜索 *name*、*name.cat* 和 /nlslib/\$LANG/*name.cat* 中的请求的消息目录。对于 gettext(), %N 将自动映射到 messages。

如果 NLSPATH 未设置或为 NULL，catopen() 和 gettext() 将调用 [setlocale\(3C\)](#)，后者将检查 LANG 和 LC_* 变量以定位消息目录。

Solaris 中对 %L 的扩展解释包括对已接受语言环境名称别名的支持，如 [gettext\(1\)](#)、[gettext\(3C\)](#)、[catopen\(3C\)](#)、[setlocale\(3C\)](#) 和 [locale_alias\(5\)](#) 中所述。

NLSPATH 通常在系统范围内设置（在 /etc/profile 中），这样可以使与消息目录相关的位置和命名约定对程序和用户都是透明的。

PATH

[sh\(1\)](#)、[time\(1\)](#)、[nice\(1\)](#)、[nohup\(1\)](#) 和其他实用程序在按不完整的路径名搜索文件时应用的目录前缀序列。前缀由冒号 (:) 分隔。[login\(1\)](#) 用于设置 PATH=/usr/bin。有关详细信息，请参见 [sh\(1\)](#)。

SEV_LEVEL

定义严重性级别，并且在标准格式错误消息中将字符串与严重性级别相关联并输出其字符串；参见 [addseverity\(3C\)](#)、[fmtmsg\(1\)](#) 以及 [fmtmsg\(3C\)](#)。

TERM

要为其准备输出的终端的种类。此信息由可能利用该终端的特殊功能的命令（例如 [vi\(1\)](#)）使用。

TZ

时区信息。此环境变量的内容由函数 `ctime(3C)`、`localtime(3C)`、`strftime(3C)` 以及 `mktime(3C)` 用来覆盖缺省时区。TZ 值为以下两种格式（为清晰可见，插入了空格）之一：

```
:characters
```

或

```
std offset dst offset, rule
```

如果 TZ 属于第一种格式（即如果第一个字符为冒号(:)）或者 TZ 不属于第二种格式，则 TZ 将指定时区数据库文件的路径（相对于 `/usr/share/lib/zoneinfo/`），如果存在前导冒号，将忽略它。

否则，如果 TZ 属于第二种格式，其展开形式如下：

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

std 和 *dst*

指示作为标准 (*std*) 时区或备用 (*dst*，例如夏时制) 时区指定的字符串，不少于三个字节且不超过 `{TZNAME_MAX}`。只有 *std* 是必需的；如果缺少 *dst*，则在此时区中不应用备用时间。其中每个字段都以两种格式之一出现，即括起或不括起：

- 在括起的格式中，第一个字符是小于号('<') 字符，最后一个字符是大于号('>') 字符。两个尖括号字符之间的所有字符都是当前语言环境的可移植字符集中的字母数字字符、加号('+') 字符或减号('-') 字符。这种情况下，*std* 和 *dst* 字段不包含尖括号字符。
- 在不括起的格式中，这些字段中的所有字符都是当前语言环境的可移植字符集中的字母字符。

如果任意字段的长度小于三个字节（缺少 *dst* 的情况除外）、大于 `{TZNAME_MAX}` 个字节或者包含指定字符之外的其他字符，则这些字段的解释是未明确指定的。

offset

指示本地时间与世界标准时间之间的差值。此时差的格式如下：

```
hh[:mm[:ss]]
```

分钟 (*mm*) 和秒 (*ss*) 是可选的。小时 (*hh*) 是必需的，可以是一位数。*offset*（放在 *std* 后面）是必需的。如果 *offset* 未出现在 *dst* 后面，将假定夏时制时间比标准时间早一小时。可以使用一位数或多位数。此值始终解释为十进制数字。小时必须介于 0 到 24 之间，如果出现分钟（和秒），分钟（和秒）必须介于 0 到 59 之间。如果值超出范围，可能会导致不可预测的行为。如果开头带有 -，表示时区在本初子午线的东侧。否则，时区位于本初子午线的西侧（通过开头的 "+" 符号表示，该符号是可选的）。

start/time, end/time

指示何时进行夏时制切换，其中，*start/time* 说明何时从标准时间更改为夏时制时间，而 *end/time* 说明何时从夏时制时间更改回标准时间。每个 *time* 字段都用本地时间说明发生更改的时间。

start 和 *end* 为以下格式之一：

Jn

儒略日 n ($1 \leq n \leq 365$)。不计算闰日。也就是说，在所有年份中，2月28日是第59日，3月1日是第60日。无法表示特殊的2月29日。

n

从零开始计算的儒略日 ($0 \leq n \leq 365$)。计算闰日，而且可以表示2月29日。

Mm.n.d

年度的第 m 月的第 n 周 ($1 \leq n \leq 5, 1 \leq m \leq 12$) 第 d 日 ($0 \leq d \leq 6$)，其中，第5周表示“第 m 月中的最后 d 日”，它可能出现在第四周或第五周。第1周是第 d 日后的第一周。第零日是周日。

实现特定的缺省值用于 *start* 和 *end*（如果未指定这些可选字段）。

time 的格式与 *offset* 相同，但前者不允许使用前导符号（“-”或“+”）。如果未指定 *time*，缺省值为 02:00:00。

另请参见

[cat\(1\)](#)、[date\(1\)](#)、[ed\(1\)](#)、[fmtmsg\(1\)](#)、[localedef\(1\)](#)、[login\(1\)](#)、[ls\(1\)](#)、[mkmsgs\(1\)](#)、[nice\(1\)](#)、[no](#)

引用名

eqnchar – special character definitions for eqn

用法概要

eqn /usr/share/lib/pub/eqnchar filename | troff options

neqn /usr/share/lib/pub/eqnchar filename | troff options

描述

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>	\bigcirc
<i>wig</i>	\sim	<i>rangle</i>	\rangle	<i>blot</i>	\blacksquare
<i>-wig</i>	\approx	<i>hbar</i>	\hbar	<i>bullet</i>	\bullet
<i>>wig</i>	\gtrsim	<i>ppd</i>	\perp	<i>prop</i>	\propto
<i><wig</i>	\lesssim	<i><-></i>	\leftrightarrow	<i>empty</i>	\emptyset
<i>=wig</i>	\equiv	<i><=></i>	\Leftrightarrow	<i>member</i>	\in
<i>star</i>	$*$	<i> <</i>	\lt	<i>nomem</i>	\notin
<i>bigstar</i>	\ast	<i> ></i>	\gt	<i>cup</i>	\cup
<i>=dot</i>	$\dot{=}$	<i>ang</i>	\angle	<i>cap</i>	\cap
<i>orsign</i>	\vee	<i>rang</i>	\sphericalangle	<i>incl</i>	\sqsubseteq
<i>andsign</i>	\wedge	<i>3dot</i>	\vdots	<i>subset</i>	\subset
<i>=del</i>	$\overset{\Delta}{=}$	<i>thf</i>	\therefore	<i>supset</i>	\supset
<i>oppA</i>	∇	<i>quarter</i>	$\frac{1}{4}$	<i>!subset</i>	\subseteq
<i>oppE</i>	\equiv	<i>3quarter</i>	$\frac{3}{4}$	<i>!supset</i>	\supseteq
<i>angstrom</i>	\AA	<i>degree</i>	$^{\circ}$		

文件

/usr/share/lib/pub/eqnchar

属性

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	text/doctools

另请参见

eqn(1), **nroff(1)**, **troff(1)**, **attributes(5)**

引用名	extendedFILE – enable extended FILE facility usage				
用法概要	<pre>\$ ulimit -n <i>N_file_descriptors</i> \$ LD_PRELOAD_32=/usr/lib/extendedFILE.so.1 <i>application</i> [<i>arg...</i>]</pre>				
描述	<p>The extendedFILE.so.1 is not a library but an enabler of the extended FILE facility.</p> <p>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).</p> <p>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</p> <pre>\$ ulimit -n <i>max_file_descriptors</i> \$ LD_PRELOAD_32=/usr/lib/extendedFILE.so.1 <i>application</i> [<i>arg...</i>]</pre> <p>where <i>max_file_descriptors</i> is the maximum number of file descriptors desired. See limit(1). The maximum value is the same as the maximum value for open(2).</p>				
环境变量	<p>The following environment variables control the behavior of the extended FILE facility.</p> <table> <tr> <td><code>_STDIO_BADFD</code></td><td>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.</td></tr> <tr> <td><code>_STDIO_BADFD_SIGNAL</code></td><td>This variable takes an integer or string representing any valid signal. See signal.h(3HEAD) 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.</td></tr> </table>	<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 signal.h(3HEAD) 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.
<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 signal.h(3HEAD) 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.				
示例	<p>示例 1 Limit the number of file descriptors and FILE standard I/O structures.</p> <p>The following example limits the number of file descriptors and FILE standard I/O structures to 1000.</p> <pre>\$ ulimit -n 1000 \$ LD_PRELOAD_32=/usr/lib/extendedFILE.so.1 <i>application</i> [<i>arg...</i>]</pre> <p>示例 2 Enable the extended FILE facility.</p> <p>The following example enables the extended FILE facility. See enable_extended_FILE_stdio(3C) for more examples.</p>				

示例 2 Enable the extended FILE facility. (续)

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

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

文件 /usr/lib/extendedFILE.so.1 enabling library

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

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

另请参见 [limit\(1\)](#), [open\(2\)](#), [enable_extended_FILE_stdio\(3C\)](#), [fdopen\(3C\)](#), [fopen\(3C\)](#), [popen\(3C\)](#), [signal.h\(3HEAD\)](#), [stdio\(3C\)](#), [attributes\(5\)](#)

警告 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.

引用名	extensions – localedef extensions description file
描述	<p>A localedef extensions description file or <i>extensions</i> file defines various extensions for the localedef(1) command.</p> <p>The localedef extensions description file provides:</p> <ul style="list-style-type: none">■ 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.
另请参见	locale(1) , localedef(1) , environ(5) , locale(5)

引用名	filesystem – file system organization						
用法概要	/						
描述	<p>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 <code>etc</code>, <code>usr</code>, and <code>var</code>, 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.</p> <p>All subdirectories of the root file system that are part of the Oracle Solaris OS, with the exception of <code>/var</code>, must be contained in the same file system as the root file system. In addition, all OS components 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.</p> <p>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 <code>uname -i</code> 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 <code>/etc/hosts</code>. 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 <code>/platform</code> and <code>/usr/platform</code>.</p>						
Root File System	<p>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 (<code>/</code>) directory also contains mount point directories where local and remote file systems can be attached to the file system hierarchy.</p> <table> <tr> <td><code>/</code></td><td>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.</td></tr> <tr> <td><code>/boot</code></td><td>Directory that contains files and executables that are needed for booting the system. Directory that contains files and executables that are needed for booting the system.</td></tr> <tr> <td><code>/bin</code></td><td>Symbolic link to the <code>/usr/bin</code> directory that contains system executables and scripts.</td></tr> </table>	<code>/</code>	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.	<code>/boot</code>	Directory that contains files and executables that are needed for booting the system. Directory that contains files and executables that are needed for booting the system.	<code>/bin</code>	Symbolic link to the <code>/usr/bin</code> directory that contains system executables and scripts.
<code>/</code>	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.						
<code>/boot</code>	Directory that contains files and executables that are needed for booting the system. Directory that contains files and executables that are needed for booting the system.						
<code>/bin</code>	Symbolic link to the <code>/usr/bin</code> directory that contains system executables and scripts.						

<code>/dev</code>	Directory that contains special device files. Typically, device files are built to match the kernel and hardware configuration of the system.
<code>/devices</code>	Mount point directory for the <code>devfs</code> file system that manages the device name space.
<code>/etc</code>	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.
<code>/export/home</code> or <code>/home</code>	Directory or file system mount point for user home directories, which store user files. By default, the <code>/home</code> directory is an automounted file system.
<code>/kernel</code>	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, <code>/kernel/genunix</code> . See kernel(1M) . An approved installation location for bundled Oracle Solaris software and for add-on system software.
<code>/lib</code>	Directory that contains core system libraries. Historically, this directory contained essential library components for system startup.
<code>/media</code>	Directory for accessing removable media that is automatically mounted.
<code>/mnt</code>	Default temporary mount point directory for file systems. This empty directory is used to temporarily mount a file system.
<code>/net</code>	Temporary mount point directory for file systems that are mounted by the automounter.
<code>/opt</code>	Directory for unbundled application packages.
<code>/platform</code>	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 <code>/ (root)</code> .
<code>/proc</code>	Mount point directory for the process file system.
<code>/root</code>	Home directory for the root user.
<code>/rpool</code>	Mount point directory for the ZFS boot-related components. By default, the root pool is named <code>rpool</code> during installation.
<code>/sbin</code>	Symbolic link to the <code>/usr/sbin</code> directory.

<code>/system</code>	Mount point directory for the contract (CTFS) and object (OBJFS) file systems.
<code>/tmp</code>	Directory that contains temporary files that are removed during a boot operation.
<code>/usr</code>	Directory that contains platform-dependent and platform-independent binaries and files. The <code>/usr/share</code> subdirectory contains platform-independent files. The rest of the <code>/usr</code> directory contains platform-dependent files.
<code>/usr/bin</code>	Directory that contains platform-dependent, user-invoked executables. These are commands that users expect to be run as part of their normal <code>\$PATH</code> . An approved installation location for bundled Oracle Solaris software. The analogous location for unbundled system software or for applications is <code>/opt/package/bin</code> .
<code>/usr/lib</code>	Directory that augments the contents of <code>/lib</code> with additional system libraries, and other supporting files that are required by programs at runtime.
<code>/usr/java</code>	Directory that contains Java files and executables.
<code>/usr/sbin</code>	Contains essential executables used in the booting process and in manual system recovery. Historically, this directory was needed to recover the system before the <code>/usr</code> file system was mounted. In this Oracle Solaris release, <code>/usr</code> is a directory, not a separate file system and is available when the root file system is mounted.
<code>/var</code>	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.
<code>/var/tmp</code>	<p>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 <code>/var/tmp</code> to clear all of the files except editor temporary files by setting the <code>clean_vartmp</code> property value of the <code>rmtmpfiles</code> service. This is done with the following commands:</p> <pre># svccfg -s svc:/system/rmtmpfiles setprop\ options/clean_vartmp = "true" # svcadm refresh svc:/system/rmtmpfiles:default</pre> <p>The <code>solaris.smf.value.rmtmpfiles</code> authorization is required to modify this property.</p>

另请参见

[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\(7FS\)](#),

引用名

fmri – Fault Managed Resource Identifier

描述

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 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
nvlist	Unordered list of nvpairs
nvlist array	Array of nvlist
fmri	An nvlist 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 FMRIs 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

scheme	string	Committed
version	uint8	Committed

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.

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 identified as 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 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
cpu	0, 1	No	Logical processor identification
dev	0	No	Device resources
fmd	0	No	Fault Management Daemon modules
hc	0	Yes	Hardware Components
mem	0	No	Memory modules and cache
mod	0	No	Kernel modules
pkg	0, 1	Yes	Software packages
svc	0	No	SMF services & service instances
sw	0	No	Software objects
zfs	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
scheme	string	Committed, value "cpu"
version	uint8	Committed, value 0
cpuid	uint32	Committed
[serial]	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

scheme	string	Committed, value "cpu"
version	uint8	Committed, value 1
cpuid	uint32	Committed
[serial]	string	Private
[cpumask]	uint8	Private
[cpufru]	string	Private
[cacheindex]	uint32	Private
[cacheway]	uint32	Private
[cachebit]	uint16	Private
[cachetype]	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 cpuf ru 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

scheme	string	Committed, value "dev"
version	uint8	Committed, value 0
device-path	string	Committed
[devid]	string	Private
[target-port-l0id]	string	Private

device-path

A filesystem path within the Oracle Solaris /devices tree that identifies a device node, such as pci@0,0/pci108e,cb84@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.

Member Name	Data Type	Stability

scheme	string	Committed, value "fmd"
version	uint8	Committed, value 0
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 hc VERSION 0

The hc scheme is used to identify hardware components. In most cases these FMRI's 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>
```

Current variant of hc version 0:

The authority nvlist is 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
[devid]	string	Private, 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.

[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 string (see [hostid\(1\)](#) in Oracle Solaris); not generally used. Not necessarily invariant.
- [devid]
The device id.

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
[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.

[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>]

                [: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 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	arrayPrivate
[physaddr]	uint64	Private
[offset]	uint64	Private

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-id]	int32	Private

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.

mod-id
The module id in the current Oracle Solaris instance, as per [modinfo\(1M\)](#).

String form:

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

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

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.

另请参见

[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](#)

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

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.

另请参见

[find\(1\)](#), [ksh\(1\)](#), [fnmatch\(3C\)](#), [regex\(5\)](#)

引用名

formats – file format notation

描述

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:

<i>characters</i>	Characters that are not <i>escape sequences</i> or <i>conversion specifications</i> , as described below, are copied to the output.
<i>escape sequences</i>	Represent non-graphic characters.
<i>conversion specifications</i>	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
<code>\n</code>	newline	Moves the printing position to the start of the next line.
<code>\r</code>	carriage-return	Moves the printing position to the start of the current line.
<code>\t</code>	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.
<code>\v</code>	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 <i>flags</i> 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

	<i>precision</i> is omitted from the argument, six digits are written after the radix character; if the <i>precision</i> is explicitly 0, no radix character appears.
<i>e,E</i>	The floating point number argument is written in the style <code>[-]d.ddde±dd</code> (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 <code>LC_NUMERIC</code> 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 <code>E</code> conversion character produces a number with <code>E</code> instead of <code>e</code> 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.
<i>g,G</i>	The floating point number argument is written in style <code>f</code> or <code>e</code> (or in style <code>E</code> in the case of a <code>G</code> conversion character), with the precision specifying the number of significant digits. The style used depends on the value converted: style <code>g</code> 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.
<i>c</i>	The integer argument is converted to an unsigned char and the resulting byte is written.
<i>s</i>	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 <i>precision</i> 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.
<i>%</i>	Write a <code>%</code> 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.

示例

示例 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>
```

示例 2 To show pi written to 5 decimal places:

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

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

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

另请参见

[awk\(1\)](#), [printf\(1\)](#), [printf\(3C\)](#), [scanf\(3C\)](#)

引用名

fsattr – extended file attributes

描述

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. Implementation are therefore permitted to reject operations that are not supported. For example, the implementation for the UFS file system allows only regular files as attributes (for example, no sub-directories) and rejects attempts to place attributes on attributes.

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

link	Any attempt to create links between attribute and non-attribute space is rejected to prevent security-related or otherwise sensitive attributes from being exposed, and therefore manipulable, as regular files.
rename	Any attempt to rename between attribute and non-attribute space is rejected to prevent an already linked file from being renamed and thereby circumventing the link restriction above.
mkdir	
symlink	
mknod	Any attempt to create a “non-regular” file in attribute space is rejected to reduce the functionality, and therefore exposure and risk, of the initial implementation.

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

Shell-level API

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

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

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

- `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)`, `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:

示例 1 List extended attributes on a file.

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

示例 2 Open an extended attribute.

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

or

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

示例 3 Read from an extended attribute.

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

示例 4 Create an extended attribute.

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

or

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

示例 5 Write to an extended attribute.

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

示例 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.

另请参见

[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\(2\)](#), [attropen\(3C\)](#), [standards\(5\)](#)

引用名	grub – Solaris 上的 GRandom Unified Bootloader 软件
描述	<p>Solaris 操作系统的当前发行版附带了 GRUB (GRandom Unified Bootloader) 软件。GRUB 由自由软件基金会开发并提供支持。</p> <p>GRUB 手册概述（可从 www.gnu.org 访问）对 GRUB 进行了介绍：</p> <p>简单地说，引导装载程序就是在计算机启动时最先运行的软件程序。引导装载程序负责装入控件并将其传输到操作系统内核软件（如 Linux 或 GNU Mach），而内核又会初始化操作系统（例如，GNU [编辑注释：或 Solaris] 系统）的其余部分。</p> <p>GNU GRUB 是一种十分强大的引导装载程序，可通过链式装入的方式装入各种免费以及专用的操作系统。GRUB 设计用于解决引导个人计算机的复杂性问题；该程序和此手册均与 GNU 计算机平台紧密相关，但将来可能会解决移植到其他平台的问题。[编辑注释：Oracle 已将 GRUB 移植到 Solaris 操作系统。]</p> <p>灵活性是 GRUB 的主要特点之一；GRUB 可识别文件系统和内核可执行文件格式，这使您可以随意装入任意操作系统，无需记录内核在磁盘上的物理位置。因此，您只需指定内核的文件名以及内核所在的驱动器和分区即可装入内核。</p> <p>在 Solaris 计算机中，x86 平台上支持 GRUB。Solaris 附带的 GRUB 软件新增了两个实用程序，它们未在开源分发软件中提供：</p> <p>bootadm(1M) 用于管理引导归档文件以及更改 GRUB 菜单。</p> <p>installgrub(1M) 从磁盘装入引导程序。</p> <p>在 Solaris 手册页中对这两个实用程序进行了介绍。</p> <p>除了 Solaris 特定的这两个实用程序之外，在 GRUB 手册中还对 GRUB 软件进行了介绍，可从 Oracle Web 站点获取该手册的 PDF 版本。grub(8) 开源手册页也可从该站点获取。此手册页对 GRUB shell 进行了介绍。</p>
另请参见	<p>boot(1M)、bootadm(1M)、installgrub(1M)</p> <p>《安装 Oracle Solaris 11 系统》</p> <p>《Oracle Solaris 管理：常见任务》</p> <p>http://www.gnu.org/software/grub</p>

引用名 gss_auth_rules – overview of GSS authorization

描述 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\(3GSS\)](#).

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

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

另请参见 [ftp\(1\)](#), [ssh\(1\)](#), [gsscred\(1M\)](#), [gss_compare_name\(3GSS\)](#), [passwd\(4\)](#), [attributes\(5\)](#), [krb5_auth_rules\(5\)](#)

引用名 hal – 硬件抽象层概述

描述 硬件抽象层 (Hardware Abstraction Layer, HAL) 提供了连接到系统的各种硬件的视图。当硬件配置通过热插拔或其他机制发生更改时，此视图也将随之进行动态更新。HAL 将硬件表示为设备对象。设备对象由唯一标识符进行标识并带有一组键/值对，这些键/值对称为设备属性。有些属性源自实际硬件，有些属性由设备信息文件（.fdi 文件）合并而来，有些属性与实际设备配置相关。

HAL 通过 D-Bus 提供简单易用的 API。D-Bus 是一种 IPC 框架，在其功能中提供了使应用程序可以相互通话的系统范围的消息总线。具体来说，D-Bus 提供了异步通知，使 HAL 可以在添加和删除设备以及设备的属性发生更改时通知消息总线上的其他对等方。

在 Solaris 操作系统上，HAL 由守护进程、[hald\(1M\)](#) 和一组实用程序（用于添加和删除设备以及修改设备的属性）提供支持。

另请参见 [hald\(1M\)](#)、[fdi\(4\)](#)

请参见以下网址上的 HAL 页，包括 HAL 规范：<http://freedesktop.org>。

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

描述 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		

- 文件

/usr/lib/iconv/*.so

conversion modules

/usr/lib/iconv/*.t

conversion tables

/usr/lib/iconv/iconv_data

list of conversions supported by conversion tables
- 另请参见

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

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

描述 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	KOI8-R	MS 1251	KOI8-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		

- 文件
- /usr/lib/iconv/*.so

conversion modules
- /usr/lib/iconv/*.t

conversion tables
- /usr/lib/iconv/iconv_data

list of conversions supported by conversion tables
- 另请参见
- [iconv\(1\)](#), [iconv\(3C\)](#), [iconv\(5\)](#)

引用名
描述

iconv – code set conversion tables
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	

文件

/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.

另请参见

[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\(5\)](#)

引用名 iconv_646 – code set conversion tables for ISO 646

描述 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

- 文件
- /usr/lib/iconv/*.so

conversion modules
- /usr/lib/iconv/*.t

conversion tables
- /usr/lib/iconv/iconv_data

list of conversions supported by conversion tables
- 另请参见
- iconv(1), iconv(3C), iconv(5)

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

描述 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		

文件

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

另请参见

[iconv\(1\)](#), [iconv\(3C\)](#), [iconv\(5\)](#)

引用名 iconv_8859-1 – code set conversion tables for ISO 8859-1 (Latin 1)

描述 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

文件	/usr/lib/iconv/*.so	conversion modules
	/usr/lib/iconv/*.t	conversion tables
	/usr/lib/iconv/iconv_data	list of conversions supported by conversion tables
另请参见	iconv(1) , iconv(3C) , iconv(5)	

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

描述 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		

- 文件

/usr/lib/iconv/*.so

conversion modules

/usr/lib/iconv/*.t

conversion tables

/usr/lib/iconv/iconv_data

list of conversions supported by conversion tables
- 另请参见

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

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

描述 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		

- 文件

/usr/lib/iconv/*.so

conversion modules

/usr/lib/iconv/*.t

conversion tables

/usr/lib/iconv/iconv_data

list of conversions supported by conversion tables
- 另请参见

[iconv\(1\)](#), [iconv\(3C\)](#), [iconv\(5\)](#)
- 170

手册页第 5 部分：标准、环境和宏 • 最后修订 18 Apr 1997

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

描述 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		

- 文件

/usr/lib/iconv/*.so

conversion modules

/usr/lib/iconv/*.t

conversion tables

/usr/lib/iconv/iconv_data

list of conversions supported by conversion tables
- 另请参见

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

引用名 iconv_koi8-r – code set conversion tables for KOI8-R

描述 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			
KOI8-R	MS 1251	KOI8-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			
KOI8-R	Mac Cyrillic	KOI8-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			
KOI8-R	Mac Cyrillic	KOI8-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		

- 文件

/usr/lib/iconv/*.so

conversion modules

/usr/lib/iconv/*.t

conversion tables

/usr/lib/iconv/iconv_data

list of conversions supported by conversion tables
- 另请参见

[iconv\(1\)](#), [iconv\(3C\)](#), [iconv\(5\)](#)

引用名 iconv_mac_cyr – code set conversion tables for Macintosh Cyrillic

描述 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	KOI8-R	Mac Cyrillic	KOI8-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

- 文件

/usr/lib/iconv/*.so

conversion modules

/usr/lib/iconv/*.t

conversion tables

/usr/lib/iconv/iconv_data

list of conversions supported by conversion tables
- 另请参见

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

引用名 iconv_maz – code set conversion tables for Mazovia

描述 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		

文件

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

另请参见

[iconv\(1\)](#), [iconv\(3C\)](#), [iconv\(5\)](#)

引用名 iconv_pc_cyr – code set conversion tables for Alternative PC Cyrillic

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

PC Cyrillic to MS 1251

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

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		

- 文件

/usr/lib/iconv/*.so

conversion modules

/usr/lib/iconv/*.t

conversion tables

/usr/lib/iconv/iconv_data

list of conversions supported by conversion tables
- 另请参见

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

引用名 iconv_unicode – code set conversion tables for Unicode

描述 The following code set conversions are supported:

CODE SET CONVERSIONS SUPPORTED			

FROM Code Set Code	FROM Filename Element	TO Code Set Target Code	TO Filename Element
ISO 8859-1 (Latin 1)	8859-1	UTF-8	UTF-8
ISO 8859-2 (Latin 2)	8859-2	UTF-8	UTF-8
ISO 8859-3 (Latin 3)	8859-3	UTF-8	UTF-8
ISO 8859-4 (Latin 4)	8859-4	UTF-8	UTF-8
ISO 8859-5 (Cyrillic)	8859-5	UTF-8	UTF-8
ISO 8859-6 (Arabic)	8859-6	UTF-8	UTF-8
ISO 8859-7 (Greek)	8859-7	UTF-8	UTF-8
ISO 8859-8 (Hebrew)	8859-8	UTF-8	UTF-8
ISO 8859-9 (Latin 5)	8859-9	UTF-8	UTF-8
ISO 8859-10 (Latin 6)	8859-10	UTF-8	UTF-8
Japanese EUC	eucJP	UTF-8	UTF-8
Chinese/PRC EUC (GB 2312-1980)	gb2312	UTF-8	UTF-8
ISO-2022	iso2022	UTF-8	UTF-8
Korean EUC	ko_KR-euc	Korean UTF-8	ko_KR-UTF-8
ISO-2022-KR	ko_KR-iso2022-7	Korean UTF-8	ko_KR-UTF-8
Korean Johap (KS C 5601-1987)	ko_KR-johap	Korean UTF-8	ko_KR-UTF-8
Korean Johap (KS C 5601-1992)	ko_KR-johap92	Korean UTF-8	ko_KR-UTF-8
Korean UTF-8	ko_KR-UTF-8	Korean EUC	ko_KR-euc
Korean UTF-8	ko_KR-UTF-8	Korean Johap (KS C 5601-1987)	ko_KR-johap
Korean UTF-8	ko_KR-UTF-8	Korean Johap (KS C 5601-1992)	ko_KR-johap92
KOI8-R (Cyrillic)	KOI8-R	UCS-2	UCS-2
KOI8-R (Cyrillic)	KOI8-R	UTF-8	UTF-8
PC Kanji (SJIS)	PCK	UTF-8	UTF-8
PC Kanji (SJIS)	SJIS	UTF-8	UTF-8
UCS-2	UCS-2	KOI8-R (Cyrillic)	KOI8-R
UCS-2	UCS-2	UCS-4	UCS-4

CODE SET CONVERSIONS SUPPORTED			

FROM Code Set Code	FROM Filename Element	TO Code Set Target Code	TO Filename

Element		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 (GB 2312-1980)	gb2312
UTF-8	UTF-8	ISO-2022	iso2022
UTF-8	UTF-8	KOI8-R (Cyrillic)	KOI8-R
UTF-8	UTF-8	PC Kanji (SJIS)	PCK
UTF-8	UTF-8	PC Kanji (SJIS)	SJIS
UTF-8	UTF-8	UCS-2	UCS-2
UTF-8	UTF-8	UCS-4	UCS-4
UTF-8	UTF-8	UTF-16	UTF-16
UTF-8	UTF-8	UTF-7	UTF-7
UTF-8	UTF-8	Chinese/PRC EUC (GB 2312-1980)	zh_CN.euc

CODE SET CONVERSIONS SUPPORTED

FROM Code Set Code	TO Code Set		
	FROM Filename Element	Target Code	TO Filename Element
UTF-8	UTF-8	ISO 2022-CN	zh_CN.iso2022-7
UTF-8	UTF-8	Chinese/Taiwan Big5	zh_TW-big5
UTF-8	UTF-8	Chinese/Taiwan EUC (CNS 11643-1992)	zh_TW-euc

UTF-8	UTF-8	ISO 2022-TW	zh_TW-iso2022-7
Chinese/PRC EUC (GB 2312-1980)	zh_CN.euc	UTF-8	UTF-8
ISO 2022-CN	zh_CN.iso2022-7	UTF-8	UTF-8
Chinese/Taiwan Big5	zh_TW-big5	UTF-8	UTF-8
Chinese/Taiwan EUC (CNS 11643-1992)	zh_TW-euc	UTF-8	UTF-8
ISO 2022-TW	zh_TW-iso2022-7	UTF-8	UTF-8

示例

示例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`

文件

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

另请参见

[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, Rationel 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.

附注

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.		

引用名	ieee802.11 – 802.11 kernel statistics																										
描述	<p>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.</p> <table><tr><td>tx_frags</td><td>Count of data and management fragments transmitted.</td></tr><tr><td>rx_frags</td><td>Count of data and management fragments received.</td></tr><tr><td>rx_dups</td><td>Count of duplicate frames received. Duplicates are determined by the sequence control field.</td></tr><tr><td>mcast_tx</td><td>Count of broadcast and multicast frames transmitted.</td></tr><tr><td>mcast_rx</td><td>Count of broadcast and multicast frames received.</td></tr><tr><td>tx_failed</td><td>Count of frames that could not be transmitted due to the retransmission limit being reached.</td></tr><tr><td>tx_retrans</td><td>Count of frames successfully retransmitted after one or more retransmissions.</td></tr><tr><td>tx_reretrans</td><td>Count of frames successfully retransmitted after more than one retransmission.</td></tr><tr><td>rts_success</td><td>Count of times a CTS was received in response to an RTS.</td></tr><tr><td>rts_failure</td><td>Count of times a CTS was not received in response to an RTS.</td></tr><tr><td>ack_failure</td><td>Count of times an ACK was expected but was not received.</td></tr><tr><td>fcs_errors</td><td>Count of frames received with FCS errors.</td></tr><tr><td>wep_errors</td><td>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.</td></tr></table>	tx_frags	Count of data and management fragments transmitted.	rx_frags	Count of data and management fragments received.	rx_dups	Count of duplicate frames received. Duplicates are determined by the sequence control field.	mcast_tx	Count of broadcast and multicast frames transmitted.	mcast_rx	Count of broadcast and multicast frames received.	tx_failed	Count of frames that could not be transmitted due to the retransmission limit being reached.	tx_retrans	Count of frames successfully retransmitted after one or more retransmissions.	tx_reretrans	Count of frames successfully retransmitted after more than one retransmission.	rts_success	Count of times a CTS was received in response to an RTS.	rts_failure	Count of times a CTS was not received in response to an RTS.	ack_failure	Count of times an ACK was expected but was not received.	fcs_errors	Count of frames received with FCS errors.	wep_errors	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.
tx_frags	Count of data and management fragments transmitted.																										
rx_frags	Count of data and management fragments received.																										
rx_dups	Count of duplicate frames received. Duplicates are determined by the sequence control field.																										
mcast_tx	Count of broadcast and multicast frames transmitted.																										
mcast_rx	Count of broadcast and multicast frames received.																										
tx_failed	Count of frames that could not be transmitted due to the retransmission limit being reached.																										
tx_retrans	Count of frames successfully retransmitted after one or more retransmissions.																										
tx_reretrans	Count of frames successfully retransmitted after more than one retransmission.																										
rts_success	Count of times a CTS was received in response to an RTS.																										
rts_failure	Count of times a CTS was not received in response to an RTS.																										
ack_failure	Count of times an ACK was expected but was not received.																										
fcs_errors	Count of frames received with FCS errors.																										
wep_errors	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.																										
另请参见	kstat(1M)																										

引用名 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

描述 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.

cap_autoneg	Capable of auto-negotiation
cap_1000fdx	Capable of 1000 full duplex operation
cap_1000hdx	Capable of 1000 half duplex operation
cap_100fdx	Capable of 100 full duplex operation
cap_100hdx	Capable of 100 half duplex operation
cap_10fdx	Capable of 10 full duplex operation
cap_10hdx	Capable of 10 half duplex operation
cap_rem_fault	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.

adv_cap_autoneg	Advertise auto-negotiation capability
adv_cap_1000fdx	Advertise 1000 full duplex capability
adv_cap_1000hdx	Advertise 1000 half duplex capability
adv_cap_100fdx	Advertise 100 full duplex capability
adv_cap_100hdx	Advertise 100 half duplex capability
adv_cap_10fdx	Advertise 10 full duplex capability
adv_cap_10hdx	Advertise 10 half duplex capability
adv_rem_fault	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.

lp_cap_autoneg	Link partner advertises auto-negotiation capability
lp_cap_1000fdx	Link partner advertises 1000 full duplex capability
lp_cap_1000hdx	Link partner advertises 1000 half duplex capability
lp_cap_100fdx	Link partner advertises 100 full duplex capability
lp_cap_100hdx	Link partner advertises 100 half duplex capability
lp_cap_10fdx	Link partner advertises 10 full duplex capability
lp_cap_10hdx	Link partner advertises 10 half duplex capability
lp_rem_fault	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.

xcvr_addr	MII address in the 0 to 31 range of the physical layer device in use for a given Ethernet device
xcvr_id	MII transceiver manufacturer and device ID
xcvr_inuse	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.

link_up	1	Link is up
	0	Link is down
link_duplex	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 bot 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.

另请参见

[dladm\(1M\)](#), [driver.conf\(4\)](#), [bge\(7D\)](#), [d1pi\(7P\)](#), [eri\(7D\)](#), [gld\(7D\)](#), [hme\(7D\)](#), [qfe\(7d\)](#)

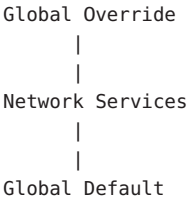
附注

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.

引用名	ipfilter – IP 包过滤软件
描述	<p>IP 过滤器是 Solaris 系统上提供包过滤功能的软件。在设置正确的系统上，IP 过滤器可用于构建防火墙。</p> <p>Solaris IP 过滤器随 Solaris 操作系统一起安装。但是，缺省情况下不启用包过滤。有关启用和激活 IP 过滤器功能的过程，请参见 ipf(1M)。</p>
服务	<p>ipfilterSMF 服务支持 start、stop、restart 和 refresh 方法。这些方法可以通过使用 svcadm(1M) 进行调用。</p> <p>start 装入 ipfilter 内核模块并根据配置激活任何防火墙或 NAT 规则。</p> <p>stop 清除应用的所有防火墙和 NAT 规则以及创建的任何活动会话信息。仅当能够进入主机的任何网络通信流量不存在风险时，才能停止启用联网功能的服务。</p> <p>restart 停止然后启动 ipfilter 服务。对活动防火墙使用该方法会导致一定时段的暴露，在该时段内，通信流量可能在未经过滤的情况下进入和/或通过防火墙。</p> <p>refresh 装入当前配置并从旧配置切换为新配置，而在这期间系统一直都在积极使用安全策略，没有间断。</p>
基于主机的防火墙	<p>为简化 IP 过滤器配置管理，创建了一个防火墙框架，以便用户可以通过在系统和 service 级别指定防火墙策略来配置 IP 过滤器。在指定用户定义的防火墙策略后，该框架会生成一组 IP 过滤器规则来执行所需的系统行为。用户可以指定系统和 service 防火墙策略，以允许或拒绝来自特定主机、子网和接口的网络通信流量。这些策略将转换成一组活动的 IPF 规则来执行指定的防火墙策略。</p> <p>如果用户选择不使用该框架，仍可以指定他们自己的 ipf 规则文件。有关如何启用定制规则以及 ipf(4) 找出 ipf 规则语法的信息，请参见 ipf(1M)。</p>
型号	<p>本节介绍基于主机的防火墙框架。有关如何配置防火墙策略的详细信息，请参见 svc.ipfd(1M)。</p> <p>有一种具有不同优先级的三层方法可帮助用户实现所需行为。</p> <p>全局缺省 全局缺省—缺省的系统范围防火墙策略。所有服务都将自动继承此策略，除非有服务修改了其防火墙策略。</p> <p>网络服务 其优先级高于全局缺省。无论全局缺省策略如何，服务的策略都将允许/禁止其特定端口的通信。</p> <p>全局覆盖 另一种系统范围的策略，优先于网络服务层中特定服务的需求。</p>



防火墙策略包括一种防火墙模式和一组可选的网络源。网络源包括 IP 地址、子网和本地网络接口，所有这些网络源都可供系统从中接收传入通信。基本防火墙模式组包括：

- 无
 无防火墙，允许所有传入通信。
- 拒绝
 允许所有传入通信，只拒绝来自指定源的通信。
- 允许
 拒绝所有传入通信，只允许来自指定源的通信。

有关层的详细说明

第一个系统范围层全局缺省定义应用于**任何**传入通信的防火墙策略；例如，允许或阻止来自某一 IP 地址的所有通信。这使得很容易设置阻止所有传入通信或阻止来自不需要源的所有传入通信的策略。

网络服务层包含向远程客户机提供服务的本地程序（例如，telnetd、sshd 和 httpd）的防火墙策略。其中每个程序（即一个网络服务）都有其自己的防火墙策略来控制对其服务的访问。最初，服务的策略设置为继承全局缺省策略，即“使用全局缺省”模式。这使得很容易在全局缺省层设置一个可由所有服务继承的策略。

如果某一服务的策略不同于全局缺省策略，则该服务的策略具有较高的优先级。如果全局缺省策略设置为阻止来自子网的所有通信，则 SSH 服务可配置为允许从该子网中的特定主机进行访问。所有网络服务的所有策略的集合构成了网络服务层。

第二个系统范围层全局覆盖的防火墙策略也应用于任何传入网络通信。此策略具有最高优先级，将覆盖其他层中的策略，具体来说覆盖网络服务的需求。例如，无论服务的策略如何，都能理想地阻止已知的恶意源。

用户交互

此框架利用 IP 过滤器功能，仅在启用 `svc:/network/ipfilter` 时才有效，而在禁用 `network/ipfilter` 时将无效。同样，网络服务的防火墙策略也仅在启用该服务时才有效，而在禁用该服务时也将无效。具有有效防火墙的系统具有 IP 过滤器规则（用于正在运行/已启用的每个网络服务）以及系统范围的策略（防火墙模式不是无）。

用户可通过设置系统范围策略和各网络服务的策略来配置防火墙。有关配置防火墙策略的方法，请参见 `svc.ipfd(1M)`。

防火墙框架由策略配置和一个机制组成，用以从策略生成 IP 过滤器规则并应用这些规则来获取所需的 IP 过滤器配置。下面对设计和用户交互进行了快速汇总：

- 系统范围策略存储在 `network/ipfilter` 中

- 网络服务的策略存储在各 SMF 服务中
- 用户通过启用 `network/ipfilter`（请参见 [ipf\(1M\)](#)）来激活防火墙
- 用户通过启用/禁用网络服务来激活/取消激活该服务的防火墙
- 对系统范围或每服务防火墙策略进行更改会导致系统防火墙规则的更新

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
接口稳定性	Committed（已确定）

另请参见 [svcs\(1\)](#)、[ipf\(1M\)](#)、[ipnat\(1M\)](#)、[svcadm\(1M\)](#)、[svc.ipfd\(1M\)](#)、[ipf\(4\)](#)、[ipnat\(4\)](#)、[attributes\(5\)](#)
《Oracle Solaris 管理：IP 服务》

附注 `ipfilter` 服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：
`svc:/network/ipfilter:default`

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

IP 过滤器启动配置文件存储在 `/etc/ipf` 中。

引用名	isalist – Solaris 软件可识别的本机指令集																				
描述	<p>本文列出了由 <code>isalist(1)</code> 以及 <code>SI_ISALIST</code> 命令（属于 <code>sysinfo(2)</code>）返回的可能的指令集名称。</p> <p>该列表按某种方式对指令集系列进行了排序，即排在后面的名称通常比前面的名称的运行速度快；请注意，这与 <code>isalist(1)</code> 和 <code>sysinfo(2)</code> 列出的顺序相反。在下面的值列表中，带编号的条目通常表示性能渐增；带字母的条目表示互斥或无法排序。</p> <p>此功能已过时，可能会在以后的 Solaris 版本中删除。以下列表不会影响现代处理器已创建的所有扩展。要了解处理指令集扩展的更好方法，请参见 <code>getisax(2)</code>。</p>																				
SPARC 平台	<p>在适当时，指出了 Oracle 的 C 4.0 编译器带有给定值的 <code>-xarch</code> 选项的对应项。其他编译器可能也有类似的选项。</p> <table><tr><td>1a. sparc</td><td>指示 SPARC V8 指令集，其定义详见《SPARC Architecture Manual》，第 8 版，Prentice-Hall, Inc. 出版，1992 年。在特定系统上，某些指令（如整数相乘和相除、FSMULD 及对四元操作数执行的所有浮点运算）可由内核模拟。</td></tr><tr><td>1b. sparcv7</td><td>与 sparc 相同。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v7</code> 选项生成的代码相对应。</td></tr><tr><td>2. sparcv8-fsmuld</td><td>与 sparc 类似，但必须在硬件中执行整数相乘和相除。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v8a</code> 选项生成的代码相对应。</td></tr><tr><td>3. sparcv8</td><td>与 sparcv8-fsmuld 类似，但也必须在硬件中执行 FSMULD。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v8</code> 选项生成的代码相对应。</td></tr><tr><td>4. sparcv8plus</td><td>指示 SPARC V8 指令集以及 SPARC V9 指令集（其定义详见《SPARC Architecture Manual》，第 9 版，Prentice-Hall 出版，1994 年）中可按照 V8+ 技术规范使用的那些指令。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v8plus</code> 选项生成的代码相对应。</td></tr><tr><td>5a. sparcv8plus+vis</td><td>与 sparcv8plus 类似，但新增了可按照 V8+ 技术规范使用的那些 UltraSPARC I 可视化指令。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v8plusa</code> 选项生成的代码相对应。</td></tr><tr><td>5b. sparcv8plus+fmuladd</td><td>与 sparcv8plus 类似，但新增了 Fujitsu SPARC64 浮点相乘-相加和相乘-相减指令。</td></tr><tr><td>6. sparcv9</td><td>指示 SPARC V9 指令集，其定义详见《SPARC Architecture Manual》，第 9 版，Prentice-Hall 出版，1994 年。</td></tr><tr><td>7a. sparcv9+vis</td><td>与 sparcv9 类似，但新增了 UltraSPARC I 可视化指令。</td></tr><tr><td>7b. sparcv9+vis2</td><td>与 sparcv9 类似，但新增了 UltraSPARC III 可视化指令。</td></tr></table>	1a. sparc	指示 SPARC V8 指令集，其定义详见《SPARC Architecture Manual》，第 8 版，Prentice-Hall, Inc. 出版，1992 年。在特定系统上，某些指令（如整数相乘和相除、FSMULD 及对四元操作数执行的所有浮点运算）可由内核模拟。	1b. sparcv7	与 sparc 相同。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v7</code> 选项生成的代码相对应。	2. sparcv8-fsmuld	与 sparc 类似，但必须在硬件中执行整数相乘和相除。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v8a</code> 选项生成的代码相对应。	3. sparcv8	与 sparcv8-fsmuld 类似，但也必须在硬件中执行 FSMULD。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v8</code> 选项生成的代码相对应。	4. sparcv8plus	指示 SPARC V8 指令集以及 SPARC V9 指令集（其定义详见《SPARC Architecture Manual》，第 9 版，Prentice-Hall 出版，1994 年）中可按照 V8+ 技术规范使用的那些指令。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v8plus</code> 选项生成的代码相对应。	5a. sparcv8plus+vis	与 sparcv8plus 类似，但新增了可按照 V8+ 技术规范使用的那些 UltraSPARC I 可视化指令。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v8plusa</code> 选项生成的代码相对应。	5b. sparcv8plus+fmuladd	与 sparcv8plus 类似，但新增了 Fujitsu SPARC64 浮点相乘-相加和相乘-相减指令。	6. sparcv9	指示 SPARC V9 指令集，其定义详见《SPARC Architecture Manual》，第 9 版，Prentice-Hall 出版，1994 年。	7a. sparcv9+vis	与 sparcv9 类似，但新增了 UltraSPARC I 可视化指令。	7b. sparcv9+vis2	与 sparcv9 类似，但新增了 UltraSPARC III 可视化指令。
1a. sparc	指示 SPARC V8 指令集，其定义详见《SPARC Architecture Manual》，第 8 版，Prentice-Hall, Inc. 出版，1992 年。在特定系统上，某些指令（如整数相乘和相除、FSMULD 及对四元操作数执行的所有浮点运算）可由内核模拟。																				
1b. sparcv7	与 sparc 相同。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v7</code> 选项生成的代码相对应。																				
2. sparcv8-fsmuld	与 sparc 类似，但必须在硬件中执行整数相乘和相除。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v8a</code> 选项生成的代码相对应。																				
3. sparcv8	与 sparcv8-fsmuld 类似，但也必须在硬件中执行 FSMULD。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v8</code> 选项生成的代码相对应。																				
4. sparcv8plus	指示 SPARC V8 指令集以及 SPARC V9 指令集（其定义详见《SPARC Architecture Manual》，第 9 版，Prentice-Hall 出版，1994 年）中可按照 V8+ 技术规范使用的那些指令。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v8plus</code> 选项生成的代码相对应。																				
5a. sparcv8plus+vis	与 sparcv8plus 类似，但新增了可按照 V8+ 技术规范使用的那些 UltraSPARC I 可视化指令。此选项与使用 Oracle C 4.0 编译器的 <code>-xarch=v8plusa</code> 选项生成的代码相对应。																				
5b. sparcv8plus+fmuladd	与 sparcv8plus 类似，但新增了 Fujitsu SPARC64 浮点相乘-相加和相乘-相减指令。																				
6. sparcv9	指示 SPARC V9 指令集，其定义详见《SPARC Architecture Manual》，第 9 版，Prentice-Hall 出版，1994 年。																				
7a. sparcv9+vis	与 sparcv9 类似，但新增了 UltraSPARC I 可视化指令。																				
7b. sparcv9+vis2	与 sparcv9 类似，但新增了 UltraSPARC III 可视化指令。																				

	7c. sparcv9+fmuladd	与 sparcv9 类似，但新增了 Fujitsu SPARC64 浮点相乘-相加和相乘-相减指令。
x86 平台	1. i386	Intel 80386 指令集，其说明详见《i386 Microprocessor Programmer's Reference Manual》。
	2. i486	Intel 80486 指令集，其说明详见《i486 Microprocessor Programmer's Reference Manual》。（实际上，这包括 i386 以及 CMPXCHG、BSWAP 和 XADD 指令。）
	3. Pentium	Intel Pentium 指令集，其说明详见《iPentium Processor User's Manual》。（实际上，这包括 i486、CPU_ID 指令以及 CPU_ID 指令指明存在的所有功能。）
	4. pentium+mmx	与 pentium 类似，包含保证存在的 MMX 指令。
	5. pentium_pro	Intel PentiumPro 指令集，其说明详见《PentiumPro Family Developer's Manual》。（实际上，这包括 pentium 以及保证存在的 CMOVcc、FCMOVcc、FCOMI 和 RDPMC 指令。）
	6. pentium_pro+mmx	与 pentium_pro 类似，包含保证存在的 MMX 指令。
	7. amd64	AMD Opteron 指令集，其说明详见AMD64 Architecture Programmer's Manual。

另请参见 [isalist\(1\)](#)、[getisax\(2\)](#)、[sysinfo\(2\)](#)

引用名

kerberos – overview of Solaris Kerberos implementation

描述

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:

<code>primary/instance@REALM</code>	
primary	A user, a host, or a service.
instance	A qualification of the primary. If the primary is a host — indicated by the keyword <code>host</code> — then the instance is the fully-qualified domain name of that host. If the primary is a user or service, then the instance is optional. Some instances, such as <code>admin</code> or <code>root</code> , are privileged.
realm	The Kerberos equivalent of a domain; in fact, in most cases the realm is directly mapped to a DNS domain name. 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\(3NSL\)](#)).

示例

示例 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.

另请参见

[kdestroy\(1\)](#), [kinit\(1\)](#), [klist\(1\)](#), [kpasswd\(1\)](#), [krb5.conf\(4\)](#), [krb5envvar\(5\)](#)

《Oracle Solaris 管理：安全服务》

附注

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 管理：安全服务》.

引用名

krb5_auth_rules – overview of Kerberos V5 authorization

描述

When kerberized versions of the ftp, rcp, rlogin, rsh, 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.

- 文件
- ~/k5login

Per user-account authorization file.
- /etc/passwd

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

属性

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

另请参见

[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\(5\)](#)

引用名

krb5envvar – Kerberos environment variables

描述

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
```

属性

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/security/kerberos-5
Interface Stability	Uncommitted

另请参见

[chmod\(1\)](#), [kinit\(1\)](#), [klist\(1\)](#), [ksh\(1\)](#), [kadmin\(1M\)](#), [kadmind\(1M\)](#), [krb5.conf\(4\)](#), [attributes\(5\)](#), [kerberos\(5\)](#)

引用名	kssl, KSSL – kernel SSL proxy
描述	<p>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.</p> <p>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.</p> <p>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.</p> <p>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).</p> <p>KSSL provides SSL processing for records passed with TCP over both IPv4 and IPv6.</p> <p>KSSL supports the following protocols: SSLv3, TLSv1.0</p>
另请参见	<p>ksslcfg(1M), smf(5)</p> <p>T. Dierks, C. Allen, RFC 2246, The TLS Protocol Version 1.0, The Internet Society, 1999.</p>

引用名	labels – Solaris Trusted Extensions label attributes
描述	<p>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).</p>
Mandatory Policy	<p>Various mandatory policies might be delivered in the lifetime of Solaris Trusted Extensions.</p> <p>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 Mesegeur model of Non-Inteference.</p> <p>For this MAC policy, two labels are always defined: <code>admin_low</code> and <code>admin_high</code>. The site's security administrator defines all other labels in <code>label_encodings(4)</code>. <code>admin_low</code> is associated with all normal user readable (viewable) Trusted Extensions objects. <code>admin_high</code> is associated with all other Trusted Extensions objects. Only administrative users have MAC read (view) access to <code>admin_high</code> objects and only administrative users have MAC write (modify) access to <code>admin_low</code> objects or <code>admin_high</code> objects.</p>
Human Readable Labels	<p>Users interact with labels as strings. Graphical user interfaces and command line interfaces present the strings as defined in <code>label_encodings(4)</code>. 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 <code>opaque_m_label_t</code>) by a subject whose label allows read (view) access to that label.</p>
Internal Text Labels	<p>In order to store labels in publicly accessible (<code>admin_low</code>) 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.</p>
Labels and Applications	<p>Applications interact with labels as opaque (<code>m_label_t</code>) structures. The semantics of these opaque structures are defined by a string to <code>m_label_t</code> translation. This translation is defined in <code>label_encodings(4)</code>. Various Application Programming Interfaces (API) translate between strings and <code>m_label_t</code> structures. Various APIs test access of subject-related labels to object-related labels.</p>
属性	<p>See attributes(5) for description of the following attributes:</p>

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 `bslabel_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.

另请参见

[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](#)》

附注

The functionality described on this manual page is available only if the system is configured with Trusted Extensions.

引用名 largefile – large file status of utilities

描述 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:

adb	aliasadm	awk	bdiff	cat
chgrp	chmod	chown	cksum	cmp
compress	cp	csch	csplit	cut
dd	dircmp	du	egrep	fgrep
file	find	ftp	getconf	grep
gzip	head	join	jsh	ksh88
ksh	ln	ls	mailcompat	mailstats
mdb	mkdir	mkfifo	more	mv
nawk	page	paste	pathchk	pg
praliases	rcp	remsh	rksh88	rksh
rm	rmdir	rsh	sed	sh
sort	split	sum	tail	tar
tee	test	touch	tr	uncompress
uudcode	uuencode	vacation	wc	zcat

The following `/usr/xpg4/bin` utilities are large file aware:

awk	cp	chgrp	chown	du
egrep	fgrep	file	grep	ln

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
mvdire	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:

<code>/usr/lib/autofs/automountd</code>	<code>/usr/sbin/mount</code>
<code>/usr/lib/nfs/rquotad</code>	

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

另请参见

[lf64\(5\)](#), [lfcompile\(5\)](#), [lfcompile64\(5\)](#)

引用名

ldap – LDAP as a naming repository

描述

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 shadowAccount	ou=people,dc=...
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.

文件

- `/var/ldap/ldap_client_cred`
- `/var/ldap/ldap_client_file` 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 `ldapclient(1M)` to update them.
- `/etc/nsswitch.conf` Configuration file for the name-service switch.
- `/etc/nsswitch.ldap` Sample configuration file for the name-service switch configured with LDAP and files.
- `/etc/pam.conf` PAM framework configuration file.

另请参见

`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》

<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>

off_t;	off64_t;
ino_t;	ino64_t;
blkcnt_t;	blkcnt64_t;
fsblkcnt_t;	fsblkcnt64_t;
fsfilcnt_t;	fsfilcnt64_t;

<unistd.h>

_LFS64_LARGEFILE
_LFS64_STDIO

<sys/unistd.h>

_CS_LFS64_CFLAGS
_CS_LFS64_LDFLAGS
_CS_LFS64_LIBS
_CS_LFS64_LINTFLAGS

System Interfaces The following tables display the standard API and the corresponding transitional interfaces for 64-bit file offsets. The interfaces are grouped by header. The interface name and the affected data types are displayed in courier font.

<aio.h>

int aio_cancel(..., struct aiocb *);	int aio_cancel64(..., struct aiocb64 *);
int aio_error(const struct aiocb *);	int aio_error64(const struct aiocb64 *);
int aio_fsync(..., struct aiocb *);	int aio_fsync64(..., struct aiocb64 *);
int aio_read(struct aiocb *);	int aio_read64(struct aiocb64 *);
int aio_return(struct aiocb *);	int aio_return64(struct aiocb64 *);
int aio_suspend(int aio_suspend64(

```

const struct aiocb *, ...);
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);
```

另请参见

[lfcompile\(5\)](#), [lfcompile64\(5\)](#)

引用名

lfcompile – large file compilation environment for 32-bit applications

描述

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\(5\)](#)).

示例

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.

示例 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)
```

示例 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          \
```

示例 3 Compile a program with a “default” `off_t` that uses `fseeko()` and `ftello()`.

```
$ c89 -D_LARGEFILE_SOURCE a.c
```

另请参见

[csh\(1\)](#), [getconf\(1\)](#), [ksh\(1\)](#), [yacc\(1\)](#), [sh\(1\)](#), [fseek\(3C\)](#), [ftell\(3C\)](#), [lf64\(5\)](#), [lfcompile64\(5\)](#), [standards\(5\)](#)

附注

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.

已知问题

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.

引用名

lfcompile64 – transitional compilation environment

描述

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 (2³¹ 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
LFS64_CFLAGS	obtain compilation flags necessary to enable the transitional compilation environment
LFS64_LDFLAGS	obtain link editor options
LFS64_LIBS	obtain link library names
LFS64_LINTFLAGS	obtain lint options

示例

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.

示例 1 An example of compiling a program using transitional interfaces such as `lseek64()` and `fopen64()`:

```
$ cc89 -D_LARGEFILE64_SOURCE      \
      $(getconf LFS64_CFLAGS) a.c \
      $(getconf LFS64_LDFLAGS)    \
      $(getconf LFS64_LIBS)
```

示例 2 An example of running `lint` on a program using transitional interfaces:

```
% lint -D_LARGEFILE64_SOURCE      \
      `getconf LFS64_LINTFLAG`S ... \
      `getconf LFS64_LIB`S
```

另请参见

[getconf\(1\)](#), [lseek\(2\)](#), [fopen\(3C\)](#), [lf64\(5\)](#), [standards\(5\)](#)

locale – subset of a user's environment that depends on language and cultural conventions

描述

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 `localdef(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 `localdef` 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 `set locale()` 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.

`alpha` Define characters to be classified as letters.

	<p>In the POSIX locale, all characters in the classes <code>upper</code> and <code>lower</code> are included.</p> <p>In a locale definition file, no character specified for the keywords <code>cntrl</code>, <code>digit</code>, <code>punct</code>, or <code>space</code> can be specified. Characters classified as either <code>upper</code> or <code>lower</code> are automatically included in this class.</p>
<code>digit</code>	<p>Define the characters to be classified as numeric digits.</p> <p>In the POSIX locale, only</p> <p>0 1 2 3 4 5 6 7 8 9</p> <p>are included.</p> <p>In a locale definition file, only the digits 0, 1, 2, 3, 4, 5, 6, 7, 8, and 9 can be specified, and in contiguous ascending sequence by numerical value. The digits 0 to 9 of the portable character set are automatically included in this class.</p> <p>The definition of character class <code>digit</code> requires that only ten characters; the ones defining digits can be specified; alternative digits (for example, Hindi or Kanji) cannot be specified here.</p>
<code>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>
<code>punct</code>	<p>Define characters to be classified as punctuation characters.</p>

	<p>In the POSIX locale, neither the space character nor any characters in classes <code>alpha</code>, <code>digit</code>, or <code>cntrl</code> are included.</p> <p>In a locale definition file, no character specified for the keywords <code>upper</code>, <code>lower</code>, <code>alpha</code>, <code>digit</code>, <code>cntrl</code>, <code>xdigit</code> or as the space character can be specified.</p>
<code>graph</code>	<p>Define characters to be classified as printable characters, not including the space character.</p> <p>In the POSIX locale, all characters in classes <code>alpha</code>, <code>digit</code>, and <code>punct</code> are included; no characters in class <code>cntrl</code> are included.</p> <p>In a locale definition file, characters specified for the keywords <code>upper</code>, <code>lower</code>, <code>alpha</code>, <code>digit</code>, <code>xdigit</code>, and <code>punct</code> are automatically included in this class. No character specified for the keyword <code>cntrl</code> can be specified.</p>
<code>print</code>	<p>Define characters to be classified as printable characters, including the space character.</p> <p>In the POSIX locale, all characters in class <code>graph</code> are included; no characters in class <code>cntrl</code> are included.</p> <p>In a locale definition file, characters specified for the keywords <code>upper</code>, <code>lower</code>, <code>alpha</code>, <code>digit</code>, <code>xdigit</code>, <code>punct</code>, and the space character are automatically included in this class. No character specified for the keyword <code>cntrl</code> can be specified.</p>
<code>xdigit</code>	<p>Define the characters to be classified as hexadecimal digits.</p> <p>In the POSIX locale, only:</p> <p><code>0 1 2 3 4 5 6 7 8 9 A B C D E F a b c d e f</code></p> <p>are included.</p> <p>In a locale definition file, only the characters defined for the class <code>digit</code> can be specified, in contiguous ascending sequence by numerical value, followed by one or more sets of six characters representing the hexadecimal digits 10 to 15 inclusive, with each set in ascending order (for example <code>A, B, C, D, E, F, a, b, c, d, e, f</code>). The digits <code>0</code> to <code>9</code>, the upper-case letters <code>A</code> to <code>F</code> and the lower-case letters <code>a</code> to <code>f</code> of the portable character set are automatically included in this class.</p> <p>The definition of character class <code>xdigit</code> requires that the characters included in character class <code>digit</code> be included here also.</p>
<code>blank</code>	<p>Define characters to be classified as blank characters.</p>

In the POSIX locale, only the space and tab characters are included.

In a locale definition file, the characters space and tab are automatically included in this class.

`charclass` 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 `LC_CTYPE` definition. A character class name consists of at least one and at most `{CHARCLASS_NAME_MAX}` 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 `LC_CTYPE` keywords defined in this document.

`charclass-name` 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 `charclass` 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 `charclass-name` can be used as the *property* argument to the [wctype\(3C\)](#) function, in regular expression and shell pattern-matching bracket expressions, and by the [tr\(1\)](#) command.

`toupper` Define the mapping of lower-case letters to upper-case letters.

In the POSIX locale, at a minimum, the 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

are mapped to the corresponding 26 upper-case characters:

A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

In a locale definition file, the operand consists of character pairs, separated by semicolons. The characters in each character pair are separated by a comma and the pair enclosed by parentheses. The first character in each pair is the lower-case letter, the second the corresponding upper-case letter. Only characters specified for the keywords `lower` and `upper` can be specified. The lower-case letters a to z, and their corresponding upper-case letters A to Z, of the portable character set are automatically included in this mapping, but only when the `toupper` keyword is omitted from the locale definition.

`tolower` Define the mapping of upper-case letters to lower-case letters.

In the POSIX locale, at a minimum, the 26 upper-case characters:

A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

are mapped to the corresponding 26 lower-case characters:

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

In a locale definition file, the operand consists of character pairs, separated by semicolons. The characters in each character pair are separated by a comma and the pair enclosed by parentheses. The first character in each pair is the upper-case letter, the second the corresponding lower-case letter. Only characters specified for the keywords `lower` and `upper` can be specified. If the `tolower` keyword is omitted from the locale definition, the mapping will be the reverse mapping of the one specified for `toupper`.

LC_COLLATE

The `LC_COLLATE` category provides a collation sequence definition for numerous utilities (such as `sort(1)`, `uniq(1)`, and so forth), regular expression matching (see `regex(5)`), and the `strcoll(3C)`, `strxfrm(3C)`, `wcscoll(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 recompared according to the relative subsequent weights, until either a pair of collating elements compare unequal or the weights are exhausted.

The following keywords are recognized in a collation sequence definition. They are described in detail in the following sections.

	copy	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.
	collating-element	Define a collating-element symbol representing a multi-character collating element. This keyword is optional.
	collating-symbol	Define a collating symbol for use in collation order statements. This keyword is optional.
	order_start	Define collation rules. This statement is followed by one or more collation order statements, assigning character collation values and collation weights to collating elements.
	order_end	Specify the end of the collation-order statements.
collating-element keyword		<p>In addition to the collating elements in the character set, the <code>collating-element</code> keyword is used to define multi-character collating elements. The syntax is:</p> <pre>"collating-element %s from \"%s\"\\n", <collating-symbol>, <string></pre> <p>The <code><collating-symbol></code> 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 <code><collating-element></code> defined via this keyword is only recognized with the LC_COLLATE category.</p> <p>Example:</p> <pre>collating-element <ch> from "<c><h>" collating-element <e-acute> from "<acute><e>" collating-element <ll> from "\\l"</pre>
collating-symbol keyword		<p>This keyword will be used to define symbols for use in collation sequence statements; that is, between the <code>order_start</code> and the <code>order_end</code> keywords. The syntax is:</p> <pre>"collating-symbol %s\\n", <collating-symbol></pre> <p>The <code><collating-symbol></code> 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.</p> <p>A <code>collating-symbol</code> defined via this keyword is only recognized with the LC_COLLATE category.</p> <p>Example:</p>

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;...;%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.

LC_MONETARY	<p>The LC_MONETARY category defines the rules and symbols that are used to format monetary numeric information. This information is available through the <code>localeconv(3C)</code> function</p> <p>The following items are defined in this category of the locale. The item names are the keywords recognized by the <code>localedef(1)</code> utility when defining a locale. They are also similar to the member names of the <code>lconv</code> structure defined in <code><locale.h></code>. The <code>localeconv</code> function returns <code>{CHAR_MAX}</code> for unspecified integer items and the empty string <code>("")</code> for unspecified or size zero string items.</p> <p>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 <code>("")</code>, or integer keywords set to <code>-1</code>, are used to indicate that the value is not available in the locale.</p> <table> <tr> <td><code>int_curr_symbol</code></td><td>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.</td></tr> <tr> <td><code>currency_symbol</code></td><td>The string used as the local currency symbol.</td></tr> <tr> <td><code>mon_decimal_point</code></td><td>The operand is a string containing the symbol that is used as the decimal delimiter (radix character) in monetary formatted quantities.</td></tr> <tr> <td><code>mon_thousands_sep</code></td><td>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.</td></tr> <tr> <td><code>mon_grouping</code></td><td> <p>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 <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.</p> <p>The following is an example of the interpretation of the <code>mon_grouping</code> keyword. Assuming that the value to be formatted is <code>123456789</code> and the <code>mon_thousands_sep</code> is <code>' '</code>, then the following table shows the result. The third column shows the equivalent string in the ISO C standard that would be used by the <code>localeconv</code> function to accommodate this grouping.</p> </td></tr> </table>	<code>int_curr_symbol</code>	The international currency symbol. The operand is a four-character string, with the first three characters containing the alphabetic international currency symbol in accordance with those specified in the ISO 4217 standard. The fourth character is the character used to separate the international currency symbol from the monetary quantity.	<code>currency_symbol</code>	The string used as the local currency symbol.	<code>mon_decimal_point</code>	The operand is a string containing the symbol that is used as the decimal delimiter (radix character) in monetary formatted quantities.	<code>mon_thousands_sep</code>	The operand is a string containing the symbol that is used as a separator for groups of digits to the left of the decimal delimiter in formatted monetary quantities.	<code>mon_grouping</code>	<p>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 <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.</p> <p>The following is an example of the interpretation of the <code>mon_grouping</code> keyword. Assuming that the value to be formatted is <code>123456789</code> and the <code>mon_thousands_sep</code> is <code>' '</code>, then the following table shows the result. The third column shows the equivalent string in the ISO C standard that would be used by the <code>localeconv</code> function to accommodate this grouping.</p>
<code>int_curr_symbol</code>	The international currency symbol. The operand is a four-character string, with the first three characters containing the alphabetic international currency symbol in accordance with those specified in the ISO 4217 standard. The fourth character is the character used to separate the international currency symbol from the monetary quantity.										
<code>currency_symbol</code>	The string used as the local currency symbol.										
<code>mon_decimal_point</code>	The operand is a string containing the symbol that is used as the decimal delimiter (radix character) in monetary formatted quantities.										
<code>mon_thousands_sep</code>	The operand is a string containing the symbol that is used as a separator for groups of digits to the left of the decimal delimiter in formatted monetary quantities.										
<code>mon_grouping</code>	<p>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 <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.</p> <p>The following is an example of the interpretation of the <code>mon_grouping</code> keyword. Assuming that the value to be formatted is <code>123456789</code> and the <code>mon_thousands_sep</code> is <code>' '</code>, then the following table shows the result. The third column shows the equivalent string in the ISO C standard that would be used by the <code>localeconv</code> function to accommodate this grouping.</p>										

mon_grouping	Formatted Value	ISO C String
3;-1	123456'789	"\3\177"
3	123'456'789	"\3"
3;2;-1	1234'56'789	"\3\2\177"
3;2	12'34'56'789	"\3\2"
-1	1234567898	"\177"

In these examples, the octal value of {CHAR_MAX} is 177.

positive_sign	A string used to indicate a non-negative-valued formatted monetary quantity.
negative_sign	A string used to indicate a negative-valued formatted monetary quantity.
int_frac_digits	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 int_curr_symbol.
frac_digits	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 currency_symbol.
p_cs_precedes	<p>In an application conforming to the SUSv3 standard, an integer set to 1 if the currency_symbol precedes the value for a monetary quantity with a non-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 currency_symbol or int_currency_symbol precedes the value for a monetary quantity with a non-negative value, and set to 0 if the symbol succeeds the value.</p>
p_sep_by_space	<p>In an application conforming to the SUSv3 standard, an integer set to 0 if no space separates the currency_symbol 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.</p> <p>In an application <i>not</i> conforming to the SUSv3 standard, an integer set to 0 if no space separates the currency_symbol or int_curr_symbol 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.</p>

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"> 0 Parentheses enclose the quantity and the <code>currency_symbol</code>. 1 The sign string precedes the quantity and the <code>currency_symbol</code>. 2 The sign string succeeds the quantity and the <code>currency_symbol</code>. 3 The sign string precedes the <code>currency_symbol</code>. 4 The sign string succeeds the <code>currency_symbol</code>. <p>In an application <i>not</i> conforming to the SUSv3 standard:</p> <ul style="list-style-type: none"> 0 Parentheses enclose the quantity and the <code>currency_symbol</code> or <code>int_curr_symbol</code>. 1 The sign string precedes the quantity and the <code>currency_symbol</code> or <code>int_curr_symbol</code>.

	<div>2 The sign string succeeds the quantity and the currency_symbol or int_curr_symbol.</div> <div>3 The sign string precedes the currency_symbol or int_curr_symbol.</div> <div>4 The sign string succeeds the currency_symbol or int_curr_symbol.</div>
n_sign_posn	An integer set to a value indicating the positioning of the negative_sign for a negative formatted monetary quantity.
int_p_cs_precedes	An integer set to 1 if the int_curr_symbol precedes the value for a monetary quantity with a non-negative value, and set to 0 if the symbol succeeds the value.
int_n_cs_precedes	An integer set to 1 if the int_curr_symbol precedes the value for a monetary quantity with a negative value, and set to 0 if the symbol succeeds the value.
int_p_sep_by_space	An integer set to 0 if no space separates the int_curr_symbol 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.
int_n_sep_by_space	An integer set to 0 if no space separates the int_curr_symbol 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.
int_p_sign_posn	<div>An integer set to a value indicating the positioning of the positive_sign for a positive monetary quantity formatted with the international format. The following integer values are recognized for int_p_sign_posn and int_n_sign_posn:</div> <div>0 Parentheses enclose the quantity and the int_curr_symbol.</div> <div>1 The sign string precedes the quantity and the int_curr_symbol.</div> <div>2 The sign string precedes the quantity and the int_curr_symbol.</div> <div>3 The sign string precedes the int_curr_symbol.</div> <div>4 The sign string succeeds the int_curr_symbol.</div>
int_n_sign_posn	An integer set to a value indicating the positioning of the negative_sign 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 [strftime\(3C\)](#), [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 `strftime()`, `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	<p>Define how years are counted and displayed for each era in a locale. The operand consists of semicolon-separated strings. Each string is an era description segment with the format:</p> <p><i>direction:offset:start_date:end_date:era_name:era_format</i></p> <p>according to the definitions below. There can be as many era description segments as are necessary to describe the different eras.</p> <p>The start of an era might not be the earliest point. For example, the Christian era B.C. starts on the day before January 1, A.D. 1, and increases with earlier time.</p>

	<p><i>direction</i> Either a + or a – character. The + character indicates that years closer to the <i>start_date</i> have lower numbers than those closer to the <i>end_date</i>. The – character indicates that years closer to the <i>start_date</i> have higher numbers than those closer to the <i>end_date</i>.</p> <p><i>offset</i> The number of the year closest to the <i>start_date</i> in the era, corresponding to the %Eg and %Ey field descriptors.</p> <p><i>start_date</i> A date in the form <i>yyyy/mm/dd</i>, where <i>yyyy</i>, <i>mm</i>, and <i>dd</i> are the year, month and day numbers respectively of the start of the era. Years prior to A.D. 1 are represented as negative numbers.</p> <p><i>end_date</i> The ending date of the era, in the same format as the <i>start_date</i>, or one of the two special values –* or +*. The value –* indicates that the ending date is the beginning of time. The value +* indicates that the ending date is the end of time.</p> <p><i>era_name</i> A string representing the name of the era, corresponding to the %EC field descriptor.</p> <p><i>era_format</i> A string for formatting the year in the era, corresponding to the %EG and %EY field descriptors.</p>
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	<p>The following information can be accessed. These correspond to constants defined in <code><langinfo.h></code> and used as arguments to the nl_langinfo(3C) function.</p> <p>ABDAY_x The abbreviated weekday names (for example Sun), where <i>x</i> is a number from 1 to 7.</p>

DAY_x	The full weekday names (for example Sunday), where <i>x</i> is a number from 1 to 7.								
ABMON_x	The abbreviated month names (for example Jan), where <i>x</i> is a number from 1 to 12.								
MON_x	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	<p>The era description segments, which describe how years are counted and displayed for each era in a locale. Each era description segment has the format:</p> <p><i>direction:offset:start_date:end_date:era_name:era_format</i></p> <p>according to the definitions below. There will be as many era description segments as are necessary to describe the different eras. Era description segments are separated by semicolons.</p> <p>The start of an era might not be the earliest point For example, the Christian era B.C. starts on the day before January 1, A.D. 1, and increases with earlier time.</p> <table><tr><td><i>direction</i></td><td>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>.</td></tr><tr><td><i>offset</i></td><td>The number of the year closest to the <i>start_date</i> in the era.</td></tr><tr><td><i>start_date</i></td><td>A date in the form <i>yyyy/mm/dd</i>, where <i>yyyy</i>, <i>mm</i>, and <i>dd</i> are the year, month and day numbers respectively of the start of the era. Years prior to AD 1 are represented as negative numbers.</td></tr><tr><td><i>end_date</i></td><td>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</td></tr></table>	<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.	<i>start_date</i>	A date in the form <i>yyyy/mm/dd</i> , where <i>yyyy</i> , <i>mm</i> , and <i>dd</i> are the year, month and day numbers respectively of the start of the era. Years prior to AD 1 are represented as negative numbers.	<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
<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.								
<i>start_date</i>	A date in the form <i>yyyy/mm/dd</i> , where <i>yyyy</i> , <i>mm</i> , and <i>dd</i> are the year, month and day numbers respectively of the start of the era. Years prior to AD 1 are represented as negative numbers.								
<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.

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 `strftime(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_AMP	%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
yesexpr	YESEXPR	"^[yY]"
noexpr	NOEXPR	"^[nN]"
yesstr	YESSTR	"yes"
nostr	NOSTR	"no"

In an application conforming to the SUSv3 standard, the information on `yesstr` and `nostr` is not available.

另请参见

`date(1)`, `locale(1)`, `localedef(1)`, `sort(1)`, `tr(1)`, `uniq(1)`, `localeconv(3C)`, `nl_langinfo(3C)`, `setlocale(3C)`, `strcoll(3C)`, `strftime(3C)`, `strptime(3C)`, `strxfrm(3C)`, `wscoll(3C)`, `wcsftime(3C)`, `wcsxfrm(3C)`, `wctype(3C)`, `attributes(5)`, `charmap(5)`, `extensions(5)`, `regex(5)`

引用名 locale_alias – locale name aliases and their corresponding canonical locale names

描述 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 at below:

Locale Name Alias	Canonical Locale Name
-----	-----
AR_AA	ar_AA.UTF-8
AR_AA.UTF-8	ar_AA.UTF-8
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_LB	ar_LB.UTF-8
AR_LB.UTF-8	ar_LB.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_SY	ar_SY.UTF-8
AR_SY.UTF-8	ar_SY.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
CY_GB	cy_GB.UTF-8
CY_GB.UTF-8	cy_GB.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	el_GR.UTF-8
EL_GR.UTF-8	el_GR.UTF-8
EN_AU	en_AU.UTF-8
EN_AU.UTF-8	en_AU.UTF-8
EN_BE	en_BE.UTF-8
EN_BE.UTF-8	en_BE.UTF-8
EN_BE.UTF-8@euro	en_BE.UTF-8
EN_BE@euro	en_BE.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
EN_ZA	en_ZA.UTF-8
EN_ZA.UTF-8	en_ZA.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
NO_NO	no_NO.UTF-8
NO_NO.UTF-8	no_NO.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
SH_SP	sh_SP.UTF-8
SH_SP.UTF-8	sh_SP.UTF-8
SH_YU	sh_YU.UTF-8
SH_YU.UTF-8	sh_YU.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
SR_SP	sr_SP.UTF-8
SR_SP.UTF-8	sr_SP.UTF-8
SR_YU	sr_YU.UTF-8
SR_YU.UTF-8	sr_YU.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
UR_IN	ur_IN.UTF-8
UR_IN.UTF-8	ur_IN.UTF-8
UR_PK	ur_PK.UTF-8
UR_PK.UTF-8	ur_PK.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
aa_DJ	aa_DJ.ISO8859-1
aa_DJ.iso88591	aa_DJ.ISO8859-1
aa_DJ.utf8	aa_DJ.UTF-8
aa_ER	aa_ER.UTF-8
aa_ER.utf8	aa_ER.UTF-8
aa_ET	aa_ET.UTF-8
aa_ET.utf8	aa_ET.UTF-8
af_ZA	af_ZA.ISO8859-1
af_ZA.iso88591	af_ZA.ISO8859-1
af_ZA.utf8	af_ZA.UTF-8
am_ET	am_ET.UTF-8
am_ET.utf8	am_ET.UTF-8
an_ES	an_ES.ISO8859-15
an_ES.iso885915	an_ES.ISO8859-15
an_ES.utf8	an_ES.UTF-8
ar	ar_EG.ISO8859-6
ar_AA	ar_AA.ISO8859-6
ar_AE	ar_AE.ISO8859-6
ar_AE.iso88596	ar_AE.ISO8859-6
ar_AE.utf8	ar_AE.UTF-8
ar_BH	ar_BH.ISO8859-6
ar_BH.iso88596	ar_BH.ISO8859-6
ar_BH.utf8	ar_BH.UTF-8
ar_DZ	ar_DZ.ISO8859-6
ar_DZ.iso88596	ar_DZ.ISO8859-6
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_IN	ar_IN.UTF-8
ar_IN.utf8	ar_IN.UTF-8
ar_IQ	ar_IQ.ISO8859-6
ar_IQ.iso88596	ar_IQ.ISO8859-6
ar_IQ.utf8	ar_IQ.UTF-8
ar_JO	ar_JO.ISO8859-6
ar_JO.iso88596	ar_JO.ISO8859-6
ar_JO.utf8	ar_JO.UTF-8
ar_KW	ar_KW.ISO8859-6
ar_KW.iso88596	ar_KW.ISO8859-6
ar_KW.utf8	ar_KW.UTF-8
ar_LB	ar_LB.ISO8859-6
ar_LB.iso88596	ar_LB.ISO8859-6
ar_LB.utf8	ar_LB.UTF-8
ar_LY	ar_LY.ISO8859-6
ar_LY.iso88596	ar_LY.ISO8859-6
ar_LY.utf8	ar_LY.UTF-8

ar_MA	ar_MA.ISO8859-6
ar_MA.iso88596	ar_MA.ISO8859-6
ar_MA.utf8	ar_MA.UTF-8
ar_OM	ar_OM.ISO8859-6
ar_OM.iso88596	ar_OM.ISO8859-6
ar_OM.utf8	ar_OM.UTF-8
ar_QA	ar_QA.ISO8859-6
ar_QA.iso88596	ar_QA.ISO8859-6
ar_QA.utf8	ar_QA.UTF-8
ar_SA	ar_SA.ISO8859-6
ar_SA.iso88596	ar_SA.ISO8859-6
ar_SA.utf8	ar_SA.UTF-8
ar_SD	ar_SD.ISO8859-6
ar_SD.iso88596	ar_SD.ISO8859-6
ar_SD.utf8	ar_SD.UTF-8
ar_SY	ar_SY.ISO8859-6
ar_SY.iso88596	ar_SY.ISO8859-6
ar_SY.utf8	ar_SY.UTF-8
ar_TN	ar_TN.ISO8859-6
ar_TN.iso88596	ar_TN.ISO8859-6
ar_TN.utf8	ar_TN.UTF-8
ar_YE	ar_YE.ISO8859-6
ar_YE.iso88596	ar_YE.ISO8859-6
ar_YE.utf8	ar_YE.UTF-8
as_IN.utf8	as_IN.UTF-8
az_AZ.utf8	az_AZ.UTF-8
be_BY	be_BY.ISO8859-5
be_BY.cp1251	be_BY.ANSI1251
be_BY.utf8	be_BY.UTF-8
bg_BG	bg_BG.ISO8859-5
bg_BG.cp1251	bg_BG.ANSI1251
bg_BG.utf8	bg_BG.UTF-8
bn_BD	bn_BD.UTF-8
bn_BD.utf8	bn_BD.UTF-8
bn_IN	bn_IN.UTF-8
bn_IN.utf8	bn_IN.UTF-8
br_FR	br_FR.ISO8859-1
br_FR.iso88591	br_FR.ISO8859-1
br_FR.iso885915@euro	br_FR.ISO8859-15
br_FR.utf8	br_FR.UTF-8
br_FR@euro	br_FR.ISO8859-15
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_AD	ca_AD.ISO8859-15
ca_AD.iso885915	ca_AD.ISO8859-15
ca_AD.utf8	ca_AD.UTF-8

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.IBM-1252	ca_ES.ANSI1252
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
ca_FR	ca_FR.ISO8859-15
ca_FR.iso885915	ca_FR.ISO8859-15
ca_FR.utf8	ca_FR.UTF-8
ca_IT	ca_IT.ISO8859-15
ca_IT.iso885915	ca_IT.ISO8859-15
ca_IT.utf8	ca_IT.UTF-8
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
cy_GB	cy_GB.ISO8859-14
cy_GB.iso885914	cy_GB.ISO8859-14
cy_GB.utf8	cy_GB.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	de_BE.ISO8859-1
de_BE.iso88591	de_BE.ISO8859-1
de_BE.iso885915@euro	de_BE.ISO8859-15
de_BE.utf8	de_BE.UTF-8
de_BE@euro	de_BE.ISO8859-15
de_CH	de_CH.ISO8859-1
de_CH.8859-15	de_CH.ISO8859-15
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.IBM-1252	de_DE.ANSI1252
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	de_LU.ISO8859-15
de_LU.8859-15	de_LU.ISO8859-15
de_LU.8859-15@euro	de_LU.ISO8859-15
de_LU.iso88591	de_LU.ISO8859-1
de_LU.iso885915@euro	de_LU.ISO8859-15
de_LU.utf8	de_LU.UTF-8
de_LU@euro	de_LU.ISO8859-15
dz_BT	dz_BT.UTF-8
dz_BT.utf8	dz_BT.UTF-8
el	el_GR.ISO8859-7
el.UTF-8	el_CY.UTF-8
el.sun_eu_greek	el_GR.ISO8859-7
el_CY	el_CY.ISO8859-7
el_CY.iso88597	el_CY.ISO8859-7
el_CY.utf8	el_CY.UTF-8
el_GR	el_GR.ISO8859-7
el_GR.ISO8859-7@euro	el_GR.ISO8859-7
el_GR.iso88597	el_GR.ISO8859-7
el_GR.utf8	el_GR.UTF-8
en_AU	en_AU.ISO8859-1
en_AU.8859-15	en_AU.ISO8859-15
en_AU.iso88591	en_AU.ISO8859-1
en_AU.utf8	en_AU.UTF-8
en_BE	en_BE.ISO8859-15
en_BE.8859-15	en_BE.ISO8859-15
en_BE.8859-15@euro	en_BE.ISO8859-15
en_BE@euro	en_BE.ISO8859-15
en_BW	en_BW.ISO8859-1
en_BW.iso88591	en_BW.ISO8859-1
en_BW.utf8	en_BW.UTF-8
en_CA	en_CA.ISO8859-1
en_CA.8859-15	en_CA.ISO8859-15
en_CA.iso88591	en_CA.ISO8859-1
en_CA.utf8	en_CA.UTF-8
en_DK	en_DK.ISO8859-1
en_DK.iso88591	en_DK.ISO8859-1
en_DK.utf8	en_DK.UTF-8
en_GB	en_GB.ISO8859-1
en_GB.8859-15	en_GB.ISO8859-15

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	en_HK.ISO8859-15
en_HK.8859-15	en_HK.ISO8859-15
en_HK.iso88591	en_HK.ISO8859-1
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	en_IN.ISO8859-15
en_IN.8859-15	en_IN.ISO8859-15
en_IN.utf8	en_IN.UTF-8
en_NZ	en_NZ.ISO8859-1
en_NZ.8859-15	en_NZ.ISO8859-15
en_NZ.iso88591	en_NZ.ISO8859-1
en_NZ.utf8	en_NZ.UTF-8
en_PH	en_PH.ISO8859-15
en_PH.8859-15	en_PH.ISO8859-15
en_PH.iso88591	en_PH.ISO8859-1
en_PH.utf8	en_PH.UTF-8
en_SG	en_SG.ISO8859-15
en_SG.8859-15	en_SG.ISO8859-15
en_SG.iso88591	en_SG.ISO8859-1
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_ZA	en_ZA.ISO8859-15
en_ZA.8859-15	en_ZA.ISO8859-15
en_ZA.iso88591	en_ZA.ISO8859-1
en_ZA.utf8	en_ZA.UTF-8
en_ZW	en_ZW.ISO8859-1
en_ZW.iso88591	en_ZW.ISO8859-1
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.8859-15	es_AR.ISO8859-15
es_AR.iso88591	es_AR.ISO8859-1
es_AR.utf8	es_AR.UTF-8
es_BO	es_BO.ISO8859-1
es_BO.8859-15	es_BO.ISO8859-15
es_BO.iso88591	es_BO.ISO8859-1
es_BO.utf8	es_BO.UTF-8
es_CL	es_CL.ISO8859-1
es_CL.8859-15	es_CL.ISO8859-15
es_CL.iso88591	es_CL.ISO8859-1
es_CL.utf8	es_CL.UTF-8
es_CO	es_CO.ISO8859-1
es_CO.8859-15	es_CO.ISO8859-15
es_CO.iso88591	es_CO.ISO8859-1
es_CO.utf8	es_CO.UTF-8
es_CR	es_CR.ISO8859-1
es_CR.8859-15	es_CR.ISO8859-15
es_CR.iso88591	es_CR.ISO8859-1
es_CR.utf8	es_CR.UTF-8
es_DO	es_DO.ISO8859-15
es_DO.8859-15	es_DO.ISO8859-15
es_DO.iso88591	es_DO.ISO8859-1
es_DO.utf8	es_DO.UTF-8
es_EC	es_EC.ISO8859-1
es_EC.8859-15	es_EC.ISO8859-15
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.IBM-1252	es_ES.ANSI1252
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.8859-15	es_GT.ISO8859-15
es_GT.iso88591	es_GT.ISO8859-1
es_GT.utf8	es_GT.UTF-8
es_HN	es_HN.ISO8859-15
es_HN.8859-15	es_HN.ISO8859-15
es_HN.iso88591	es_HN.ISO8859-1
es_HN.utf8	es_HN.UTF-8
es_MX	es_MX.ISO8859-1
es_MX.8859-15	es_MX.ISO8859-15
es_MX.iso88591	es_MX.ISO8859-1
es_NI.utf8	es_NI.UTF-8

es_PA	es_PA.ISO8859-1
es_PA.8859-15	es_PA.ISO8859-15
es_PA.iso88591	es_PA.ISO8859-1
es_PA.utf8	es_PA.UTF-8
es_PE	es_PE.ISO8859-1
es_PE.8859-15	es_PE.ISO8859-15
es_PE.iso88591	es_PE.ISO8859-1
es_PE.utf8	es_PE.UTF-8
es_PR	es_PR.ISO8859-15
es_PR.8859-15	es_PR.ISO8859-15
es_PR.iso88591	es_PR.ISO8859-1
es_PR.utf8	es_PR.UTF-8
es_PY	es_PY.ISO8859-1
es_PY.8859-15	es_PY.ISO8859-15
es_PY.iso88591	es_PY.ISO8859-1
es_PY.utf8	es_PY.UTF-8
es_SV	es_SV.ISO8859-1
es_SV.8859-15	es_SV.ISO8859-15
es_SV.iso88591	es_SV.ISO8859-1
es_SV.utf8	es_SV.UTF-8
es_US	es_US.ISO8859-15
es_US.8859-15	es_US.ISO8859-15
es_US.iso88591	es_US.ISO8859-1
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.8859-15	es_VE.ISO8859-15
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.iso88591	et_EE.ISO8859-1
et_EE.iso885915	et_EE.ISO8859-15
et_EE.utf8	et_EE.UTF-8
eu_ES	eu_ES.ISO8859-1
eu_ES.iso88591	eu_ES.ISO8859-1
eu_ES.iso885915@euro	eu_ES.ISO8859-15
eu_ES.utf8	eu_ES.UTF-8
eu_ES@euro	eu_ES.ISO8859-15
fa_IR	fa_IR.UTF-8
fa_IR.utf8	fa_IR.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.IBM-1252	fi_FI.ANSI1252
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
fo_FO	fo_FO.ISO8859-1
fo_FO.iso88591	fo_FO.ISO8859-1
fo_FO.utf8	fo_FO.UTF-8
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.IBM-1252	fr_BE.ANSI1252
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.8859-15	fr_CA.ISO8859-15
fr_CA.iso88591	fr_CA.ISO8859-1
fr_CA.utf8	fr_CA.UTF-8
fr_CH	fr_CH.ISO8859-1
fr_CH.8859-15	fr_CH.ISO8859-15
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.IBM-1252	fr_FR.ANSI1252
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	fr_LU.ISO8859-15
fr_LU.8859-15	fr_LU.ISO8859-15
fr_LU.8859-15@euro	fr_LU.ISO8859-15
fr_LU.iso88591	fr_LU.ISO8859-1
fr_LU.iso885915@euro	fr_LU.ISO8859-15
fr_LU.utf8	fr_LU.UTF-8
fr_LU@euro	fr_LU.ISO8859-15
fy_NL	fy_NL.UTF-8
fy_NL.utf8	fy_NL.UTF-8
ga_IE	ga_IE.ISO8859-1

ga_IE.iso88591	ga_IE.IS08859-1
ga_IE.iso885915@euro	ga_IE.IS08859-15
ga_IE.utf8	ga_IE.UTF-8
ga_IE@euro	ga_IE.IS08859-15
gd_GB	gd_GB.IS08859-15
gd_GB.iso885915	gd_GB.IS08859-15
gd_GB.utf8	gd_GB.UTF-8
gl_ES	gl_ES.IS08859-1
gl_ES.iso88591	gl_ES.IS08859-1
gl_ES.iso885915@euro	gl_ES.IS08859-15
gl_ES.utf8	gl_ES.UTF-8
gl_ES@euro	gl_ES.IS08859-15
gu_IN	gu_IN.UTF-8
gu_IN.utf8	gu_IN.UTF-8
gv_GB	gv_GB.IS08859-1
gv_GB.iso88591	gv_GB.IS08859-1
gv_GB.utf8	gv_GB.UTF-8
he	he_IL.IS08859-8
he_IL	he_IL.IS08859-8
he_IL.iso88598	he_IL.IS08859-8
he_IL.utf8	he_IL.UTF-8
hi_IN	hi_IN.UTF-8
hi_IN.utf8	hi_IN.UTF-8
hr_HR	hr_HR.IS08859-2
hr_HR.iso88592	hr_HR.IS08859-2
hr_HR.utf8	hr_HR.UTF-8
hu	hu_HU.IS08859-2
hu_HU	hu_HU.IS08859-2
hu_HU.iso88592	hu_HU.IS08859-2
hu_HU.utf8	hu_HU.UTF-8
hy_AM	hy_AM.UTF-8
hy_AM.utf8	hy_AM.UTF-8
id_ID	id_ID.IS08859-15
id_ID.8859-15	id_ID.IS08859-15
id_ID.iso88591	id_ID.IS08859-1
id_ID.utf8	id_ID.UTF-8
is_IS	is_IS.IS08859-1
is_IS.8859-15	is_IS.IS08859-15
is_IS.iso88591	is_IS.IS08859-1
is_IS.utf8	is_IS.UTF-8
it	it_IT.IS08859-1
it.IS08859-15	it_IT.IS08859-15
it.UTF-8	it_IT.UTF-8
it_CH	it_CH.IS08859-15
it_CH.8859-15	it_CH.IS08859-15
it_CH.iso88591	it_CH.IS08859-1
it_CH.utf8	it_CH.UTF-8
it_IT	it_IT.IS08859-1

it_IT.8859-15	it_IT.ISO8859-15
it_IT.8859-15@euro	it_IT.ISO8859-15
it_IT.IBM-1252	it_IT.ANSI1252
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
iw_IL	iw_IL.ISO8859-8
iw_IL.iso88598	iw_IL.ISO8859-8
iw_IL.utf8	iw_IL.UTF-8
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
kl_GL	kl_GL.ISO8859-1
kl_GL.iso88591	kl_GL.ISO8859-1
kl_GL.utf8	kl_GL.UTF-8
km_KH	km_KH.UTF-8
km_KH.utf8	km_KH.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	ku_TR.ISO8859-9
ku_TR.iso88599	ku_TR.ISO8859-9
ku_TR.utf8	ku_TR.UTF-8
kw_GB	kw_GB.ISO8859-1
kw_GB.iso88591	kw_GB.ISO8859-1
kw_GB.utf8	kw_GB.UTF-8
ky_KG	ky_KG.UTF-8
ky_KG.utf8	ky_KG.UTF-8
lg_UG.utf8	lg_UG.UTF-8
lo_LA	lo_LA.UTF-8
lo_LA.utf8	lo_LA.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
mg_MG	mg_MG.ISO8859-15
mg_MG.iso885915	mg_MG.ISO8859-15
mg_MG.utf8	mg_MG.UTF-8
mi_NZ	mi_NZ.ISO8859-13
mi_NZ.iso885913	mi_NZ.ISO8859-13
mi_NZ.utf8	mi_NZ.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
mn_MN	mn_MN.UTF-8
mn_MN.utf8	mn_MN.UTF-8
mr_IN	mr_IN.UTF-8
mr_IN.utf8	mr_IN.UTF-8
ms_MY	ms_MY.ISO8859-15
ms_MY.8859-15	ms_MY.ISO8859-15
ms_MY.iso88591	ms_MY.ISO8859-1
ms_MY.utf8	ms_MY.UTF-8
mt_MT	mt_MT.ISO8859-3
mt_MT.iso88593	mt_MT.ISO8859-3
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
ne_NP	ne_NP.UTF-8
ne_NP.utf8	ne_NP.UTF-8
nl	nl_NL.ISO8859-1
nl.ISO8859-15	nl_NL.ISO8859-15
nl_BE	nl_BE.ISO8859-1
nl_BE.8859-15	nl_BE.ISO8859-15
nl_BE.8859-15@euro	nl_BE.ISO8859-15
nl_BE.IBM-1252	nl_BE.ANSI1252
nl_BE.ISO8859-15@euro	nl_BE.ISO8859-15
nl_BE.iso88591	nl_BE.ISO8859-1
nl_BE.utf8	nl_BE.UTF-8
nl_BE@euro	nl_BE.ISO8859-15
nl_NL	nl_NL.ISO8859-1
nl_NL.8859-15	nl_NL.ISO8859-15
nl_NL.8859-15@euro	nl_NL.ISO8859-15
nl_NL.IBM-1252	nl_NL.ANSI1252
nl_NL.ISO8859-15@euro	nl_NL.ISO8859-15
nl_NL.iso88591	nl_NL.ISO8859-1
nl_NL.utf8	nl_NL.UTF-8
nl_NL@euro	nl_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.8859-15	no_NO.ISO8859-15
no_NO.ISO8859-1@bokmal	nb_NO.ISO8859-1
no_NO.ISO8859-1@nynorsk	nn_NO.ISO8859-1
no_NO.iso88591	no_NO.ISO8859-1
no_NO.utf8	no_NO.UTF-8
no_NY	nn_NO.ISO8859-1
nr_ZA	nr_ZA.UTF-8
nr_ZA.utf8	nr_ZA.UTF-8
oc_FR	oc_FR.ISO8859-1
oc_FR.iso88591	oc_FR.ISO8859-1
oc_FR.utf8	oc_FR.UTF-8
om_ET	om_ET.UTF-8
om_ET.utf8	om_ET.UTF-8
om_KE	om_KE.ISO8859-1
om_KE.iso88591	om_KE.ISO8859-1
om_KE.utf8	om_KE.UTF-8
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
pa_PK	pa_PK.UTF-8
pa_PK.utf8	pa_PK.UTF-8
pl	pl_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.8859-15	pt_BR.ISO8859-15
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.IBM-1252	pt_PT.ANSI1252
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	ru_UA.KOI8-U
ru_UA.koi8u	ru_UA.KOI8-U
ru_UA.utf8	ru_UA.UTF-8
rw_RW	rw_RW.UTF-8
rw_RW.utf8	rw_RW.UTF-8
se_NO	se_NO.UTF-8
se_NO.utf8	se_NO.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
sh_SP	sh_SP.ISO8859-2
sh_YU	sh_YU.ISO8859-2
si_LK	si_LK.UTF-8
si_LK.utf8	si_LK.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
so_DJ	so_DJ.ISO8859-1
so_DJ.iso88591	so_DJ.ISO8859-1
so_DJ.utf8	so_DJ.UTF-8
so_ET	so_ET.UTF-8
so_ET.utf8	so_ET.UTF-8
so_KE	so_KE.ISO8859-1
so_KE.iso88591	so_KE.ISO8859-1
so_KE.utf8	so_KE.UTF-8
so_SO	so_SO.ISO8859-1
so_SO.iso88591	so_SO.ISO8859-1
so_SO.utf8	so_SO.UTF-8
sq_AL	sq_AL.ISO8859-2
sq_AL.8859-15	sq_AL.ISO8859-15
sq_AL.iso88591	sq_AL.ISO8859-1
sq_AL.utf8	sq_AL.UTF-8
sr_CS	sr_RS.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
sr_SP	sr_RS.ISO8859-5
sr_YU	sr_RS.ISO8859-5
sr_YU.ISO8859-5	sr_RS.ISO8859-5
ss_ZA	ss_ZA.UTF-8
ss_ZA.utf8	ss_ZA.UTF-8
st_ZA	st_ZA.ISO8859-1
st_ZA.iso88591	st_ZA.ISO8859-1
st_ZA.utf8	st_ZA.UTF-8
sv	sv_SE.ISO8859-1
sv.ISO8859-15	sv_SE.ISO8859-15
sv.UTF-8	sv_SE.UTF-8
sv_FI	sv_FI.ISO8859-1
sv_FI.iso88591	sv_FI.ISO8859-1
sv_FI.iso885915@euro	sv_FI.ISO8859-15
sv_FI.utf8	sv_FI.UTF-8
sv_FI@euro	sv_FI.ISO8859-15
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
tg_TJ	tg_TJ.KOI8-T
tg_TJ.koi8t	tg_TJ.KOI8-T
tg_TJ.utf8	tg_TJ.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
ti_ER	ti_ER.UTF-8
ti_ER.utf8	ti_ER.UTF-8
ti_ET	ti_ET.UTF-8
ti_ET.utf8	ti_ET.UTF-8
tl_PH	tl_PH.ISO8859-1
tl_PH.iso88591	tl_PH.ISO8859-1
tl_PH.utf8	tl_PH.UTF-8
tn_ZA	tn_ZA.UTF-8
tn_ZA.utf8	tn_ZA.UTF-8
tr	tr_TR.ISO8859-9
tr_CY	tr_CY.ISO8859-9
tr_CY.iso88599	tr_CY.ISO8859-9

tr_CY.utf8	tr_CY.UTF-8
tr_TR	tr_TR.ISO8859-9
tr_TR.iso88599	tr_TR.ISO8859-9
tr_TR.utf8	tr_TR.UTF-8
ts_ZA	ts_ZA.UTF-8
ts_ZA.utf8	ts_ZA.UTF-8
tt_RU.utf8	tt_RU.UTF-8
uk_UA	uk_UA.KOI8-U
uk_UA.koi8u	uk_UA.KOI8-U
uk_UA.utf8	uk_UA.UTF-8
ur_PK	ur_PK.UTF-8
ur_PK.utf8	ur_PK.UTF-8
uz_UZ	uz_UZ.ISO8859-1
uz_UZ.iso88591	uz_UZ.ISO8859-1
ve_ZA	ve_ZA.UTF-8
ve_ZA.utf8	ve_ZA.UTF-8
vi_VN	vi_VN.UTF-8
vi_VN.tcvn	vi_VN.TCVN5712-1
vi_VN.utf8	vi_VN.UTF-8
wa_BE	wa_BE.ISO8859-1
wa_BE.iso88591	wa_BE.ISO8859-1
wa_BE.iso885915@euro	wa_BE.ISO8859-15
wa_BE.utf8	wa_BE.UTF-8
wa_BE@euro	wa_BE.ISO8859-15
xh_ZA	xh_ZA.ISO8859-1
xh_ZA.iso88591	xh_ZA.ISO8859-1
xh_ZA.utf8	xh_ZA.UTF-8
yi_US	yi_US.ANSI1255
yi_US.cp1255	yi_US.ANSI1255
yi_US.utf8	yi_US.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	zh_SG.EUC
zh_SG.gb2312	zh_SG.EUC
zh_SG.gbk	zh_SG.GBK
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
zu_ZA	zu_ZA.ISO8859-1
zu_ZA.iso88591	zu_ZA.ISO8859-1
zu_ZA.utf8	zu_ZA.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.IS08859-1	es_EC
es_ES.IS08859-1	es, es_ES
es_ES.IS08859-15	es.IS08859-15, es_ES.IS08859-15@euro
es_ES.UTF-8	es.UTF-8, es_ES.UTF-8@euro
es_GT.IS08859-1	es_GT
es_MX.IS08859-1	es_MX
es_NI.IS08859-1	es_NI
es_PA.IS08859-1	es_PA
es_PE.IS08859-1	es_PE
es_PY.IS08859-1	es_PY
es_SV.IS08859-1	es_SV
es_UY.IS08859-1	es_UY
es_VE.IS08859-1	es_VE
et_EE.IS08859-15	et, et_EE
fi_FI.IS08859-1	fi, fi_FI
fi_FI.IS08859-15	fi.IS08859-15, fi_FI.IS08859-15@euro
fr_BE.IS08859-1	fr_BE
fr_BE.IS08859-15	fr_BE.IS08859-15@euro
fr_BE.UTF-8	fr_BE.UTF-8@euro
fr_CA.IS08859-1	fr_CA
fr_CH.IS08859-1	fr_CH
fr_FR.IS08859-1	fr, fr_FR
fr_FR.IS08859-15	fr.IS08859-15, fr_FR.IS08859-15@euro
fr_FR.UTF-8	fr.UTF-8, fr_FR.UTF-8@euro
he_IL.IS08859-8	he, he_IL
hr_HR.IS08859-2	hr_HR
hu_HU.IS08859-2	hu, hu_HU
is_IS.IS08859-1	is_IS
it_IT.IS08859-1	it, it_IT
it_IT.IS08859-15	it.IS08859-15, it_IT.IS08859-15@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.IS08859-13	lt, lt_LT
lv_LV.IS08859-13	lv, lv_LV
mk_MK.IS08859-5	mk_MK
nb_NO.IS08859-1	no, no_NO, no_NO.IS08859-1@bokmal
nl_BE.IS08859-1	nl_BE
nl_BE.IS08859-15	nl_BE.IS08859-15@euro
nl_NL.IS08859-1	nl, nl_NL
nl_NL.IS08859-15	nl.IS08859-15, nl_NL.IS08859-15@euro
nn_NO.IS08859-1	no_NO.IS08859-1@nynorsk, no_NY
pl_PL.IS08859-2	pl, pl_PL
pl_PL.UTF-8	pl.UTF-8
pt_BR.IS08859-1	pt_BR
pt_PT.IS08859-1	pt, pt_PT
pt_PT.IS08859-15	pt.IS08859-15, pt_PT.IS08859-15@euro

ro_R0.IS08859-2	ro_R0
ru_RU.IS08859-5	ru, ru_RU
ru_RU.KOI8-R	ru.koi8-r
ru_RU.UTF-8	ru.UTF-8
sk_SK.IS08859-2	sk_SK
sl_SI.IS08859-2	sl_SI
sq_AL.IS08859-2	sq_AL
sr_ME.IS08859-5	sr_SP, sr_YU, sr_YU.IS08859-5
sr_ME.UTF-8	sr_CS, sr_CS.UTF-8
sr_RS.IS08859-5	sr_SP, sr_YU, sr_YU.IS08859-5
sr_RS.UTF-8	sr_CS, sr_CS.UTF-8
sv_SE.IS08859-1	sv, sv_SE
sv_SE.IS08859-15	sv.IS08859-15
sv_SE.UTF-8	sv.UTF-8
th_TH.TIS620	th, th_TH, th_TH.IS08859-11
tr_TR.IS08859-9	tr, tr_TR
zh_CN.EUC	zh
zh_CN.GBK	zh.GBK
zh_CN.UTF-8	zh.UTF-8
zh_TW.EUC	zh_TW

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

另请参见 [gettext\(1\)](#), [catopen\(3C\)](#), [gettext\(3C\)](#), [setlocale\(3C\)](#), [attributes\(5\)](#), [environ\(5\)](#)

引用名

man – macros to format Reference Manual pages

用法概要

nroff -man *filename*...

troff -man *filename*...

描述

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>	
.B <i>t</i>	no	<i>t</i> =n.t.l.*	Text is in bold font.
.BI <i>t</i>	no	<i>t</i> =n.t.l.	Join words, alternating bold and italic.
.BR <i>t</i>	no	<i>t</i> =n.t.l.	Join words, alternating bold and roman.
.DT	no	.5i 1i...	Restore default tabs.
.HP <i>i</i>	yes	<i>i</i> =p.i.*	Begin paragraph with hanging indent. Set prevailing indent to <i>i</i> .
.I <i>t</i>	no	<i>t</i> =n.t.l.	Text is italic.
.IB <i>t</i>	no	<i>t</i> =n.t.l.	Join words, alternating italic and bold.
.IP <i>x i</i>	yes	<i>x</i> =""	Same as .TP with tag <i>x</i> .
.IR <i>t</i>	no	<i>t</i> =n.t.l.	Join words, alternating italic and roman.

<i>Request</i>	<i>Cause</i>	<i>If no</i>	<i>Explanation</i>
	<i>Break</i>	<i>Argument</i>	
.IX <i>t</i>	no	-	Index macro, for SunSoft internal use.
.LP	yes	-	Begin left-aligned paragraph. Set prevailing indent to .5i.
.P	yes	-	Same as .LP.
.PD <i>d</i>	no	<i>d</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, 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.

文件

`/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`

另请参见

[man\(1\)](#), [nroff\(1\)](#), [troff\(1\)](#), [whatis\(1\)](#)

Dale Dougherty and Tim O'Reilly, *Unix Text Processing*

引用名 mansun – macros to format Reference Manual pages

用法概要 nroff -mansun *filename*...

troff -mansun *filename*...

描述 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` ‘*®*’, ‘(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 1i...</code>	Restore default tabs.
<code>.HP i</code>	yes	<code>i=p.i.*</code>	Begin paragraph with hanging indent. Set prevailing indent to <i>i</i> .
<code>.I t</code>	no	<code>t=n.t.l.</code>	Text is italic.
<code>.IB t</code>	no	<code>t=n.t.l.</code>	Join words, alternating italic and bold.
<code>.IP x i</code>	yes	<code>x=""</code>	Same as <code>.TP</code> with tag <i>x</i> .

<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 <i>title</i> [1-8]	The name of the command or function, which serves as the title of the manual page. This is followed by the number of the section in which it appears.
.SH NAME	The name, or list of names, by which the command is called, followed by a dash and then a one-line summary of the action performed. All in Roman font, this section contains no <code>troff(1)</code> commands or escapes, and no macro requests. It is used to generate the <code>windex</code> database, which is used by the <code>whatis(1)</code> command.
.SH SYNOPSIS	<p>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.</p> <p>Syntactic symbols appear in Roman face:</p> <p>[] An argument, when surrounded by brackets is optional.</p> <p> Arguments separated by a vertical bar are exclusive. You can supply only one item from such a list.</p> <p>. . . Arguments followed by an ellipsis can be repeated. When an ellipsis follows a bracketed set, the expression within the brackets can be repeated.</p> <p>Functions: If required, the data declaration, or <code>#include</code> directive, is shown first, followed by the function declaration. Otherwise, the function declaration is shown.</p>
.SH DESCRIPTION	<p>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?"</p> <p>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.</p> <p>If a command interprets either subcommands or an input grammar, its command interface or input grammar is normally described in a <code>USAGE</code></p>

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.

文件

`/usr/share/lib/tmac/ansun`
`/usr/share/man/windex`

另请参见

[man\(1\)](#), [nroff\(1\)](#), [troff\(1\)](#), [whatis\(1\)](#)
Dale Dougherty and Tim O'Reilly, *Unix Text Processing*

引用名 me – macros for formatting papers

用法概要 nroff -me [*options*] *filename*...

troff -me [*options*] *filename*...

描述 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 Value</i>	<i>Cause Break</i>	<i>Explanation</i>
. (c	-	yes	Begin centered block.
. (d	-	no	Begin delayed text.
. (f	-	no	Begin footnote.
. (l	-	yes	Begin list.
. (q	-	yes	Begin major quote.
. (xx	-	no	Begin indexed item in index <i>x</i> .
. (z	-	no	Begin floating keep.

<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.
. ++ m H	-	no	Define paper section. m 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 T	-	yes	Begin chapter (or appendix, etc., as set by . ++). T 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 x y	-	yes	Precede equation; break out and add space. Equation number is y. The optional argument x may be I to indent equation (default), L 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 pic picture.
.PS	-	yes	Begin pic picture.
.TE	-	yes	End table.
.TH	-	yes	End heading section of table.
.TS x	-	yes	Begin table; if x is H table has repeated heading.
.ac A N	-	no	Set up for ACM style output. A is the Author's name(s), N is the total number of pages. Must be given before the first initialization.
.b x	no	no	Print x in boldface; if no argument switch to boldface.
.ba $+n$	0	yes	Augments the base indent by n . This indent is used to set the indent on regular text (like paragraphs).
.bc	no	yes	Begin new column.
.bi x	no	no	Print x in bold italics (nofill only).
.bu	-	yes	Begin bulleted paragraph.
.bx x	no	no	Print x in a box (nofill only).
.ef 'x'y'z	""	no	Set even footer to $x y z$.
.eh 'x'y'z	""	no	Set even header to $x y z$.
.fo 'x'y'z	""	no	Set footer to $x y z$.
.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 <i>x y z</i> .
.hl	-	yes	Draw a horizontal line.
.i x	no	no	Italicize <i>x</i> ; if <i>x</i> missing, italic text follows.
.ip x y	no	yes	Start indented paragraph, with hanging tag <i>x</i> . Indentation is <i>y</i> ens (default 5).
.lp	yes	yes	Start left-blocked paragraph.
.lo	-	no	Read in a file of local macros of the form . * <i>x</i> . Must be given before initialization.
.np	1	yes	Start numbered paragraph.
.of 'x'y'z	""	no	Set odd footer to <i>x y z</i> .
.oh 'x'y'z	""	no	Set odd header to <i>x y z</i> .
.pd	-	yes	Print delayed text.
.pp	no	yes	Begin paragraph. First line indented.
.r	yes	no	Roman text follows.
.re	-	no	Reset tabs to default values.
.sc	no	no	Read in a file of special characters and diacritical marks. Must be given before initialization.
.sh <i>n x</i>	-	yes	Section head follows, font automatically bold. <i>n</i> is level of section, <i>x</i> is title of section.
.sk	no	no	Leave the next page blank. Only one page is remembered ahead.

<i>Request</i>	<i>Initial</i>	<i>Cause</i>	<i>Explanation</i>
	<i>Value</i>	<i>Break</i>	
<code>.sm <i>x</i></code>	-	no	Set <i>x</i> in a smaller pointsize.
<code>.sz +<i>n</i></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 <i>x</i></code>	-	no	Underline argument (even in <code>troff</code>). (Nofill only).
<code>.uh</code>	-	yes	Like <code>.sh</code> but unnumbered.
<code>.xp <i>x</i></code>	-	no	Print index <i>x</i> .

文件

`/usr/share/lib/tmac/e`
`/usr/share/lib/tmac/*.me`

另请参见

[col\(1\)](#), [eqn\(1\)](#), [nroff\(1\)](#), [refer\(1\)](#), [tbl\(1\)](#), [troff\(1\)](#)

引用名	mech_spnego – Simple and Protected GSS-API Negotiation Mechanism						
用法概要	/usr/lib/gss/mech_spnego.so.1						
描述	<p>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. <code>mech_spnego.so.1</code> 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)).</p> <p>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.</p> <p>When SPNEGO is used, it selects the list of mechanisms to advertise by reading the GSS mechanism configuration file, <code>/etc/gss/mech</code> (see mech(4)), and by listing all active mechanisms except for itself.</p>						
选项	<p>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 <code>/etc/gss/mech</code>.</p> <p>This is an example (from <code>/etc/gss/mech</code>):</p> <pre>spnego 1.3.6.1.5.5.2 mech_spnego.so [msinterop]</pre> <p>Without the "[msinterop]" option, <code>mech_spnego</code> 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.</p>						
Interfaces	<code>mech_spnego.so.1</code> has no public interfaces. It is only activated and used through the GSS-API interface provided by <code>libgss.so.1</code> (see libgss(3LIB)).						
文件	<table><tr><td><code>/usr/lib/gss/mech_spnego.so.1</code></td><td>shared object file</td></tr><tr><td><code>/usr/lib/sparcv9/gss/mech_spnego.so.1</code></td><td>SPARC 64-bit shared object file</td></tr><tr><td><code>/usr/lib/amd64/gss/mech_spnego.so.1</code></td><td>x86 64-bit shared object file</td></tr></table>	<code>/usr/lib/gss/mech_spnego.so.1</code>	shared object file	<code>/usr/lib/sparcv9/gss/mech_spnego.so.1</code>	SPARC 64-bit shared object file	<code>/usr/lib/amd64/gss/mech_spnego.so.1</code>	x86 64-bit shared object file
<code>/usr/lib/gss/mech_spnego.so.1</code>	shared object file						
<code>/usr/lib/sparcv9/gss/mech_spnego.so.1</code>	SPARC 64-bit shared object file						
<code>/usr/lib/amd64/gss/mech_spnego.so.1</code>	x86 64-bit shared object file						
属性	See attributes(5) for descriptions of the following attributes:						

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUWNspnego
MT Level	Safe

另请参见

[Intro\(3\)](#), [libgss\(3LIB\)](#), [mech\(4\)](#), [attributes\(5\)](#)

《Oracle Solaris 11 开发者安全性指南》

引用名 mm – text formatting (memorandum) macros

用法概要 nroff -mm [options] filename...

troff -mm [options] filename...

描述 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 [l]	–	y,y	two column format l=line length
.AE	–	y	end abstract
.AL [t][i][s]	t=1;i=.Li;s=0	y	Start automatic list type t=[1,A,a,I,i] 1=arabic numbers; A=uppercase letters a=lowercase letters; I=uppercase Roman numerals; i=lowercase Roman numerals indentation i; separation s

Macro Name	Initial Value	Break? Reset?	Explanation
.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>]	<i>f</i> =10; <i>r</i> =1	n	set footnote style format <i>f</i> =[0-11]; renumber <i>r</i> =[0,1]
.FE	–	y	end footnote
.FG <i>x</i> [<i>n</i>]	<i>n</i> =1	y	figure title; figure <i>x</i> ; number <i>n</i>
.FS	–	n	start footnote
.H <i>l</i> [<i>t</i>]	–	y	produce numbered heading level <i>l</i> =[1-7]; title <i>t</i>
.HU <i>t</i>	–	y	produce unnumbered heading; title <i>t</i>
.I <i>x</i>	–	n	underline <i>x</i>
.IB <i>x y</i>	–	n	underline <i>x</i> and embolden <i>y</i>
.IR <i>x y</i>	–	n	underline <i>x</i> and use Roman font on <i>y</i>
.LE [<i>s</i>]	<i>s</i> =0	y	end list; separation <i>s</i>
.LI [<i>m</i>] [<i>p</i>]	–	y	start new list item; mark <i>m</i>
			prefix <i>p</i> (mark only)
.ML <i>m</i> [<i>i</i>] [<i>s</i>]	<i>s</i> =0	y	start marked list; mark <i>m</i> indentation <i>i</i> ; separation <i>s</i> =[0,1]
.MT <i>x</i>		y	memo title; title <i>x</i>
.ND <i>x</i>		n	no date in page footer; <i>x</i> is date on cover
.NE	–	y	end block text
.NS	–	y	start block text
.OF <i>x</i>	–	n	odd footer appears at the bottom of odd-numbered pages; <i>x</i> ="l' c' r" <i>l</i> =left; <i>c</i> =center; <i>r</i> =right
.OF <i>x</i>	–	n	odd header appears at the top of odd-numbered pages; <i>x</i> ="l' c' r" <i>l</i> =left; <i>c</i> =center; <i>r</i> =right
.OP	–	y	skip to the top of an odd-number page
.P [<i>t</i>]	<i>t</i> =0	y,y	begin paragraph; <i>t</i> =[0,1] 0=justified; 1=indented
.PF <i>x</i>	–	n	page footer appears at the bottom of every page; <i>x</i> ="l' c' r" <i>l</i> =left; <i>c</i> =center; <i>r</i> =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.

文件

/usr/share/lib/tmac/m
/usr/share/lib/tmac/mm.[nt] nroff and troff definitions of mm.

属性

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	text/doctools

另请参见

[col\(1\)](#), [eqn\(1\)](#), [nroff\(1\)](#), [refer\(1\)](#), [tbl\(1\)](#), [troff\(1\)](#), [attributes\(5\)](#)

已知问题

Floating keeps and regular keeps are diverted to the same space, so they cannot be mixed together with predictable results.

引用名 ms – text formatting macros

用法概要 nroff -ms [*options*] *filename*...

 troff -ms [*options*] *filename*...

描述 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 Whippy 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
*Q	quote (" in nroff, " in troff)
*U	unquote (" in nroff, " in troff)
*-	dash (- in nroff, - in troff)
*(MO	month (month of the year)
*(DY	day (current date)
**	automatically numbered footnote
*'	acute accent (before letter)
*^	grave accent (before letter)
*^	circumflex (before letter)
*,	cedilla (before letter)
*:	umlaut (before letter)
*~	tilde (before letter)

When using the extended accent mark definitions available with .AM, these strings should come after, rather than before, the letter to be accented.

文件

```
/usr/share/lib/tmac/s
/usr/share/lib/tmac/ms.???
```

另请参见

```
col(1),eqn(1),nroff(1),refer(1),tbl(1),troff(1)
```

已知问题

Floating keeps and regular keeps are diverted to the same space, so they cannot be mixed together with predictable results.

引用名

mutex – concepts relating to mutual exclusion locks

描述

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
mutex_init	Initialize a mutex.
mutex_destroy	Destroy a mutex.
mutex_lock	Lock a mutex.
mutex_trylock	Attempt to lock a mutex.
mutex_unlock	Unlock a mutex.
pthread_mutex_init	Initialize a mutex.
pthread_mutex_destroy	Destroy a mutex.
pthread_mutex_lock	Lock a mutex.
pthread_mutex_trylock	Attempt to lock a mutex.
pthread_mutex_unlock	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.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

另请参见 [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\(5\)](#)

附注 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).

引用名

mwac, MWAC – Mandatory Write Access Control

描述

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.

另请参见

[ln\(1\)](#), [zoneadm\(1M\)](#), [zonecfg\(1M\)](#), [link\(2\)](#), [pathconf\(2\)](#)

《Oracle Solaris 管理：安全服务》

引用名

nfssec – overview of NFS security modes

描述

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 RFC 2695: Authentication Mechanisms for ONC RPC).
<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.

<code>none</code>	Use null authentication (<code>AUTH_NONE</code>). NFS clients using <code>AUTH_NONE</code> have no identity and are mapped to the anonymous user <code>nobody</code> by NFS 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 <code>AUTH_NONE</code> . In this case, if the file system is shared with <code>sec=none</code> , users from the client are mapped to the anonymous user. The NFS security mode <code>none</code> is supported by share_nfs(1M) .
<code>sec=mode[:mode]. . .</code>	<p>Sharing uses one or more of the specified security modes. The <i>mode</i> in the <code>sec=mode</code> option must be a node name supported on the client. If the <code>sec=</code> option is not specified, the default security mode used is <code>AUTH_SYS</code>. Multiple <code>sec=</code> options can be specified on the command line, although each mode can appear only once.</p> <p>Each <code>sec=</code> option specifies modes that apply to any subsequent <code>window=</code>, <code>rw</code>, <code>ro</code>, <code>rw=</code>, <code>ro=</code> and <code>root=</code> options that are provided before another <code>sec=</code> option. Each additional <code>sec=</code> resets the security mode context, so that more <code>window=</code>, <code>rw</code>, <code>ro</code>, <code>rw=</code>, <code>ro=</code> and <code>root=</code> 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.

示例

示例 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
```

示例 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
```

示例 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
```

示例 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
```

文件

/etc/nfssec.conf NFS security service configuration file

属性

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

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

另请参见

[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

附注

/etc/nfssec.conf lists the NFS security services. Do not edit this file. It is not intended to be user-configurable. See [kclient\(1M\)](#).

引用名

nis, NIS, NIS+ – network information name service

描述

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.

引用名	nwam – 网络自动配置管理
描述	<p>可以根据系统定义的缺省策略规则或用户定义的策略规则自动管理网络配置。此管理由 <code>nwamd</code> 守护进程执行。</p> <p>NWAM 工具由 <code>network/physical:default</code> SMF 服务提供。如果 <code>netcfg/active_ncp</code> 属性设置为 NWAM NCP 名称，则会启用 NWAM 配置管理。或者，如果 <code>netcfg/active_ncp</code> 属性设置为 <code>DefaultFixed</code>，则会执行传统的网络配置。</p>
操作	<p>在自动 NCP 中实施的缺省配置策略用于让所有以物理方式连接的以太网链路处于活动状态；如果未提供以太网链路，则使单个无线链路处于活动状态。DHCP 用于获取所有活动链路的 IP 地址。可以通过创建备用网络配置文件 (Network Configuration Profiles, NCP) 更改该缺省策略。<code>nwam-manager</code> GUI 工具或 <code>nwamd(1M)</code> 命令可用于创建和修改 NCP。</p> <p>如果使用的是早期版本的 NWAM 服务（这会创建 <code>/etc/nwam/llp</code> 配置文件），则在升级到当前版本的 NWAM 服务后会自动将该文件中的配置包含到用户 NCP 中。</p> <p>由 NWAM 工具提供的自动行为要求通过 NWAM 服务传送网络配置管理。当条件发生更改以致激活了新配置文件或者当 NWAM 服务重新启动时，对网络配置所做的任何手动更改都可能会丢失。如果需要进行持久性更改，则建议方法是创建外部网络修饰符 (ENM)，可根据需要通过 NWAM 服务将其激活和取消激活。</p>
服务属性	<p>以下列表的格式为：</p> <pre>property_group/property_name property_type default_value</pre> <p>说明</p> <p><code>nwamd/debug</code> boolean false 使用 <code>daemon.debug</code> 启用调试日志。</p> <p><code>nwamd/autoconf</code> boolean false 指示在没有更好（更优）选择的情况下是否应自动连接打开的 WLAN。</p> <p><code>nwamd/ncu_wait_time</code> count 60 在尝试下一个可用 NCU 之前等待 NCU（或链路/接口 NCU 对，视需要而定）启动的秒数。初启活动不会被取消，并且最终可能会成功，此时会根据指定的配置条件激活更优的 NCU 并可能会禁用备用 NCU。</p> <p><code>nwamd/condition_check_interval</code> count 120 对有条件激活对象执行定期条件检查的间隔秒数。最小值为 30 秒。</p> <p><code>nwamd/scan_interval</code> count 120 执行定期无线扫描的间隔秒数。</p> <p><code>nwamd/scan_level</code> astring weak 一个信号强度阈值；如果当前连接的 AP 低于此信号级别并以更高的信号强度提供了对等 AP（属于同一 ESSID），则将断开现有连接以便于连接到信号更强的 AP。</p>

`nwamd/strict_bssid` boolean `false`
如果为 `true`，则必须同时匹配 ESSID 和 BSSID 才能连接到以前连接的 WLAN。如果为 `false`，则只需匹配 ESSID。

`netcfg/active_ncp` string `Automatic`
当前活动的 NCP。此属性不应由用户设置；它由 NWAM 服务在内部使用以在多次重新启动之间保持持久性。更改当前活动 NCP 的正确用户界面是使用 `nwam-manager` GUI 或 `netadm` 命令的 `enable` 子命令。

如果启用保留的 NCP 名称 `DefaultFixed`，将禁用 NWAM 自动配置管理，并启用传统的固定网络配置。

属性 有关下列属性的描述，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	<code>system/core-os</code>
接口稳定性	<code>Volatile</code> （可变）

另请参见 [svcs\(1\)](#)、[netcfgd\(1M\)](#)、[nwamd\(1M\)](#)、[svadm\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

另请参见 JDS/GNOME 手册页集合中的 `nwam-manager(1M)`。

附注 联网服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

`svc:/network/physical:default`

可以使用 [svadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

引用名	openssl – OpenSSL cryptographic and Secure Sockets Layer toolkit
描述	<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 cryptoadm(1M) 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, and CKM_SHA512.</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 fork(2) 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> <p>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.</p>

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 `mllocked-memory` to keep the PIN in a locked page using `mllock(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 --license library/security/openssl`.

示例

示例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
```

属性

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	library/security/openssl, library/security/openssl
Interface Stability	Volatile

另请参见

[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`

引用名 pam_allow – PAM authentication, account, session and password management PAM module to allow operations

用法概要 pam_allow.so.1

描述 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.

错误 PAM_SUCCESS is always returned.

示例 示例 1 Allowing ssh none
The following example is a pam.conf fragment that illustrates a sample for allowing ssh none authentication:

```
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
```

示例 2 Allowing Kiosk Automatic Login Service
The following is example is a pam.conf fragment that illustrates a sample for allowing gdm kiosk auto login:

```
gdm-autologin auth required pam_unix_cred.so.1
gdm-autologin auth sufficient pam_allow.so.1
```

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

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

另请参见 [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam_sm\(3PAM\)](#), [syslog\(3C\)](#), [pam.conf\(4\)](#), [attributes\(5\)](#), [pam_deny\(5\)](#), [pam_unix_cred\(5\)](#)

附注 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).

引用名	pam_authtok_check – authentication and password management module												
用法概要	pam_authtok_check.so.1												
描述	<p>pam_authtok_check provides functionality to the Password Management stack. The implementation of <code>pam_sm_chauthtok()</code> performs a number of checks on the construction of the newly entered password. <code>pam_sm_chauthtok()</code> is invoked twice by the PAM framework, once with flags set to <code>PAM_PRELIM_CHECK</code>, and once with flags set to <code>PAM_UPDATE_AUTHTOK</code>. This module only performs its checks during the first invocation. This module expects the current authentication token in the <code>PAM_OLDAUTHTOK</code> item, the new (to be checked) password in the <code>PAM_AUTHTOK</code> item, and the login name in the <code>PAM_USER</code> item. The checks performed by this module are:</p> <table><tr><td>length</td><td>The password length should not be less than the minimum specified in <code>/etc/default/passwd</code>.</td></tr><tr><td>circular shift</td><td>The password should not be a circular shift of the login name. This check may be disabled in <code>/etc/default/passwd</code>.</td></tr><tr><td>complexity</td><td>The password should contain at least the minimum number of characters described by the parameters <code>MINALPHA</code>, <code>MINNONALPHA</code>, <code>MINDIGIT</code>, and <code>MINSPECIAL</code>. Note that <code>MINNONALPHA</code> describes the same character classes as <code>MINDIGIT</code> and <code>MINSPECIAL</code> combined; therefore the user cannot specify both <code>MINNONALPHA</code> and <code>MINSPECIAL</code> (or <code>MINDIGIT</code>). The user must choose which of the two options to use. Furthermore, the <code>WHITESPACE</code> parameter determines whether whitespace characters are allowed. If unspecified <code>MINALPHA</code> is 2, <code>MINNONALPHA</code> is 1 and <code>WHITESPACE</code> is yes</td></tr><tr><td>variation</td><td>The old and new passwords must differ by at least the <code>MINDIFF</code> value specified in <code>/etc/default/passwd</code>. 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.</td></tr><tr><td>dictionary check</td><td>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 <code>DICTIONLIST</code>. 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 <code>DICTIONDBDIR</code> (defaults to <code>/var/passwd</code>). See mkpwdict(1M) for information on pre-generating the database. If neither <code>DICTIONLIST</code> nor <code>DICTIONDBDIR</code> is specified, no dictionary check is made.</td></tr><tr><td>upper/lower case</td><td>The password must contain at least the minimum of upper- and lower-case letters specified by the <code>MINUPPER</code> and <code>MINLOWER</code> values in <code>/etc/default/passwd</code>. If unspecified, the defaults are 0.</td></tr></table>	length	The password length should not be less than the minimum specified in <code>/etc/default/passwd</code> .	circular shift	The password should not be a circular shift of the login name. This check may be disabled in <code>/etc/default/passwd</code> .	complexity	The password should contain at least the minimum number of characters described by the parameters <code>MINALPHA</code> , <code>MINNONALPHA</code> , <code>MINDIGIT</code> , and <code>MINSPECIAL</code> . Note that <code>MINNONALPHA</code> describes the same character classes as <code>MINDIGIT</code> and <code>MINSPECIAL</code> combined; therefore the user cannot specify both <code>MINNONALPHA</code> and <code>MINSPECIAL</code> (or <code>MINDIGIT</code>). The user must choose which of the two options to use. Furthermore, the <code>WHITESPACE</code> parameter determines whether whitespace characters are allowed. If unspecified <code>MINALPHA</code> is 2, <code>MINNONALPHA</code> is 1 and <code>WHITESPACE</code> is yes	variation	The old and new passwords must differ by at least the <code>MINDIFF</code> value specified in <code>/etc/default/passwd</code> . 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 <code>DICTIONLIST</code> . 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 <code>DICTIONDBDIR</code> (defaults to <code>/var/passwd</code>). See mkpwdict(1M) for information on pre-generating the database. If neither <code>DICTIONLIST</code> nor <code>DICTIONDBDIR</code> 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 <code>MINUPPER</code> and <code>MINLOWER</code> values in <code>/etc/default/passwd</code> . If unspecified, the defaults are 0.
length	The password length should not be less than the minimum specified in <code>/etc/default/passwd</code> .												
circular shift	The password should not be a circular shift of the login name. This check may be disabled in <code>/etc/default/passwd</code> .												
complexity	The password should contain at least the minimum number of characters described by the parameters <code>MINALPHA</code> , <code>MINNONALPHA</code> , <code>MINDIGIT</code> , and <code>MINSPECIAL</code> . Note that <code>MINNONALPHA</code> describes the same character classes as <code>MINDIGIT</code> and <code>MINSPECIAL</code> combined; therefore the user cannot specify both <code>MINNONALPHA</code> and <code>MINSPECIAL</code> (or <code>MINDIGIT</code>). The user must choose which of the two options to use. Furthermore, the <code>WHITESPACE</code> parameter determines whether whitespace characters are allowed. If unspecified <code>MINALPHA</code> is 2, <code>MINNONALPHA</code> is 1 and <code>WHITESPACE</code> is yes												
variation	The old and new passwords must differ by at least the <code>MINDIFF</code> value specified in <code>/etc/default/passwd</code> . 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 <code>DICTIONLIST</code> . 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 <code>DICTIONDBDIR</code> (defaults to <code>/var/passwd</code>). See mkpwdict(1M) for information on pre-generating the database. If neither <code>DICTIONLIST</code> nor <code>DICTIONDBDIR</code> 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 <code>MINUPPER</code> and <code>MINLOWER</code> values in <code>/etc/default/passwd</code> . 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_AUTHTOK_CHECK flag set, force_check ignores this flag. The PAM_NO_AUTHTOK_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

返回值 If the password in PAM_AUTHTOK passes all tests, PAM_SUCCESS is returned. If any of the tests fail, PAM_AUTHTOK_ERR is returned.

文件 /etc/default/passwd See [passwd\(1\)](#) for a description of the contents.

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

ATTRIBUTETYPE	ATTRIBUTEVALUE
Interface Stability	Committed
MT Level	MT-Safe with exceptions

另请参见 [passwd\(1\)](#), [pam\(3PAM\)](#), [mkpasswd\(1M\)](#), [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\(5\)](#)

附注 The interfaces in [libpam\(3LIB\)](#) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

引用名	pam_authtok_get – authentication and password management module								
用法概要	pam_authtok_get.so.1								
描述	<p>The <code>pam_authtok_get</code> service module provides password prompting functionality to the PAM stack. It implements <code>pam_sm_authenticate()</code> and <code>pam_sm_chauthtok()</code>, providing functionality to both the Authentication Stack and the Password Management Stack.</p>								
Authentication Service	<p>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.</p>								
Password Management Service	<p>Due to the nature of the PAM Password Management stack traversal mechanism, the pam_sm_chauthtok(3PAM) function is called twice. Once with the PAM_PRELIM_CHECK flag, and one with the PAM_UPDATE_AUTHTOK flag.</p> <p>In the first (PRELIM) invocation, the implementation of pam_sm_chauthtok(3PAM) moves the contents of the PAM_AUTHTOK (current authentication token) to PAM_OLDAUTHTOK, and subsequently prompts the user for a new password. This new password is stored in PAM_AUTHTOK.</p> <p>If a previous module has set PAM_OLDAUTHTOK prior to the invocation of <code>pam_authtok_get</code>, this module turns into a NO-OP and immediately returns PAM_SUCCESS.</p> <p>In the second (UPDATE) invocation, the user is prompted to Re-enter his password. The <code>pam_sm_chauthtok</code> implementation verifies this reentered password with the password stored in PAM_AUTHTOK. If the passwords match, the module returns PAM_SUCCESS.</p> <p>The following option can be passed to the module:</p> <p><code>debug</code> syslog(3C) debugging information at the LOG_DEBUG level</p>								
错误	<p>The authentication service returns the following error codes:</p> <table><tr><td>PAM_SUCCESS</td><td>Successfully obtains authentication token</td></tr><tr><td>PAM_SYSTEM_ERR</td><td>Fails to retrieve username, username is NULL or empty</td></tr></table> <p>The password management service returns the following error codes:</p> <table><tr><td>PAM_SUCCESS</td><td>Successfully obtains authentication token</td></tr><tr><td>PAM_AUTHTOK_ERR</td><td>Authentication token manipulation error</td></tr></table>	PAM_SUCCESS	Successfully obtains authentication token	PAM_SYSTEM_ERR	Fails to retrieve username, username is NULL or empty	PAM_SUCCESS	Successfully obtains authentication token	PAM_AUTHTOK_ERR	Authentication token manipulation error
PAM_SUCCESS	Successfully obtains authentication token								
PAM_SYSTEM_ERR	Fails to retrieve username, username is NULL or empty								
PAM_SUCCESS	Successfully obtains authentication token								
PAM_AUTHTOK_ERR	Authentication token manipulation error								
属性	See attributes(5) for descriptions of the following attributes:								

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

另请参见

[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\(5\)](#)

附注

The interfaces in [libpam\(3LIB\)](#) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

引用名

pam_authtok_store – password management module

用法概要

pam_authtok_store.so.1

描述

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 <code>LOG_DEBUG</code> level
server_policy	If the account authority for the user, as specified by <code>PAM_USER</code> , is a server, do not encrypt the authentication token before updating.

错误

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.

属性

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

ATTRIBUTETYPE	ATTRIBUTEVALUE
Interface Stability	Committed
MT Level	MT-Safe with exceptions

另请参见

[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\(5\)](#)

附注

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.

引用名 pam_deny – PAM authentication, account, session and password management PAM module to deny operations

用法概要 pam_deny.so.1

描述 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

错误 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.

示例 示例 1 Disallowing ssh none authentication

```
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
```

示例 2 Disallowing any service not explicitly defined

```
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
```

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

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

另请参见 [su\(1M\)](#), [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam_sm_authenticate\(3PAM\)](#), [syslog\(3C\)](#), [pam.conf\(4\)](#), [nsswitch.conf\(4\)](#), [attributes\(5\)](#), [pam_authtok_check\(5\)](#), [pam_authtok_get\(5\)](#), [pam_authtok_store\(5\)](#), [pam_dhkeys\(5\)](#), [pam_passwd_auth\(5\)](#), [pam_unix_account\(5\)](#), [pam_unix_auth\(5\)](#), [pam_unix_session\(5\)](#), [privileges\(5\)](#)

附注

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.

引用名	pam_dhkeys – authentication Diffie-Hellman keys management module														
用法概要	pam_dhkeys.so.1														
描述	<p>The pam_dhkeys.so.1 service module provides functionality to two PAM services: Secure RPC authentication and Secure RPC authentication token management.</p> <p>Secure RPC authentication differs from regular Unix authentication because ONC RPCs use Secure RPC as the underlying security mechanism.</p> <p>The following options may be passed to the module:</p> <p>debug syslog(3C) debugging information at LOG_DEBUG level</p> <p>nowarn Turn off warning messages</p>														
Authentication Services	<p>If the user has Diffie-Hellman keys, <code>pam_sm_authenticate()</code> 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 <code>pam_sm_setcred()</code> is called with PAM_ESTABLISH_CRED and the user's secure RPC credentials need to be established, these credentials are set. This is equivalent to running keylogin(1).</p> <p>If the credentials could not be set and PAM_SILENT is not specified, a diagnostic message is displayed. If <code>pam_setcred()</code> is called with PAM_DELETE_CRED, the user's secure RPC credentials are unset. This is equivalent to running keylogout(1).</p> <p>PAM_REINITIALIZE_CRED and PAM_REFRESH_CRED are not supported and return PAM_IGNORE.</p>														
Authentication Token Management	<p>The <code>pam_sm_chauthtok()</code> implementation checks whether the old login password decrypts the users secret keys. If it doesn't this module prompts the user for an old Secure RPC password and stores it in a pam data item called SUNW_OLDRPCPASS. This data item can be used by the store module to effectively update the users secret keys.</p>														
错误	<p>The authentication service returns the following error codes:</p> <table><tr><td>PAM_SUCCESS</td><td>Credentials set successfully.</td></tr><tr><td>PAM_IGNORE</td><td>Credentials not needed to access the password repository.</td></tr><tr><td>PAM_USER_UNKNOWN</td><td>PAM_USER is not set, or the user is unknown.</td></tr><tr><td>PAM_AUTH_ERR</td><td>No secret keys were set. PAM_AUTHTOK is not set, no credentials are present or there is a wrong password.</td></tr><tr><td>PAM_BUF_ERR</td><td>Module ran out of memory.</td></tr></table> <p>The authentication token management returns the following error codes:</p> <table><tr><td>PAM_SUCCESS</td><td>Old rpc password is set in SUNW_OLDRPCPASS</td></tr><tr><td>PAM_USER_UNKNOWN</td><td>User in PAM_USER is unknown.</td></tr></table>	PAM_SUCCESS	Credentials set successfully.	PAM_IGNORE	Credentials not needed to access the password repository.	PAM_USER_UNKNOWN	PAM_USER is not set, or the user is unknown.	PAM_AUTH_ERR	No secret keys were set. PAM_AUTHTOK is not set, no credentials are present or there is a wrong password.	PAM_BUF_ERR	Module ran out of memory.	PAM_SUCCESS	Old rpc password is set in SUNW_OLDRPCPASS	PAM_USER_UNKNOWN	User in PAM_USER is unknown.
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.														
PAM_SUCCESS	Old rpc password is set in SUNW_OLDRPCPASS														
PAM_USER_UNKNOWN	User in PAM_USER is unknown.														

PAM_BUF_ERR	Module ran out of memory.
-------------	---------------------------

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

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

```
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_authtok_check(5),pam_authtok_get(5),pam_authtok_store(5),
pam_passwd_auth(5),pam_unix_account(5),pam_unix_auth(5),pam_unix_session(5)
```

The interfaces in `libpam(3LIB)` are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

引用名	pam_dial_auth – authentication management PAM module for dialups								
用法概要	pam_dial_auth.so.1								
描述	<p>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.</p> <p>Authentication service modules must implement both <code>pam_sm_authenticate()</code> and <code>pam_sm_setcred()</code>. <code>pam_sm_setcred()</code> in this module always returns <code>PAM_IGNORE</code>.</p> <p>The value of the <code>PAM_TTY</code> 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.</p> <p>The following option may be passed in to this service module:</p> <p><code>debug</code> syslog(3C) debugging information at <code>LOG_DEBUG</code> level.</p>								
错误	<p>If dialups(4) is not present, <code>PAM_IGNORE</code> is returned. Upon successful completion of <code>pam_sm_authenticate()</code>, <code>PAM_SUCCESS</code> is returned. The following error codes are returned upon error:</p> <table><tr><td><code>PAM_AUTH_ERR</code></td><td>Authentication failure.</td></tr><tr><td><code>PAM_SERVICE_ERR</code></td><td>Error in the calling service, <code>PAM_TTY</code> is not set.</td></tr><tr><td><code>PAM_SYSTEM_ERR</code></td><td>System error (d_passwd(4) is not present).</td></tr><tr><td><code>PAM_USER_UNKNOWN</code></td><td>No account is present for <i>user</i>.</td></tr></table>	<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 (d_passwd(4) is not present).	<code>PAM_USER_UNKNOWN</code>	No account is present for <i>user</i> .
<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 (d_passwd(4) is not present).								
<code>PAM_USER_UNKNOWN</code>	No account is present for <i>user</i> .								
属性	See attributes(5) for a description of the following attributes:								
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>MT Level</td><td>MT-Safe with exceptions</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT Level	MT-Safe with exceptions				
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
MT Level	MT-Safe with exceptions								
另请参见	pam(3PAM) , pam_authenticate(3PAM) , pam_sm_authenticate(3PAM) , d_passwd(4) , dialups(4) , libpam(3LIB) , pam.conf(4) , attributes(5) , pam_authtok_check(5) , pam_authtok_get(5) , pam_authtok_store(5) , pam_dhkeys(5) , pam_passwd_auth(5) , pam_unix_account(5) , pam_unix_auth(5) , pam_unix_session(5)								
附注	The interfaces in libpam(3LIB) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.								

引用名 pam_krb5 – authentication, account, session, and password management PAM modules for Kerberos V5

用法概要 /usr/lib/security/pam_krb5.so.1

描述 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.

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_AUTHTOK`

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_AUTHTOK`

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.

错误

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.
<code>PAM_IGNORE</code>	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.

示例

示例 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`.

<code>gdm auth requisite</code>	<code>pam_authtok_get.so.1</code>
<code>gdm auth required</code>	<code>pam_dhkeys.so.1</code>
<code>gdm auth required</code>	<code>pam_unix_cred.so.1</code>
<code>gdm auth sufficient</code>	<code>pam_krb5.so.1</code>
<code>gdm auth required</code>	<code>pam_unix_auth.so.1</code>

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.

示例1 Authenticating Users Through Kerberos as First Choice Using Password-based Authentication (续)

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
```

When changing Kerberos based user's password, use `kpasswd(1)`. When changing a non-Kerberos user's password, it is recommended that the repository is specified (`-r`) with the `passwd(1)` command.

示例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 binding      pam_krb5.so.1
gdm auth required     pam_unix_auth.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;

示例 2 Authenticating Users Through Kerberos Only Using Password-based Authentication (续)

- 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
other password required      pam_authtok_store.so.1
```

示例 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 required       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.

示例 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`.

示例 4 Authenticating Users Through Kerberos PKINIT as First Choice (续)

```
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
```

示例 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
```

示例 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
```

示例 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 PKINT 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
```

示例 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
```

示例 8 Authenticating Users Through Kerberos PKINIT as a Requirement (续)

```
login auth required      pam_krb5.so.1 pkinit
login auth required      pam_unix_auth.so.1
```

示例 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
```

示例 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
```

示例 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
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

另请参见 [kdestroy\(1\)](#), [kinit\(1\)](#), [kpasswd\(1\)](#), [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\(5\)](#)

附注 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.

引用名	pam_krb5_migrate – authentication PAM module for the KerberosV5 auto-migration of users feature						
用法概要	/usr/lib/security/pam_krb5_migrate.so.1						
描述	<p>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.</p>						
KerberosV5 Auto-migrate Authentication Module	<p>The KerberosV5 auto-migrate authentication component provides the <code>pam_sm_authenticate(3PAM)</code> function to migrate a user who does not have a corresponding krb5 principal account to the default Kerberos realm of the client.</p> <p><code>pam_sm_authenticate(3PAM)</code> uses a host-based client service principal, present in the local keytab (/etc/krb5/krb5.keytab) to authenticate to <code>kadmind(1M)</code> (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 <code>kadm5.acl(4)</code> file. <code>kadmind(1M)</code> checks for the appropriate privilege and validates the user password using PAM by calling <code>pam_authenticate(3PAM)</code> and <code>pam_acct_mgmt(3PAM)</code> for the k5migrate service.</p> <p>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.</p> <p>The authentication component always returns PAM_IGNORE and is meant to be stacked in pam.conf 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 <code>pam_krb5(5)</code> be listed in the PAM account stack of that service for proper operation (see EXAMPLES).</p>						
选项	<p>The following options can be passed to the KerberosV5 auto-migrate authentication module:</p> <table><tr><td>debug</td><td>Provides <code>syslog(3C)</code> debugging information at LOG_DEBUG level.</td></tr><tr><td>client_service=<service name></td><td>Name of the service used to authenticate to <code>kadmind(1M)</code> defaults to host. This means that the module uses host/<nodename.fqdn> as its client service principal name, KerberosV5 user principal creation operation or <service>/<nodename.fqdn> if this option is provided.</td></tr><tr><td>quiet</td><td>Do not explain KerberosV5 migration to the user.</td></tr></table>	debug	Provides <code>syslog(3C)</code> debugging information at LOG_DEBUG level.	client_service=<service name>	Name of the service used to authenticate to <code>kadmind(1M)</code> defaults to host. This means that the module uses host/<nodename.fqdn> as its client service principal name, KerberosV5 user principal creation operation or <service>/<nodename.fqdn> if this option is provided.	quiet	Do not explain KerberosV5 migration to the user.
debug	Provides <code>syslog(3C)</code> debugging information at LOG_DEBUG level.						
client_service=<service name>	Name of the service used to authenticate to <code>kadmind(1M)</code> defaults to host. This means that the module uses host/<nodename.fqdn> as its client service principal name, KerberosV5 user principal creation operation or <service>/<nodename.fqdn> if this option is provided.						
quiet	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 KerberosV5 user principals with password expiration set to now (current time).

示例

示例 1 Sample Entries from `pam.conf`

The following entries from `pam.conf(4)` 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 `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.

示例 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.

示例 3 Sample Entries in pam.conf of the Master KDC

The 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
```

属性

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

另请参见

[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\(5\)](#)

引用名	pam_ldap – authentication and account management PAM module for LDAP
用法概要	/usr/lib/security/pam_ldap.so.1
描述	<p>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.</p> <p>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 示例 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.</p>
LDAP Authentication Module	<p>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.</p> <p>The authentication method used is either defined in the client profile , or the authentication method is configured by using the ldapclient(1M) command. To determine the authentication method to use, this module first attempts to use the authentication method that is defined, for service pam_ldap, for example, <code>serviceAuthenticationMethod:pam_ldap:sasl/DIGEST-MD5</code>. If no authentication method is defined, pam_ldap uses the default authentication method. If neither are set, the authentication fails. This module skips the configured authentication method if the authentication method is set to none.</p> <p>The following options can be passed to the LDAP service module:</p> <p><code>debug</code> syslog(3C) debugging information at LOG_DEBUG level.</p> <p><code>nowarn</code> Turn off warning messages.</p>

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)`.

错误

The authentication service returns the following error codes:

- `PAM_SUCCESS` The uthentication was successful.
- `PAM_MAXTRIES` The maximum number of authentication attempts was exceeded.
- `PAM_AUTH_ERR` The authentication failed.
- `PAM_USER_UNKNOWN` No account is present for the user.
- `PAM_BUF_ERR` A memory buffer error occurred.
- `PAM_SYSTEM_ERR` A system error occurred.
- `PAM_IGNORE` The user's account was inactivated.

The account management service returns the following error codes:

- `PAM_SUCCESS` The user was allowed access to the account.
- `PAM_NEW_AUTHTOK_REQD` A new authentication token is required.
- `PAM_ACCT_EXPIRED` The user account has expired.
- `PAM_PERM_DENIED` The user was denied access to the account at this time.
- `PAM_USER_UNKNOWN` No account is present for the user.
- `PAM_BUF_ERROR` A memory buffer error occurred.
- `PAM_SYSTEM_ERR` A system error occurred.

示例

示例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
```

示例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
```

示例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
```

示例 3 Using pam_authtok_store With Password Management For Both Local and LDAP Accounts (续)

```
# 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
```

文件

/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.

属性

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

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

另请参见

[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\(5\)](#)

附注

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`.

引用名

pam_list – PAM account management module for UNIX

用法概要

pam_list.so.1

描述

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:

<code>allow=</code>	The full pathname to a file of allowed users and/or netgroups. Only one of <code>allow=</code> or <code>deny=</code> can be specified.
<code>compat</code>	Activate <code>compat</code> mode.
<code>deny=</code>	The full pathname to a file of denied users and/or netgroups. Only one of <code>deny=</code> or <code>allow=</code> can be specified.
<code>debug</code>	Provide <code>syslog(3C)</code> debugging information at the <code>LOG_AUTH LOG_DEBUG</code> level.
<code>user</code>	The module should only perform netgroup matches on the username. This is the default option.
<code>nouser</code>	The username should not be used in the netgroup match.
<code>host</code>	Only the host should be used in netgroup matches.
<code>nohost</code>	The hostname should not be used in netgroup matches.
<code>norole</code>	Return <code>PAM_IGNORE</code> if the account (<code>PAM_USER</code>) is a role. This is the default.
<code>role</code>	Evaluate the rules even if <code>PAM_USER</code> is a role account.
<code>user_host_exact</code>	The user and hostname must be in the same netgroup.

错误

The following error values are returned:

PAM_SERVICE_ERR	An invalid set of module options was given in the pam.conf(4) 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

示例

示例1 Using pam_list in default mode

/etc/pam.conf modification looks like:

```
other    account requisite    pam_roles.so.1
other    account required     pam_unix_account.so.1
other    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
```

示例2 Using pam_list with allow file

/etc/pam.conf modification looks like:

```
other    account requisite    pam_roles.so.1
other    account required     pam_unix_account.so.1
other    account required     pam_list.so.1 allow=/etc/user.allow
```

/etc/users.allow contains:

```
root
localloginname
remoteloginname
@netgroup
```

属性

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.

另请参见

[pam\(3PAM\)](#), [pam_authenticate\(3PAM\)](#), [pam_sm_acct_mgmt\(3PAM\)](#), [syslog\(3C\)](#), [libpam\(3LIB\)](#), [nsswitch.conf\(4\)](#), [pam.conf\(4\)](#), [attributes\(5\)](#)

引用名

pam_passwd_auth – authentication module for password

用法概要

pam_passwd_auth.so.1

描述

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.

错误

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

属性

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

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

另请参见

[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\(5\)](#)

附注

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\(3PAM\)](#).

引用名	pam_pkcs11 – PAM Authentication Module for the PKCS#11 token libraries
用法概要	pam_pkcs11.so [debug] [config_file=filename]
描述	<p>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.</p> <p>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:</p> <ul style="list-style-type: none"> ■ 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. <p>For the verification of the users' certificates, locally stored CA certificates as well as either online or locally accessible CRLs are used.</p>
PAM CONFIGURATION	<p>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 the /etc/pam.conf file. For details on how to configure PAM services, see pam.conf(4).</p> <p>The following example uses only pam_pkcs11 for authentication:</p> <pre>login auth requisite pam_pkcs11.so.1 login autho required pam_unix_cred.so.1</pre> <p>The following example uses pam_pkcs11 for authentication with fallback to standard UNIX authentication:</p> <pre>login auth sufficient pam_pkcs11.so.1 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 required pam_unix_auth.so.1</pre>
PAM_PKCS11 CONFIGURATION	<p>To configure the pam_pkcs11 module, you must have the following information:</p> <ul style="list-style-type: none"> ■ 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 <p>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</p>

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.

选项

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.

文件

`/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

属性 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

另请参见 [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.

引用名	pam_rhosts_auth – authentication management PAM module using ruserok()				
用法概要	/usr/lib/security/pam_rhosts_auth.so.1				
描述	<p>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.</p> <p>/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 <i>sufficient</i> so that subsequent authentication is performed only on failure of pam_sm_authenticate(). The following option may be passed in to this service module:</p> <p>debug syslog(3C) debugging information at LOG_DEBUG level.</p>				
属性	See attributes(5) for descriptions of the following attributes:				
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>MT Level</td><td>MT-Safe with exceptions</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT Level	MT-Safe with exceptions				
另请参见	pam(3PAM) , pam_authenticate(3PAM) , ruserok(3SOCKET) , syslog(3C) , libpam(3LIB) , pam.conf(4) , attributes(5)				
附注	The interfaces in libpam() are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.				

引用名 pam_roles – Solaris Roles account management module

用法概要 pam_roles.so.1

描述 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 [syslog\(3C\)](#) debugging information at the LOG_DEBUG level.

错误 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.

示例 示例1 Using the pam_roles.so.1 Module

The following are sample entries from [pam.conf\(4\)](#). These entries demonstrate 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 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.

示例 2 Allowing Remote Roles

Remote roles should only be allowed from remote services that can be trusted to provide an accurate PAM_AUSERname. This trust is a function of the protocol (such as sshd-hostbased).

The following is a sample entry for a [pam.conf\(4\)](#) file. It 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
```

属性

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

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

另请参见

[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\(5\)](#)

附注

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\(3PAM\)](#).

引用名	pam_sample – a sample PAM module																		
用法概要	/usr/lib/security/pam_sample.so.1																		
描述	<p>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.</p>																		
Sample Authentication Component	<p>The SAMPLE authentication module provides functions to test the PAM framework functionality using the <code>pam_sm_authenticate(3PAM)</code> call. The SAMPLE module implementation of the <code>pam_sm_authenticate(3PAM)</code> function compares the user entered password with the password set in the <code>pam.conf(4)</code> file, 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:</p> <table><tr><td>debug</td><td>Syslog debugging information at the LOG_DEBUG level.</td></tr><tr><td>pass=newone</td><td>Sets the password to be newone.</td></tr><tr><td>first_pass_good</td><td>The first password is always good when used with the use_first_pass or try_first_pass option.</td></tr><tr><td>first_pass_bad</td><td>The first password is always bad when used with the use_first_pass or try_first_pass option.</td></tr><tr><td>always_fail</td><td>Always returns PAM_AUTH_ERR.</td></tr><tr><td>always_succeed</td><td>Always returns PAM_SUCCESS.</td></tr><tr><td>always_ignore</td><td>Always returns PAM_IGNORE.</td></tr><tr><td>use_first_pass</td><td>Use the user's initial password (entered when the user is authenticated to the first authentication module in the stack) to authenticate with the SAMPLE module. If the passwords do not match, or if this is the first authentication module in the stack, quit and do not prompt the user for a password. It is recommended that this option only be used if the SAMPLE authentication module is designated as <i>optional</i> in the <code>pam.conf</code> configuration file.</td></tr><tr><td>try_first_pass</td><td>Use the user's initial password (entered when the user is authenticated to the first authentication module in the stack) to authenticate with the SAMPLE module. If the passwords do not match, or if this is the first authentication module in the stack, prompt the user for a password.</td></tr></table> <p>The SAMPLE module <code>pam_sm_setcred(3PAM)</code> function always returns PAM_SUCCESS.</p>	debug	Syslog debugging information at the LOG_DEBUG level.	pass=newone	Sets the password to be newone.	first_pass_good	The first password is always good when used with the use_first_pass or try_first_pass option.	first_pass_bad	The first password is always bad when used with the use_first_pass or try_first_pass option.	always_fail	Always returns PAM_AUTH_ERR.	always_succeed	Always returns PAM_SUCCESS.	always_ignore	Always returns PAM_IGNORE.	use_first_pass	Use the user's initial password (entered when the user is authenticated to the first authentication module in the stack) to authenticate with the SAMPLE module. If the passwords do not match, or if this is the first authentication module in the stack, quit and do not prompt the user for a password. It is recommended that this option only be used if the SAMPLE authentication module is designated as <i>optional</i> in the <code>pam.conf</code> configuration file.	try_first_pass	Use the user's initial password (entered when the user is authenticated to the first authentication module in the stack) to authenticate with the SAMPLE module. If the passwords do not match, or if this is the first authentication module in the stack, prompt the user for a password.
debug	Syslog debugging information at the LOG_DEBUG level.																		
pass=newone	Sets the password to be newone.																		
first_pass_good	The first password is always good when used with the use_first_pass or try_first_pass option.																		
first_pass_bad	The first password is always bad when used with the use_first_pass or try_first_pass option.																		
always_fail	Always returns PAM_AUTH_ERR.																		
always_succeed	Always returns PAM_SUCCESS.																		
always_ignore	Always returns PAM_IGNORE.																		
use_first_pass	Use the user's initial password (entered when the user is authenticated to the first authentication module in the stack) to authenticate with the SAMPLE module. If the passwords do not match, or if this is the first authentication module in the stack, quit and do not prompt the user for a password. It is recommended that this option only be used if the SAMPLE authentication module is designated as <i>optional</i> in the <code>pam.conf</code> configuration file.																		
try_first_pass	Use the user's initial password (entered when the user is authenticated to the first authentication module in the stack) to authenticate with the SAMPLE module. If the passwords do not match, or if this is the first authentication module in the stack, prompt the user for a password.																		
Sample Account Management Component	<p>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</p>																		

as option arguments to the entry for the SAMPLE account management PAM module in the `pam.conf` file. Note that the module always permits access to the root super user.

The option field syntax to limit access is shown below: `allow= name[,name]` `allow= name [allow=name]`

The example `pam.conf` show below permits only larry to login directly. `rlogin` is allowed only for don and larry. Once a user is logged in, the user can use `su` if the user are sam or eric.

login	account	require	pam_sample.so.1 allow=larry
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
Sample Session
Management
Component
属性

The SAMPLE Password Management Component function (`pam_sm_chauthtok(3PAM)`), always returns `PAM_SUCCESS`.

The SAMPLE Session Management Component functions (`pam_sm_open_session(3PAM)`, `pam_sm_close_session(3PAM)`) always return `PAM_SUCCESS`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	MT-Safe with exceptions

另请参见

`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\(5\)](#)

警告

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

附注

The interfaces in `libpam()` are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

引用名	pam_smbfs_login – PAM user credential authentication module for SMB/CIFS client login
用法概要	pam_smb_cred.so.1
描述	<p>The <code>pam_smbfs_login</code> module implements <code>pam_sm_setcred(3PAM)</code> to provide functions that act equivalently to the <code>smbadm(1M)</code> <code>add-key</code> command.</p> <p>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.</p> <p>This module permits the login password to be stored as if the <code>smbadm(1M)</code> <code>add-key</code> command was used to store a password for <code>PAM_USER</code> in the user or system default domain.</p> <p>To use this functionality, add the following line to the <code>/etc/pam.conf</code> file:</p> <pre>login auth optional pam_smbfs_login.so.1</pre> <p>Authentication service modules must implement both <code>pam_sm_authenticate(3PAM)</code> and <code>pam_sm_setcred(3PAM)</code>. In this module, <code>pam_sm_authenticate(3PAM)</code> always returns <code>PAM_IGNORE</code>.</p> <p>The <code>pam_sm_setcred(3PAM)</code> function accepts the following flags:</p> <p><code>PAM_REFRESH_CRED</code> Returns <code>PAM_IGNORE</code>.</p> <p><code>PAM_SILENT</code> Suppresses messages.</p> <p><code>PAM_ESTABLISH_CRED</code> <code>PAM_REINITIALIZE_CRED</code> Stores the authentication token for <code>PAM_USER</code> in the same manner as the <code>smbadm(1M)</code> <code>add-key</code> command.</p> <p><code>PAM_DELETE_CRED</code> Deletes the stored password for <code>PAM_USER</code> in the same manner as the <code>smbadm(1M)</code> <code>remove-key</code> command.</p> <p>The following options can be passed to the <code>pam_smbfs_login</code> module:</p> <p><code>debug</code> Produces <code>syslog(3C)</code> debugging information at the <code>LOG_AUTH</code> or <code>LOG_DEBUG</code> level.</p> <p><code>nowarn</code> Suppresses warning messages.</p> <p>Upon successful completion of <code>pam_sm_setcred(3PAM)</code>, <code>PAM_SUCCESS</code> is returned. The following error codes are returned upon error:</p> <p><code>PAM_USER_UNKNOWN</code> User is unknown.</p>
错误	

PAM_AUTHTOK_ERR
Password is bad.

PAM_AUTH_ERR
Domain is bad.

PAM_SYSTEM_ERR
System error.

属性 See [attributes\(5\)](#) for descriptions of the following attribute:

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

另请参见 [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\(7FS\)](#)

附注 The interfaces in [libpam\(3LIB\)](#) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

引用名	pam_smb_passwd – SMB password management module
用法概要	pam_smb_passwd.so.1
描述	<p>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.</p> <p>The <code>pam_sm_chauthtok()</code> function accepts the following flags:</p> <p>PAM_PRELIM_CHECK Always returns <code>PAM_IGNORE</code>.</p> <p>PAM_SILENT Suppresses messages.</p> <p>PAM_UPDATE_AUTHTOK Updates or creates a new SMB local LM/NTLM hash for the user that is specified in <code>PAM_USER</code> by using the authentication information found in <code>PAM_AUTHTOK</code>. The LM hash is only created if the <code>smbd/lmauth_level</code> property value of the <code>smb/server</code> service is set to 3 or less. <code>PAM_IGNORE</code> is returned if the user is not in the local <code>/etc/passwd</code> repository.</p> <p>The following options can be passed to the <code>pam_smb_passwd</code> module:</p> <p>debug Produces syslog(3C) debugging information at the <code>LOG_AUTH</code> or <code>LOG_DEBUG</code> level.</p> <p>nowarn Suppresses warning messages.</p>
文件	<p><code>/var/smb/smbpasswd</code> Stores SMB passwords for Solaris users.</p>
错误	<p>Upon successful completion of <code>pam_sm_chauthtok()</code>, <code>PAM_SUCCESS</code> is returned. The following error codes are returned upon error:</p> <p>PAM_AUTHTOK_ERR Authentication token manipulation error</p> <p>PAM_AUTHTOK_LOCK_BUSY SMB password file is locked</p> <p>PAM_PERM_DENIED Permissions are insufficient for accessing the SMB password file</p> <p>PAM_SYSTEM_ERR System error</p> <p>PAM_USER_UNKNOWN User is unknown</p>

属性 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

另请参见 [smbd\(1M\)](#), [syslog\(3C\)](#), [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam_chauthtok\(3PAM\)](#), [pam_sm\(3PAM\)](#), [pam_sm_chauthtok\(3PAM\)](#), [pam.conf\(4\)](#), [attributes\(5\)](#)

附注 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.

引用名	pam_tsol_account – PAM account management module for Trusted Extensions
用法概要	/usr/lib/security/pam_tsol_account.so.1
描述	<p>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.</p> <p>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.</p> <p>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:</p> <pre>other account required pam_tsol_account allow_unlabeled</pre>
选项	<p>The following options can be passed to the module:</p> <p>allow_unlabeled Allows remote connections from hosts with unlabeled template types.</p> <p>debug Provides debugging information at the LOG_DEBUG level. See syslog(3C).</p>
返回值	<p>The following values are returned:</p> <p>PAM_SUCCESS The account is valid for use at this time and label.</p> <p>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.</p> <p>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.</p>
属性	See attributes(5) for description 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 multi-threaded application uses its own PAM handle.

- 另请参见 [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\(5\)](#)
- 《Oracle Solaris 管理：安全服务》中的第 15 章“使用 PAM”
- 附注 The functionality described on this manual page is available only if the system is configured with Trusted Extensions.

引用名

用法概要

描述

pam_unix_account – PAM account management module for UNIX

pam_unix_account.so.1

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.

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

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

属性

ATTRIBUTETYPE	ATTRIBUTE VALUE
Interface Stability	Committed
MT Level	MT-Safe with exceptions

另请参见

[pam\(3PAM\)](#), [pam_authenticate\(3PAM\)](#), [syslog\(3C\)](#), [libpam\(3LIB\)](#), [pam.conf\(4\)](#), [nsswitch.conf\(4\)](#), [attributes\(5\)](#)

附注

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.

引用名 pam_unix_auth – PAM authentication module for UNIX

用法概要 pam_unix_auth.so.1

描述 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 `RETRIES` as defined in `login(1)`, the account is locked and `PAM_MAXTRIES` is returned. Currently, only the “files” repository (see `passwd(4)` and `shadow(4)`) supports 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.

错误 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.

属性

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

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

另请参见

[login\(1\)](#), [passwd\(1\)](#), [useradd\(1M\)](#), [usermod\(1M\)](#), [roleadd\(1M\)](#), [rolemod\(1M\)](#), [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_authtok_check\(5\)](#), [pam_authtok_get\(5\)](#), [pam_authtok_store\(5\)](#), [pam_dhkeys\(5\)](#), [pam_passwd_auth\(5\)](#), [pam_unix_account\(5\)](#), [pam_unix_session\(5\)](#)

附注

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.

引用名

pam_unix_cred – PAM user credential authentication module for UNIX

用法概要

pam_unix_cred.so.1

描述

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.

错误 Upon successful completion of `pam_sm_setcred()`, `PAM_SUCCESS` is returned. The following error codes are returned upon error:

- `PAM_CRED_UNAVAIL` Underlying authentication service cannot retrieve user credentials
- `PAM_CRED_EXPIRED` User credentials have expired
- `PAM_USER_UNKNOWN` User is unknown to the authentication service
- `PAM_CRED_ERR` Failure in setting user credentials
- `PAM_BUF_ERR` Memory buffer error
- `PAM_SYSTEM_ERR` System error

The following values are returned from `pam_sm_authenticate()`:

- `PAM_IGNORE` Ignores this module regardless of the control flag

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

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

另请参见 [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\(5\)](#)

附注 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.

引用名

pam_unix_session – session management PAM module for UNIX

用法概要

pam_unix_session.so.1

描述

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

错误

Upon successful completion, `PAM_SUCCESS` is returned. The following error codes are returned upon error:

<code>PAM_SESSION_ERR</code>	Cannot make or remove the entry for the specified session (<code>PAM_TTY</code> is not present).
<code>PAM_USER_UNKNOWN</code>	No account is present for <i>user</i> .

属性

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

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

另请参见

[pam\(3PAM\)](#), [pam_authenticate\(3PAM\)](#), [syslog\(3C\)](#), [libpam\(3LIB\)](#), [pam.conf\(4\)](#), [nsswitch.conf\(4\)](#), [attributes\(5\)](#), [pam_authtok_check\(5\)](#), [pam_authtok_get\(5\)](#), [pam_authtok_store\(5\)](#), [pam_dhkeys\(5\)](#), [pam_passwd_auth\(5\)](#), [pam_unix_account\(5\)](#), [pam_unix_auth\(5\)](#),

附注

The interfaces in [libpam\(3LIB\)](#) are MT-Safe only if each thread within the multi-threaded application uses its own PAM handle.

引用名	pam_zfs_key – PAM user credential module for ZFS				
用法概要	pam_zfs_key.so.1 [create] [homes=]				
描述	<p>The pam_zfs_key module implements <code>pam_sm_setcred(3PAM)</code> and <code>pam_sm_chauthtok(3PAM)</code>.</p> <p>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 <code>pam_sm_authenticate()</code> and <code>pam_sm_setcred()</code>.</p> <p><code>pam_sm_authenticate()</code> in this module always returns <code>PAM_IGNORE</code>. If they are located at a different location then the module option, <code>homes=</code> can be used to specify that. It is the ZFS dataset name, not the mountpoint, which would usually be <code>/export/home/</code>.</p> <p>The last component of the ZFS dataset name must match the value of <code>PAM_USER</code>, that is, the users login name. If the users home directory is a local ZFS filesystem with encryption enabled and the ZFS <code>keysource</code> property is set to <code>passphrase, prompt</code>, on <code>pam_sm_setcred()</code> this module manages the keys as follows:</p> <table><tr><td><code>PAM_DELETE_CRED</code></td><td>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 <code>force</code> module option is provided to attempt a forced unmount first.</td></tr><tr><td><code>PAM_ESTABLISH_CRED</code></td><td>Attempts to use the value of <code>PAM_AUTHTOK</code> to load the key for the ZFS dataset and mount it.</td></tr></table> <p>If <code>PAM_AUTHTOK</code> is not the correct passphrase, the user is prompted once for an alternate. This value is never stored in <code>PAM_AUTHTOK</code>, even if it is the correct passphrase for the ZFS dataset that is the users home directory. This is equivalent to:</p> <pre>zfs key -l rpool/export/home/\$USER</pre> <p>If no ZFS file system exists for the user and the <code>create</code> module option is provided, a new one is created. The ZFS encryption property defaults to <code>on</code> in this case unless the <code>encryption=</code> property is set for the module to override it.</p> <p>If the <code>create</code> module option is not provided and no ZFS file system exists for the user, the module returns <code>PAM_IGNORE</code>.</p> <p>The newly created ZFS file system has the following ZFS delegations specified for the user for which it is created: <code>key, keychange, mount</code>. In these cases <code>keysource</code> is always set to <code>passphrase, prompt</code>.</p>	<code>PAM_DELETE_CRED</code>	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 <code>force</code> module option is provided to attempt a forced unmount first.	<code>PAM_ESTABLISH_CRED</code>	Attempts to use the value of <code>PAM_AUTHTOK</code> to load the key for the ZFS dataset and mount it.
<code>PAM_DELETE_CRED</code>	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 <code>force</code> module option is provided to attempt a forced unmount first.				
<code>PAM_ESTABLISH_CRED</code>	Attempts to use the value of <code>PAM_AUTHTOK</code> to load the key for the ZFS dataset and mount it.				

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.

示例

示例 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
```

示例 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
```

示例 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
```

属性

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.

另请参见

[zfs\(1M\)](#), [umount2\(2\)](#), [pam.conf\(4\)](#), [libpam\(3LIB\)](#), [pam\(3PAM\)](#), [pam_sm_chauthtok\(3PAM\)](#), [pam_sm_setcred\(3PAM\)](#), [attributes\(5\)](#)

引用名	pkcs11_kernel – 内核加密框架的 PKCS#11 接口
用法概要	<code>/usr/lib/security/pkcs11_kernel.so</code> <code>/usr/lib/security/64/pkcs11_kernel.so</code>
描述	<p><code>pkcs11_kernel.so</code> 对象使用专用接口与内核加密框架进行通信，从而实现 RSA Security Inc. PKCS#11 加密令牌接口 (Cryptoki) v2.20 规范。</p> <p>每个唯一的硬件提供者均由一个 PKCS#11 插槽表示。在没有硬件内核加密框架提供者的系统中，此 PKCS#11 库不提供插槽。</p> <p>此库所提供的 PKCS#11 机制由可用硬件提供者确定。</p> <p>应用程序开发者应链接到 <code>libpkcs11.so</code>，而不是直接链接到 <code>pkcs11_kernel.so</code>。请参见 libpkcs11(3LIB)。</p> <p>除以下函数之外，列在 libpkcs11(3LIB) 中的所有标准 PKCS#11 函数都将实现：</p> <div>C_DecryptDigestUpdate C_DecryptVerifyUpdate C_DigestEncryptUpdate C_GetOperationState C_InitToken C_InitPIN C_SetOperationState C_SignEncryptUpdate C_WaitForSlotEvent</div> <p>对这些函数的调用将返回 <code>CKR_FUNCTION_NOT_SUPPORTED</code>。</p> <p>缓冲区不能大于 2 兆字节。例如，可以使用 2 兆字节的纯文本缓冲区和 2 兆字节的密文缓冲区调用 <code>C_Encrypt()</code>。</p> <p>对 <code>C_FindObjects()</code> 的调用可返回的最大对象句柄数为 512。</p> <p>可用于加密操作的最大内核内存量受 <code>project.max-crypto-memory</code> 资源控制限制。内核中对缓冲区的分配以及与会话相关的结构分配都按照此资源控制执行。</p>
返回值	在 RSA PKCS#11 v2.20 规范中定义并列出了其中每个已实现函数的返回值。请参见 http://www.rsasecurity.com 。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
接口稳定性	Committed（已确定）
MT 级别	MT-Safe with exceptions（多线程安全，但存在异常）。请参见 RSA PKCS#11 v2.20 的第 6.6.2 节

属性类型	属性值
标准	PKCS#11 v2.20

另请参见 [cryptoadm\(1M\)](#)、[rctladm\(1M\)](#)、[libpkcs11\(3LIB\)](#)、[attributes\(5\)](#)、[pkcs11_softtoken\(5\)](#)

RSA PKCS#11 v2.20 <http://www.rsasecurity.com>

附注 如果应用程序具有 PKCS#11 插槽的打开会话，则会导致相应硬件提供者不能将驱动程序设为无法装入。管理员必须关闭具有打开 PKCS#11 会话的应用程序后，硬件提供者才能将驱动程序设为无法装入。

引用名	pkcs11_kms – Oracle Key Manager 的 RSA PKCS#11 提供器
用法概要	<pre>/usr/lib/security/pkcs11_kms.so /usr/lib/security/64/pkcs11_kms.so</pre>
描述	<p>pkcs11_kms.so 对象使用 Oracle Key Manager (OKM) KMS 代理协议来实现 RSA Security Inc. PKCS#11 Cryptographic Token Interface (Cryptoki), v2.20, 规范, 从而与 Oracle Key Manager 设备 (Key Manager appliance, KMA) 对话。此提供器实现 PKCS#11 规范并使用 (专用) KMS 客户机协议与远程 OKM 进行通信。</p> <p>此提供器支持以下 PKCS#11 机制: CKM_AES_KEY_GEN、CKM_AES_CBC_PAD 和 CKM_AES_CBC。</p> <p>此提供器支持以下 PKCS#11 接口:</p> <pre>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</pre> <p>在调用时, 所有其他函数均返回 CKR_FUNCTION_NOT_SUPPORTED。</p>

先决条件

pkcs11 kms 提供器只能用于有权访问 OKM 的系统。OKM 管理员必须为每个访问 OKM 的用户（或应用程序）配置一个代理 ID。此操作通过 OKM 实用程序完成，这些实用程序属于 OKM 管理工具，并不是 Oracle Solaris 随附的。

在 OKM 管理员配置了要使用的 KMA 并将参数传递给了客户机（即 Oracle Solaris 用户或应用程序）后，即可初始化 Oracle Solaris PKCS#11 KMS 提供器以供使用。

KMS 提供器的初始化是通过 [kmscfg\(1M\)](#) 实用程序完成的。要初始化本地提供器配置文件以供进一步使用，[kmscfg](#) 要求用户至少输入配置文件的名称、OKM 代理 ID、用于保护配置文件的初始口令以及 KMA 的 IP 地址。有关详细信息，请参见 [kmscfg\(1M\)](#) 手册页。

运行完 [kmscfg](#) 并配置了本地令牌名称空间后，用户即可初始化令牌以供使用。令牌初始化是使用 [pktool\(1\)](#) 命令（如下所示）执行的：

```
$ pktool inittoken currlabel=KMS
```

用户必须提供缺省的 SO（security officer，安全官）PIN（Personal Identification Number，个人识别码）才能初始化 KMS 提供器以供使用。缺省 SO PIN 就是 OKM 管理员在最初设置 OKM 代理时所使用的 PIN。初始化令牌的用户必须知道此口令短语才能初始化提供器。

在初始化提供器后，可以更改用户 PIN 的缺省值。可再次使用 [pktool\(1\)](#) 更改该 PIN 值。

使用以下命令可更改本地 PIN：

```
$ pktool setpin token=KMS
```

为 [pktool setpin](#) 操作提供的 PIN 或通过调用 [C_Login\(\)](#) 和 [C_SetPIN\(\)](#) 函数提供的 PIN 可以是长度介于 1 到 256 之间且不带嵌入空字符的任何字符串。

访问令牌

在用户初始化令牌后，即可通过 [pktool\(1\)](#)、[decrypt\(1\)](#)、[encrypt\(1\)](#) 或通过编写 PKCS11 应用程序并指定 KMS 令牌来开始使用他们的令牌。

示例

示例 1 在 Oracle Key Manager 上创建密钥

以下命令在 Oracle Key Manager 上创建密钥：

```
$ pktool genkey token=KMS label=mykey1 keytype=aes keylen=256
```

示例 2 使用来自 Oracle Key Manager 的密钥对文件进行加密

以下命令使用来自 Oracle Key Manager 的密钥对文件进行加密：

```
$ encrypt -a aes -K mykey1 -T KMS -i input.txt -o output.enc
```

示例 3 使用来自 Oracle Key Manager 的密钥对文件进行解密
以下命令使用来自 Oracle Key Manager 的密钥对文件进行解密：

```
$ decrypt -a aes -K mykey1 -T KMS -i output.enc -o output.txt
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	/system/library/security/crypto/pkcs11_kms
接口稳定性	Committed（已确定）
MT 级别	MT-Safe with Exceptions（多线程安全，但存在异常）。参见下文。
标准	PKCS#11 v2.20

多线程安全属性的异常记录在 RSA PKCS#11 v2.20 的第 6.6.2 节中。

另请参见 [decrypt\(1\)](#)、[encrypt\(1\)](#)、[pktool\(1\)](#)、[cryptoadm\(1M\)](#)、[kmscfg\(1M\)](#)、[libpkcs11\(3LIB\)](#)、[attributes\(5\)](#)

KMS 2.2：管理指南

Oracle Key Manager (OKM) Administration Guide

附注 `pkcs11_kms.so` 使用一个专用目录存放初始化与 KMA 的连接所需的配置文件及其他数据。对于首次创建此专用目录的主机来说，此目录是本地目录。缺省情况下，KMS 令牌目录空间位于 `/var/user/kms/$USERNAME` 中。可以覆盖缺省 KMS 目录，方法是：先设置 `KMSTOKEN_DIR` 环境变量，然后再使用 [kmscfg\(1M\)](#)、[decrypt\(1\)](#)、[encrypt\(1\)](#) 和 [pktool\(1\)](#) 命令。

PKCS#11 客户机需要安装 Oracle Key Manager 软件版本 2.4。

如果 PKCS#11 客户机使用来自多个系统的相同代理 ID，则创建该代理时不应设置 `One Time Passphrase` 标记。在一些成员运行早于 OKM 软件版本 2.4 的 OKM 群集中，该选项不可用。有关创建代理的帮助，请参阅《Oracle Key Manager (OKM) Administration Guide》。

在 KCS#11 客户机使用 OKM 代理创建密钥之前，OKM 代理必须分配 `Default Key Group`。如果未向代理分配 `Default Key Group`，操作将失败，并显示 `CKR_PIN_INCORRECT` 错误。有关向代理分配密钥组的帮助，请参阅《Oracle Key Manager (OKM) Administration Guide》。

引用名	pkcs11_softtoken – 软件 RSA PKCS#11 的软令牌
用法概要	<pre> /usr/lib/security/pkcs11_softtoken.so /usr/lib/security/64/pkcs11_softtoken.so </pre>
描述	<p>pkcs11_softtoken.so 对象在软件中实现 RSA Security Inc. PKCS#11 加密令牌接口 (Cryptographic Token Interface, Cryptoki) v2.20 规范。令牌对象的持久性存储由此 PKCS#11 实现提供。</p> <p>应用程序开发者应链接到 libpkcs11.so 而不是直接链接到 pkcs11_softtoken.so。请参见 libpkcs11(3LIB)。</p> <p>将实现以下加密算法：</p> <p>DES、3DES、AES、Blowfish、RC4、MD5、SHA1、SHA256、SHA384、SHA512、RSA、和 ECC。</p> <p>除以下函数之外，列在 libpkcs11(3LIB) 中的所有标准 PKCS#11 函数都将实现：</p> <pre> C_GetObjectSize C_InitPIN C_InitToken C_WaitForSlotEvent </pre> <p>对这些函数的调用将返回 CKR_FUNCTION_NOT_SUPPORTED。</p> <p>支持以下 RSA PKCS#11 v2.20 机制：</p> <pre> 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 </pre>

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_SHA256_RSA_PKCS
CKM_SHA384_RSA_PKCS
CKM_SHA512_RSA_PKCS

CKM_MD5
CKM_SHA_1
CKM_SHA256
CKM_SHA384
CKM_SHA512

CKM_MD5_HMAC
CKM_MD5_HMAC_GENERAL
CKM_SHA_1_HMAC
CKM_SHA_1_HMAC_GENERAL
CKM_SHA256_HMAC
CKM_SHA256_HMAC_GENERAL
CKM_SHA384_HMAC
CKM_SHA384_HMAC_GENERAL

CKM_MD5_KEY_DERIVATION
CKM_SHA1_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

以下每种类型的密钥对象都具有令牌特定的某些属性，缺省情况下，这些属性将因创建对象、生成密钥/密钥对以及派生密钥而设置为 `true`。

公钥对象 CKA_ENCRYPT、CKA_VERIFY、CKA_VERIFY_RECOVER

私钥对象 CKA_DECRYPT、CKA_SIGN、CKA_SIGN_RECOVER、CKA_EXTRACTABLE

密钥对象 CKA_ENCRYPT、CKA_DECRYPT、CKA_SIGN、CKA_VERIFY、CKA_EXTRACTABLE

支持以下证书对象：

CKC_X_509 对于 CKC_X_509 证书对象，支持以下属性：CKA_SUBJECT、CKA_VALUE、CKA_LABEL、CKA_ID、CKA_ISSUER、CKA_CERTIFICATE_TYPE。

CKC_X_509_ATTR_CERT 对于 CKC_X_509_ATTR_CERT 证书对象，支持以下属性：CKA_OWNER、CKA_VALUE、CKA_LABEL、CKA_SERIAL_NUMBER、CKA_CERTIFICATE_TYPE。

在 `C_FindObjectsInit` 中执行匹配模板的对象搜索操作。匹配的对象将进行缓存以供后续 `C_FindObjects` 操作使用。

`pkcs11_softtoken.so` 对象提供了一个基于文件系统的持久性令牌对象存储，用来存储令牌对象。该令牌对象存储的缺省位置是 `getpwuid_r()` 所返回的用户起始目录。用户可以使用 `${SOFTTOKEN_DIR}` 环境变量覆盖该缺省位置。

如果令牌对象存储从未经过初始化，`C_Login()` 函数可能会返回 `CKR_OK`，但用户无法创建、生成、派生或查找任何专用令牌对象并会收到 `CKR_PIN_EXPIRED`。

用户必须使用 `pktool(1)` `setpin` 命令（其缺省口令短语 "changeme" 是旧口令短语）更改对象存储的口令短语。在初始化新创建的令牌对象存储并为其设置口令短语时，需要执行此操作。

在使用 `pktool setpin` 命令所设置的新口令短语登录到对象存储后，用户即可在此新创建的对象存储中创建和存储专用令牌对象。在 `setpin` 初始化令牌对象存储之前，允许使用 `C_Login()` 函数，但用户创建、生成、派生或查找任何专用令牌对象的所有尝试都会失败并会显示 `CKR_PIN_EXPIRED` 错误。

为 `C_Login()` 和 `C_SetPIN()` 函数提供的 PIN 可以是长度介于 1 到 256 之间且不带嵌入空字符的任何字符串。

令牌对象存储的缺省位置位于 `getpwuid_r()` 所返回的用户起始目录中。它后跟 `/.sunw/pkcs11_softtoken/`，缺省位置为 `~/.sunw/pkcs11_softtoken()`。

用户可以使用 `${SOFTTOKEN_DIR}` 环境变量覆盖该缺省位置。备用令牌对象存储的位置是 `${SOFTTOKEN_DIR}/pkcs11_softtoken/`。

- 返回值

每个已实现函数的返回值都在 RSA PKCS#11 v2.20 规范中定义和列出。请参见 <http://www.rsasecurity.com>
- 文件

`user_home_directory/.sunw/pkcs11_softtoken`

用户的缺省令牌对象存储

`${SOFTTOKEN_DIR}/pkcs11_softtoken`

备用令牌对象存储
- 属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
接口稳定性	Committed（已确定）
MT 级别	MT-Safe with exceptions（多线程安全，但存在异常）。请参见 RSA PKCS#11 v2.20 的第 6.6.2 节。
标准	PKCS#11 v2.20

- 另请参见

[pktool\(1\)](#)、[cryptoadm\(1M\)](#)、[libpkcs11\(3LIB\)](#)、[attributes\(5\)](#)、[pkcs11_kernel\(5\)](#)

RSA PKCS#11 v2.20 <http://www.rsasecurity.com>
- 406

手册页第 5 部分：标准、环境和宏 • 最后修订 2011 年 7 月 12 日

引用名	pkcs11_tpm – 可信平台模块 (Trusted Platform Modules, TPM) 的 RSA PKCS#11 令牌
用法概要	<pre>/usr/lib/security/pkcs11_tpm.so /usr/lib/security/64/pkcs11_tpm.so</pre>
描述	<p>pkcs11_tpm.so 对象使用可信计算组协议与 TPM 安全设备对话，从而实现 RSA Security Inc. PKCS#11 加密令牌接口 (Cryptographic Token Interface, Cryptoki) v2.20 规范。此提供器实现 PKCS#11 规范并在 SUNWtss 软件包中使用 TCG 软件栈 (TCG Software Stack, TSS) API。</p> <p>应用程序开发者应链接到 libpkcs11.so.1，而不是直接链接到 pkcs11_tpm.so。请参见 libpkcs11(3LIB)。</p> <p>将实现以下加密算法：RSA、SHA1 和 MD5。</p> <p>将实现 libpkcs11(3LIB) 中列出的所有标准 PKCS#11 函数，但以下函数除外：</p> <pre>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</pre> <p>支持以下 RSA PKCS#11 v2.20 机制：</p> <pre>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</pre>
基于每个用户的初始化	pkcs11_tpm 提供器只能用于安装有 TPM 设备和 SUNWtss 软件包的系统上。如果满足这些先决条件，用户可以使用 pktool(1) 创建自己的专用令牌，这样他们能够使用 TPM 设备执行操作并通过 TPM 保护的密钥保护自己的专用数据。

要准备并初始化用户的 TPM 令牌，必须执行以下步骤：

1. 初始化令牌。
2. 设置 SO（security officer，安全官）PIN（Personal Identification Number，个人识别码）。
3. 设置用户的唯一 PIN。

令牌初始化是使用 `pktool(1)` 命令（如下所示）执行的：

```
$ pktool inittoken currlabel=TPM newlabel=tpm/myname
```

- 缺省情况下，未初始化的 TPM 按名称 TPM 识别。当用户初始化自己的专用令牌时，可以将其重命名为其他某一名称（例如，tpm/joeuser）或保留为 TPM（在这种情况下，将忽略 newlabel 参数）。
- 用户必须提供缺省的 SO PIN 才能初始化其令牌。缺省 SO PIN 为 87654321。在上述步骤 2 中更改该 PIN。

在初始化令牌之后，必须更改 SO 和用户 PIN 的缺省值。可再次使用 `pktool(1)` 更改这些 PIN 值。

更改 SO PIN：

```
$ pktool setpin token=tpm/joeuser so
```

so 选项指示此 "setpin" 操作将更改 SO PIN，因此 SO PIN 必须存在。随后，用户必须输入缺省的 SO PIN (87654321) 并输入（和确认）新的 PIN。

在重置了 SO PIN 的缺省值后，还必须更改用户的唯一 PIN。

更改用户的 PIN：

```
$ pktool setpin token=tpm/joeuser
```

非 SO 用户的缺省 PIN 是 12345678。用户必须输入缺省 PIN 并输入（和确认）新的唯一 PIN。

为 `pktool setpin` 操作提供的 PIN 或通过调用 `C_Login()` 和 `C_SetPIN()` 函数提供的 PIN 可以是长度介于 1 到 256 之间且不带嵌入空字符的任何字符串。

访问令牌

在用户初始化令牌之后，即可开始通过 `pktool(1)` 或者通过编写 PKCS11 应用程序并使用前面创建的名称（在上例中为 tpm/joeuser）定位令牌来使用他们的令牌。

示例：

```
$ pktool gencert token=tpm/joeuser -i
$ pktool list token=tpm/joeuser
```

- 备注

pkcs11_tpm.so 在文件系统特定的令牌对象存储区域提供对象存储。专用对象由使用专用密钥的加密进行保护，只能通过将令牌的专用密钥装入 TPM 并完全在 TPM 中执行解密的方式进行解密。用户的专用密钥由 TPM 在用户设置其个人 PIN（参见上文）时生成。SO 和用户的密钥都存储在 TSS 持久性存储数据库中并由唯一的 UUID 值引用。所有用户令牌都有唯一的 SO 密钥和唯一的用户密钥，因此一个用户令牌的 PIN 无法解锁同一计算机上其他用户令牌的专用数据。

每个 TPM 都是唯一的，在一个 TPM 上创建的令牌密钥不能用于其他 TPM。pkcs11_tpm.so 令牌数据完全由 TPM 所在的系统进行管理并且不能移到其他系统上。如果 TPM 重置且 SRK（Storage Root Key，存储库根密钥）发生更改，则以前为该 TPM 生成的所有密钥都将不再有效。

pkcs11_tpm.so 将创建一个专用工作区来管理每个已创建令牌的管理文件。缺省情况下，此区域创建为 /var/tpm/pkcs11/\$USERNAME。不过，用户可以覆盖此区域，方法是在初始化或使用令牌之前设置 PKCS11_TPM_DIR 环境变量。
- 返回值

每个已实现函数的返回值都在 RSA PKCS#11 v2.20 规范中定义和列出。请参见 <http://www.rsasecurity.com>。
- 文件

/var/tpm/pkcs11/USERNAME
 用户的缺省令牌对象存储。

\${PKCS11_TPM_DIR}
 备用令牌对象存储。
- 属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
接口稳定性	Committed（已确定）
MT 级别	MT-Safe with Exceptions（多线程安全，但存在异常）（参见下文）
标准	PKCS#11 v2.20

- 多线程安全属性的异常记录在 RSA PKCS#11 v2.20 的第 6.6.2 节中。
- 另请参见

[pktool\(1\)](#)、[cryptoadm\(1M\)](#)、[libpkcs11\(3LIB\)](#)、[attributes\(5\)](#)

TCG 软件栈 (TCG Software Stack, TSS) 规范 <https://www.trustedcomputinggroup.org/specs/TSS>（截至出版日期）

RSA PKCS#11 v2.20, <http://www.rsasecurity.com>

引用名

privileges – process privilege model

描述

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 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`, `nounlink` 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_STORAGE_CONFIG

Allows a process to perform storage configuration, such as disk labelling.

PRIV_STORAGE_RAW

Allows a process to perform raw storage management, such as `USCSICMD`.

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 `sppptun(1M)` `sppptun(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_SUSUSER_COMPAT

Allow a process to successfully call a third party loadable module that calls the kernel `suser()` function to check for allowed access. This privilege exists only for third party loadable module compatibility and is not used by 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** and **PRIV_PROC_AUDIT** 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 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 set-uid 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 set-uid bit of set-uid root applications. The following unsafe privileges have been identified: `proc_setid`, `sys_resource` and `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 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.

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.

另请参见

`mdb(1)`, `ppriv(1)`, `add_drv(1M)`, `ifconfig(1M)`, `lockd(1M)`, `nfsd(1M)`, `pppd(1M)`, `rem_drv(1M)`, `smbd(1M)`, `sppptun(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)`, `prctl(2)`, `prctlset(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)`, `system(4)`, `user_attr(4)`, `ddi_cred(9F)`, `drv_priv(9F)`, `priv_getbyname(9F)`, `priv_policy(9F)`, `priv_policy_choice(9F)`, `priv_policy_only(9F)`

《Oracle Solaris 管理：安全服务》

附注

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.

引用名 prof – profile within a function

用法概要

```
#define MARK
#include <prof.h>

void MARK(name);
```

描述 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.

示例 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++) {
        . . .
    }
}
```

另请参见 [profil\(2\)](#), [monitor\(3C\)](#)

引用名	rbac, RBAC – role-based access control
描述	<p>The addition of role-based access control (RBAC) to the Solaris operating environment gives developers the opportunity to deliver fine-grained security in new and modified applications. RBAC is an alternative to the all-or-nothing security model of traditional superuser-based systems. With RBAC, an administrator can assign privileged functions to specific user accounts (or special accounts called roles).</p> <p>There are two ways to give applications privileges:</p> <ol style="list-style-type: none"> 1. Administrators can assign special attributes such as <code>setUID</code> to application binaries (executable files). 2. Administrators can assign special attributes such as <code>setUID</code> to applications using execution profiles. <p>Special attribute assignment along with the theory behind RBAC is discussed in detail in “Role Based Access Control” chapter of the 《Oracle Solaris 管理：安全服务》. This chapter describes what authorizations are and how to code for them.</p>
Authorizations	<p>An authorization is a unique string that represents a user's right to perform some operation or class of operations. Authorization definitions are stored in a database called <code>auth_attr(4)</code>. For programming authorization checks, only the authorization name is significant.</p> <p>Some typical values in an <code>auth_attr</code> database are shown below.</p> <pre>solaris.jobs.:::Cron and At Jobs::help=JobHeader.html solaris.jobs.grant:::Delegate Cron & At \ Administration::help=JobsGrant.html solaris.jobs.admin:::Manage All Jobs::help=AuthJobsAdmin.html solaris.jobs.user:::Cron & At User::help=JobsUser.html</pre> <p>Authorization name strings ending with the <code>grant</code> suffix are special authorizations that give a user the ability to delegate authorizations with the same prefix and functional area to other users.</p>
Creating Authorization Checks	<p>To check authorizations, use the <code>chkauthattr(3C)</code> library function, which verifies whether or not a user has a given authorization. The synopsis is:</p> <pre>int chkauthattr(const char *authname, const char *username);</pre> <p>The <code>chkauthattr()</code> function checks the <code>policy.conf(4)</code>, <code>user_attr(4)</code>, and <code>prof_attr(4)</code> databases in order for a match to the given authorization.</p> <p>If you are modifying existing code that tests for root UID, you should find the test in the code and replace it with the <code>chkauthattr()</code> function. A typical root UID check is shown in the first</p>

code segment below. An authorization check replacing it is shown in the second code segment; it uses the `solaris.jobs.admin` authorization and a variable called `real_login` representing the user.

示例1 Standard root check

```
ruid = getuid();

if ((eflag || lflag || rflag) && argc == 1) {
    if ((pwp = getpwnam(*argv)) == NULL)
        crabort(INVALIDUSER);

    if (ruid != 0) {
        if (pwp->pw_uid != ruid)
            crabort(NOTROOT);
        else
            pp = getuser(ruid);
    } else
        pp = *argv++;
} else {
```

示例2 Authorization check

```
ruid = getuid();
if ((pwp = getpwuid(ruid)) == NULL)
    crabort(INVALIDUSER);

strcpy(real_login, pwp->pw_name);

if ((eflag || lflag || rflag) && argc == 1) {
    if ((pwp = getpwnam(*argv)) == NULL)
        crabort(INVALIDUSER);

    if (!chkauthattr("solaris.jobs.admin", real_login)) {
        if (pwp->pw_uid != ruid)
            crabort(NOTROOT);
        else
            pp = getuser(ruid);
    } else
        pp = *argv++;
} else {
```

For new applications, find an appropriate location for the test and use `chkauthattr()` as shown above. Typically the authorization check makes an access decision based on the identity of the calling user to determine if a privileged action (for example, a system call) should be taken on behalf of that user.

Applications that perform a test to restrict who can perform their security-relevant functionality are generally `setuid` to root. Programs that were written prior to RBAC and that are only available to the root user may not have such checks. In most cases, the kernel requires an effective user ID of root to override policy enforcement. Therefore, authorization checking is most useful in programs that are `setuid` to root.

For instance, if you want to write a program that allows authorized users to set the system date, the command must be run with an effective user ID of root. Typically, this means that the file modes for the file would be `-rwsr-xr-x` with root ownership.

Use caution, though, when making programs `setuid` to root. For example, the effective UID should be set to the real UID as early as possible in the program's initialization function. The effective UID can then be set back to root after the authorization check is performed and before the system call is made. On return from the system call, the effective UID should be set back to the real UID again to adhere to the principle of least privilege.

Another consideration is that `LD_LIBRARY` path is ignored for `setuid` programs (see SECURITY section in [ld.so.1\(1\)](#)) and that shell scripts must be modified to work properly when the effective and real UIDs are different. For example, the `-p` flag in Bourne shell is required to avoid resetting the effective UID back to the real UID.

Using an effective UID of root instead of the real UID requires extra care when writing shell scripts. For example, many shell scripts check to see if the user is root before executing their functionality. With RBAC, these shell scripts may be running with the effective UID of root and with a real UID of a user or role. Thus, the shell script should check `euid` instead of `uid`. For example,

```
WHO='id | cut -f1 -d" "'
if [ ! "$WHO" = "uid=0(root)" ]
then
    echo "$PROG: ERROR: you must be super-user to run this script."
    exit 1
fi
```

should be changed to

```
WHO='/usr/xpg4/bin/id -n -u'
if [ ! "$WHO" = "root" ]
then
    echo "$PROG: ERROR: you are not authorized to run this script."
    exit 1
fi
```

Authorizations can be explicitly checked in shell scripts by checking the output of the [auths\(1\)](#) utility. For example,

```
for auth in `auths | tr , " ` NOTFOUND
do
```

```
    [ "$auth" = "solaris.date" ] && break          # authorization found
done

if [ "$auth" != "solaris.date" ]
then
    echo >&2 "$PROG: ERROR: you are not authorized to set the date"
    exit 1
fi
```

另请参见

[ld.so.1\(1\)](#), [chkauthattr\(3C\)](#), [auth_attr\(4\)](#), [policy.conf\(4\)](#), [prof_attr\(4\)](#), [user_attr\(4\)](#)

《Oracle Solaris 管理：安全服务》

引用名	regex – internationalized basic and extended regular expression matching
描述	<p>Regular Expressions (REs) provide a mechanism to select specific strings from a set of character strings. The Internationalized Regular Expressions described below differ from the Simple Regular Expressions described on the regex(5) manual page in the following ways:</p> <ul style="list-style-type: none">▪ both Basic and Extended Regular Expressions are supported▪ the Internationalization features—character class, equivalence class, and multi-character collation—are supported. <p>The Basic Regular Expression (BRE) notation and construction rules described in the 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).</p>
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	<p>An ordinary character is a BRE that matches itself: any character in the supported character set, except for the BRE special characters listed in BRE Special Characters, below.</p> <p>The interpretation of an ordinary character preceded by a backslash (\) is undefined, except for:</p> <ol style="list-style-type: none">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	<p>A BRE <i>special character</i> has special properties in certain contexts. Outside those contexts, or when preceded by a backslash, such a character will be a BRE that matches the special character itself. The BRE special characters and the contexts in which they have their special meaning are:</p> <p>. [\ 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.</p> <p>* The asterisk is special except when used:</p> <ul style="list-style-type: none">▪ in a bracket expression▪ as the first character of an entire BRE (after an initial ^, 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:

[:alnum:]	[:cntrl:]	[:lower:]	[:space:]
[:alpha:]	[:digit:]	[:print:]	[:upper:]
[:blank:]	[:graph:]	[:punct:]	[:xdigit:]

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> <code>[::]</code> <code>[..]</code>
escaped characters	<code>\<special character></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 specififed for BREs apply to Extended Regular Expressions (ERE) 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	<p>An <i>ERE special character</i> has special properties in certain contexts. Outside those contexts, or when preceded by a backslash, such a character is an ERE that matches the special character itself. The extended regular expression special characters and the contexts in which they have their special meaning are:</p> <ul style="list-style-type: none"> . [\ (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: <ul style="list-style-type: none"> ▪ if these characters appear first in an ERE, or immediately following a vertical-line, circumflex or left-parenthesis ▪ if a left-brace is not part of a valid interval expression. 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: <ul style="list-style-type: none"> ▪ 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	<p>The following rules will be used to construct EREs matching multiple characters from EREs matching a single character:</p> <ol style="list-style-type: none"> 1. A <i>concatenation of EREs</i> matches the concatenation of the character sequences matched by each component of the ERE. A concatenation of EREs enclosed in parentheses matches whatever the concatenation without the parentheses matches. For example, both the ERE <code>cd</code> and the ERE <code>(cd)</code> are matched by the third and fourth character of the string <code>abcdefabcdef</code>.

- 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 `cabbbbcdebbbbbbcd`. 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	[= =] [::] [. .]
escaped characters	\<special character>
bracket expression	[]

grouping	()
single-character-ERE duplication	* + ? {m,n}
concatenation	
anchoring	^ \$
alternation	

For example, the ERE `abba | cde` matches either the string `abba` or the string `cde` (rather than the string `abbade` or `abbcde`, because concatenation has a higher order of precedence than alternation).

ERE Expression
Anchoring

An ERE can be limited to matching strings that begin or end a line; this is called *anchoring*. The circumflex and dollar sign special characters are considered ERE anchors when used anywhere outside a bracket expression. This has the following effects:

1. A circumflex (^) outside a bracket expression anchors the expression or subexpression it begins to the beginning of a string; such an expression or subexpression can match only a sequence starting at the first character of a string. For example, the EREs `^ab` and `(^ab)` match `ab` in the string `abcdef`, but fail to match in the string `cdefab`, and the ERE `a^b` is valid, but can never match because the `a` prevents the expression `^b` from matching starting at the first character.
2. A dollar sign (\$) outside a bracket expression anchors the expression or subexpression it ends to the end of a string; such an expression or subexpression can match only a sequence ending at the last character of a string. For example, the EREs `ef$` and `(ef$)` match `ef` in the string `abcdef`, but fail to match in the string `cdefab`, and the ERE `e$f` is valid, but can never match because the `f` prevents the expression `e$` from matching ending at the last character.

另请参见

[localedef\(1\)](#), [regcomp\(3C\)](#), [attributes\(5\)](#), [environ\(5\)](#), [locale\(5\)](#), [regexp\(5\)](#)

引用名

regex, compile, step, advance – simple regular expression compile and match routines

用法概要

```
#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);
```

描述

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 `locs` is equal to the point in the string at sometime during the backing up process, `advance()` will break out of the loop that backs up and will return zero.

The external variables `circf`, `sed`, and `nbra` are reserved.

示例

示例 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;
```

诊断

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.

另请参见

[regex\(5\)](#)

引用名	resource_controls – resource controls available through project database
描述	<p>The resource controls facility is configured through the project database. See project(4). You can set and modify resource controls through the following utilities:</p> <ul style="list-style-type: none"> ▪ prctl(1) ▪ projadd(1M) ▪ projmod(1M) ▪ rctladm(1M) <p>In a program, you use setrctl(2) to set resource control values.</p> <p>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.</p> <p>The following are the resource controls are available:</p> <p><code>process.max-address-space</code> Maximum amount of address space, as summed over segment sizes, that is available to this process, expressed as a number of bytes.</p> <p><code>process.max-core-size</code> Maximum size of a core file created by this process, expressed as a number of bytes.</p> <p><code>process.max-cpu-time</code> Maximum CPU time that is available to this process, expressed as a number of seconds.</p> <p><code>process.max-data-size</code> Maximum heap memory available to this process, expressed as a number of bytes.</p> <p><code>process.max-file-descriptor</code> Maximum file descriptor index available to this process, expressed as an integer.</p> <p><code>process.max-file-size</code> Maximum file offset available for writing by this process, expressed as a number of bytes.</p> <p><code>process.max-msg-messages</code> Maximum number of messages on a message queue (value copied from the resource control at <code>msgget()</code> time), expressed as an integer.</p> <p><code>process.max-msg-qbytes</code> Maximum number of bytes of messages on a message queue (value copied from the resource control at <code>msgget()</code> time), expressed as a number of bytes.</p> <p><code>process.max-port-events</code> Maximum allowable number of events per event port, expressed as an integer.</p> <p><code>process.max-sem-nsems</code> Maximum number of semaphores allowed per semaphore set, expressed as an integer.</p>

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-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	1
		KB	2^10
		MB	2^20
		GB	2^30
		TB	2^40
		PB	2^50
		EB	2^60
Time	seconds	s	1
		Ks	10^3
		Ms	10^6
		Gs	10^9
		Ts	10^12
		Ps	10^15
		Es	10^18
Count	integer	none	1
		K	10^3
		M	10^6
		G	10^9
		T	10^12
		P	10^15
		Es	10^18

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.

属性

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

另请参见

`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 管理：Oracle Solaris Zones、Oracle Solaris 10 Zones 和资源管理》

引用名	sgml, solbook – Standard Generalized Markup Language	
描述	<p>Standard Generalized Markup Language (SGML) is the ISO standard 8879:1986 that describes a syntax for marking up documents with tags that describe the purpose of the text rather than the appearance on the page. This form of markup facilitates document interchange between different platforms and applications. SGML allows the management of information as data objects rather than text on a page.</p> <p>In an SGML document the main structural components are called elements. The organization and structure of a document and the meaning of elements are described in the Document Type Definition (DTD). Elements are the <i>tags</i> that identify the content. Element names may be descriptive of the content for ease of use. For example <code><para></code> for paragraphs. Elements can have 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.</p> <p>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.</p> <p>There is no single standard DTD , but the de facto standard for the computer industry is the DocBook DTD , developed and maintained by the Davenport Group. Within Sun, the SolBook DTD , which is a proper subset of DocBook DTD , is used when writing reference manual pages. The SolBook DTD contains a number of tags that are designed for the unique needs of the reference pages.</p>	
SolBook Elements	Elements are defined with a hierarchical structure that gives a structure to the document. The following is a description of some of the elements from the SolBook DTD which are used for reference pages.	
DOCTYPE	The first line in an SGML file that identifies the location of the DTD that is used to define the document. The <code><!DOCTYPE</code> string is what the SGML -aware <code>man(1)</code> command uses to identify that a file is formatted in SGML rather than <code>nroff(1)</code> .	
RefEntry	The top layer element that contains a reference page is <code><refentry></code> . All of the text and other tags must be contained within this tag.	
RefMeta	The next tag in a reference page is <code><refmeta></code> , which is a container for several other tags. They are:	
	<code><refentrytitle></code>	This is the title of the reference page. It is equivalent to the name of the reference page's file name, without the section number extension.
	<code><manvolnum></code>	This is the section number that the reference page resides in. The contents may be a text entity reference.

<code><refmiscinfo></code>	There are one or more <code><refmiscinfo></code> tags which contain <i>meta</i> information. Meta information is information about the reference page. The <code><refmiscinfo></code> tag has the <code>class</code> attribute. There are four classes that are routinely used.
<code>date</code>	This is the date that the file was last modified. By consensus this date is changed only when the technical information on the page changes and not simply for an editorial change.
<code>sectdesc</code>	This is the section title of the reference page; for example User Commands. The value of this attribute may be a text entity reference.
<code>software</code>	This is the name of the software product that the topic discussed on the reference page belongs to. For example UNIX commands are part of the SunOS x.x release. The value of this attribute may be a text entity reference.
<code>arch</code>	This is the architectural platform limitation of the subject discussed on the reference page. If there are no limitations the value used is generic. Other values are <code>sparc</code> and <code>x86</code> .
<code>copyright</code>	This attribute contains the Sun Microsystems copyright. Any other copyrights that may pertain to the individual reference page file should be entered as separate <code><refmiscinfo></code> entries. The value of this attribute may be a text entity reference.

RefNameDiv

This tag contains the equivalent information to the `.TH` macro line in an `nroff(1)` reference page. `<refnamediv>` contains three tags. These tags contain the text that is before and after the `'-` (dash) on the `NAME` line.

<code><refname></code>	These are the names of the topics that are discussed in the file. There may be more than one <code><refname></code> for a page. The first <code><refname></code> must match the name of the file and the <code><refentrytitle></code> . If there are more than one <code><refname></code> tags, each is separated by a <code>'</code> (comma). The comma is generated by the publisher of sgml files, so it should not be typed. This is referred to as <i>auto-generated</i> text.
<code><refpurpose></code>	The text after the dash on the <code>NAME</code> line is contained in this tag. This is a short summary of what the object or objects described on the reference page do or are used for. The dash is also auto-generated and should not be typed in.
<code><refdiscriptor></code>	In some cases the <code><refentrytitle></code> is a general topic descriptor of a group of related objects that are discussed on the same page. In this case the first tag after the <code><refnamediv></code> is a <code><refdiscriptor></code> . The <code><refname></code> tags follow. Only one <code><refdiscriptor></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.

<code><example></code>	This tag contains examples of source code or usage of commands. It contains a required <code><title></code> .
<code><informalexample></code>	This tag is the same as the <code><example></code> tag except the <code><title></code> is not required.

Inline Elements

The inline elements are used for tagging text.

<code><command></code>	An executable program or the entry a user makes to execute a command.
<code><function></code>	A subroutine in a program or external library.
<code><literal></code>	Contains any literal string.
<code><parameter></code>	An argument passed to a computer program by a function or routine.
<code><inlineequation></code>	An untitled mathematical equation occurring in-line.
<code><link></code>	A hypertext link to text within a book, in the case of the reference manual it is used to cross reference to another reference page.
<code><olink></code>	A hypertext link used to create cross references to books other than the reference manual.
<code><xref></code>	A cross reference to another part of the same reference page.

另请参见

[man\(1\)](#), [nroff\(1\)](#), [man\(5\)](#)

引用名	smf – 服务管理工具								
描述	<p>Solaris 服务管理工具定义了一种编程模型，用以提供持久运行的应用程序，这些应用程序称为服务。此工具还提供运行服务的基础结构。服务可以表示运行的应用程序、设备软件状态或一组其他服务。服务在框架中由服务实例对象表示，这些对象是服务对象的子项。实例对象可以继承或覆盖父服务对象的配置，这使多个服务实例可以共享配置信息。所有服务和实例对象都包含在一个范围内，该范围表示配置信息的集合。本地 Solaris 实例的配置称为“本地主机”范围，这是当前唯一支持的范围。</p> <p>每个服务实例都使用故障管理资源标识符 (fault management resource identifier, FMRI) 来命名（模式为 svc:）。例如，在系统启动时启动的 syslogd(1M) 守护进程是具有以下名称的缺省服务实例：</p> <pre>svc://localhost/system/system-log:default svc:/system/system-log:default system/system-log:default</pre> <p>许多命令也允许使用 FMRI 缩写。有关此类示例，请参见 svcs(1) 手册页。</p> <p>在上例中，'default' 是实例名称，'system/system-log' 是服务名称。服务名称可包含多个用斜杠 (/) 分隔的组成部分。所有组成部分（最后一个组成部分除外）构成了服务的类别。站点特定的服务应当用以 'site' 开头的类别来命名。</p> <p>服务实例可以启用，也可以禁用。所有服务都可以通过 svcadm(1M) 命令来启用或禁用。</p> <p>可通过 svcs(1) 命令来显示系统上管理服务实例的列表。</p>								
相关性	<p>服务实例可以与一组实体具有相关性，实体可包含服务、实例和文件。相关性控制着启动和自动停止服务的时间。当未满足已启用服务的相关性时，服务就保持脱机状态。当满足其相关性时，就会启动服务。如果启动成功，服务将转换为联机状态。不同于服务和实例，不会在创建或删除文件时动态地评估文件相关性。只对它们评估一次。</p> <p>相关性是否满足将由其分组确定：</p> <table><tr><td>require_all</td><td>当所有引用的服务正在运行（联机或降级）或者所有指示的文件都存在时满足。</td></tr><tr><td>require_any</td><td>当引用的任一服务正在运行（联机或降级）或者至少有一个指示的文件存在时满足。</td></tr><tr><td>optional_all</td><td>当引用的服务正在运行（联机或降级）或因没有管理操作而未运行（禁用、维护、不存在或脱机等待必须有管理操作才能启动的相关项）时满足。不完整的服务也满足可选的相关性。</td></tr><tr><td>exclude_all</td><td>当所有引用的服务都已禁用、处于维护状态或者引用的服务或文件不存在时满足。</td></tr></table>	require_all	当所有引用的服务正在运行（联机或降级）或者所有指示的文件都存在时满足。	require_any	当引用的任一服务正在运行（联机或降级）或者至少有一个指示的文件存在时满足。	optional_all	当引用的服务正在运行（联机或降级）或因没有管理操作而未运行（禁用、维护、不存在或脱机等待必须有管理操作才能启动的相关项）时满足。不完整的服务也满足可选的相关性。	exclude_all	当所有引用的服务都已禁用、处于维护状态或者引用的服务或文件不存在时满足。
require_all	当所有引用的服务正在运行（联机或降级）或者所有指示的文件都存在时满足。								
require_any	当引用的任一服务正在运行（联机或降级）或者至少有一个指示的文件存在时满足。								
optional_all	当引用的服务正在运行（联机或降级）或因没有管理操作而未运行（禁用、维护、不存在或脱机等待必须有管理操作才能启动的相关项）时满足。不完整的服务也满足可选的相关性。								
exclude_all	当所有引用的服务都已禁用、处于维护状态或者引用的服务或文件不存在时满足。								

在运行（联机或降级）期间，如果由 `require_all`、`require_any` 或 `optional_all` 相关性引用的服务停止或刷新，SMF（Service Management Facility，服务管理工具）将考虑服务停止的原因以及相关性的 `restart_on` 属性以确定是否停止该服务。

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

如果服务出现硬件错误或软件错误（如核心转储），则认为该服务已经因出错而停止。对于 `exclude_all` 相关性，如果引用的服务已启动，但 `restart_on` 属性不是 `none`，也将停止该服务。

服务的相关性可通过 `svcs(1)` 或 `svccfg(1M)` 列出，也可通过 `svccfg(1M)` 进行修改。

重启程序

每个服务都由一个重启程序来管理。主重启程序 `svc.startd(1M)` 管理整个服务实例集及其相关项的状态。主重启程序代表其服务操作，其操作对象是委托重启程序，委托重启程序可以为特定的应用程序类提供特定的执行环境。例如，`inetd(1M)` 是一个委托重启程序，它为其服务实例提供一个由网络连接构成的初始环境作为输入和输出文件描述符。委托给 `inetd(1M)` 的每个实例都处于联机状态。虽然特定实例的守护进程可能未运行，但该实例已经可运行。

由于当实例转为联机状态时将满足相关性，因此 `svc.startd(1M)` 将调用其他实例的启动方法或指示委托重启程序执行此类操作。这些操作可能会重叠。

可使用 `svcs(1)` 检查当前的服务集及关联的重启程序。所有重启程序使用的公共配置的说明在 `smf_restarter(5)` 中指定。

方法

每个服务或服务实例都必须定义一组启动、停止和（可选）刷新该服务的方法。有关 `svc.startd(1M)` 以及类似 `fork(2)-exec(2)` 重启程序的方法约定的更完整说明，请参见 `smf_method(5)`。

管理方法（如将传统配置信息捕获到系统信息库的方法）将在 `svccfg(1M)` 手册页中讨论。

服务的方法可使用 `svccfg(1M)` 命令列出和修改。

状态

每个服务实例始终处于某一明确定义的状态，具体取决于其相关性、其方法的执行结果以及其潜在合同事件。系统定义了以下状态：

UNINITIALIZED 这是所有服务实例的初始状态。实例将根据 `svc.startd(1M)` 或相应重启程序的计算结果转为维护、脱机或禁用状态。

OFFLINE (脱机)	实例已启用，但尚未运行或无法运行。如果重启程序成功执行服务启动方法或等效方法，实例将转为联机状态。如果执行失败，则会导致实例处于降级或维护状态。管理操作会导致实例处于初始化状态。
ONLINE	实例已启用且正在运行或可以运行。联机状态的具体特性是应用程序/模型特定的，由负责服务实例的重启程序进行定义。当满足所有相关性时，联机是正确配置的服务的预期运行状态。实例失败会导致其处于降级或维护状态。如果实例所依赖的服务失败，则会导致其处于脱机或降级状态。
DEGRADED	实例已启用且正在运行或可以运行。不过，相对于正常运行而言，实例是以有限的功能运行的。实例失败会导致其处于维护状态。如果实例所依赖的服务失败，则会导致其处于脱机或降级状态。功能恢复后，应会将实例转为联机状态。
MAINTENANCE	实例已启用，但无法运行。需要管理操作（通过 <code>svcadm clear</code> 执行）才能使实例脱离维护状态。如果管理操作正在进行中，维护状态可能是临时所处的状态。
DISABLED	实例已禁用。使服务转换为脱机状态，最终只有满足所有相关性时才会恢复联机状态。
LEGACY-RUN	此状态表示不通过服务管理工具进行管理的传统实例。处于此状态的实例已在某一时间点启动，可能正在运行，也可能未在运行。只能使用该工具观察实例，不能将实例转换为其他状态。

状态也会发生转换，但会导致返回到最初的状态。

事件通知 SMF 允许使用 SNMP（Simple Network Management Protocol，简单网络管理协议）或 SMTP（Simple Mail Transfer Protocol，简单邮件传输协议）通知状态转换。该工具会发布状态转换的信息事件，以供通知守护进程（如 `snmp-notify(1M)` 和 `sasinfo(1M)`）使用。已禁用服务的 SMF 状态转换不会生成通知，除非转换的最终状态为禁用且该转换存在通知参数。对于初始和最终状态相同的转换，也不会为其生成通知。

通知参数 FMA（Fault Management Architecture，故障管理体系结构）事件的通知参数存储在 `svc:/system/fm/notify-params:default` 中，但 SMF 状态转换所生成的信息事件除外。这些事件存储在服务中或存储在转换服务的实例中。在 `svc:/system/svc/global:default` 中，可在系统范围设置 SMF 状态转换所生成事件的通知参数。当在转换实例中找不到组合查找（如 `scf_instance_get_pg_composed(3SCF)` 中所述）时，将使用系统范围通知参数。可使用 `svccfg(1M)` 处理通知参数。可使用 DTD（document type definition，文档类型定义）中描述的 `notification_parameters` 元素在服务清单或配置文件中配置通知参数。下面列举了一个示例：

```

<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 是一个用逗号分隔的 SMF 状态转换集列表或一个用逗号分隔的 FMA 事件类列表。*events* 不能混合包含 SMF 状态转换集和 FMA 事件类。

为了方便起见，标记 `problem- {diagnosed,updated,repaired,resolved}` 描述了 FMA 子系统诊断出的问题的生命周期：包括初始诊断、中期更新和最终问题解决。这些标记是基础 FMA 协议事件类 (全部位于 `list.*` 分层结构中) 的别名，但在配置通知首选项时不能使用后者。

problem-diagnosed

FMA 子系统诊断出了新问题。诊断包含由一个或多个可疑项组成的列表，系统可能（根据需要）已将这些可疑项自动隔离以防止进一步发生错误。问题在事件有效载荷中问题由 UUID（Universally Unique Identifier，通用唯一标识符）进行标识，描述此问题解决生命周期的其他事件会引用匹配的 UUID。

problem-updated

问题诊断中的一个或多个可疑资源已修复、替换或被视为没有故障（或再次出现了故障），但列表中至少保留了一个故障资源。修复可能是 `fmadm` 命令行（`fmadm repaired`、`fmadm acquit`、`fmadm replaced`）的结果，也可能是自动检测到的（如通过检测部件序列号更改）。

problem-repaired

问题诊断中的所有可疑资源均已修复、解决或被视为无故障。在此阶段，部分或全部资源可能仍处于隔离状态。

problem-resolved

问题诊断中的所有可疑资源均已修复、解决或被视为无故障并且不再处于隔离状态（例如，曾经是可疑项并处于脱机状态的 CPU 现在又重新恢复联机状态；这种取消隔离操作通常自动执行）。

状态转换集的定义如下：

<code>to-<state></code>	将 <状态> 作为转换的最终状态的所有转换的集合。
<code>from-<state></code>	将 <状态> 作为转换的初始状态的所有转换的集合。
<code><状态></code>	将 <状态> 作为转换的初始状态的所有转换的集合。
<code>all</code>	所有转换的集合。

状态的有效值包括：维护、脱机、禁用、联机和降级。转换集定义的示例如下：maintenance、from-online、to-degraded。

属性和属性组 上述相关性、方法、委托重启程序和实例状态以服务或服务实例的属性或属性组形式表示。服务或服务实例具有任意数量的属性组，这些属性组用于存储应用程序数据。以这种方式使用特性组时，可以使应用程序的配置派生可供系统信息库提供给该工具中的所有数据的属性。应用程序也可以在框架中使用适当的 [service_bundle\(4\)](#) DTD 子集表示其配置数据。

属性查找是组合查找。如果在服务实例中未找到属性组/属性组合，大部分命令以及 [libscf\(3LIB\)](#) 的高级接口将在包含该实例的服务中搜索相同的属性组/属性组合。这样可以在服务实例之间共享公共配置。组合可视为服务实例与其父服务之间的继承关系。

属性将受到保护，以防被未经授权的进程修改。请参见 [smf_security\(5\)](#)。

通用属性组 通用属性组适用于所有服务实例。该属性组包含以下属性：

- enabled (boolean)** 指定是否启用实例。如果实例上不存在此属性，SMF 将无法告知实例的重启程序有关重启程序的存在情况。
- restarter (fmri)** 此服务的重启程序。有关更多信息，参见“重启程序”一节。如果未设置此属性，将使用缺省的系统重启程序。
- complete (astring)** 该服务是完整的，还是不应启动的部分定义。该属性在清单导入时或管理员手动创建实例时自动设置，因此应不需要修改。

层 系统信息库由标准位置中文件的管理定制、当前状态和缺省值的组合集合而成。由 SMF 管理的文件系统位置中的清单定义的服务、实例、属性组和属性始终在系统信息库中准确表示。运行时由管理员或其他程序执行的定制被捕获并存储在系统信息库中。

属性在系统信息库中可以具有不同的值，用于反映清单、配置文件和管理定制的不同设置。缺省情况下向用户和服务提供哪一个值由名为 `layers` 的简单优先级机制仲裁。

SMF 跟踪四个层。按优先级降序，它们是：

- admin** 由交互使用 SMF 命令或库进行的任何更改。此层具有最高优先级。
- site-profile（站点配置文件）** `/etc/svc/profile/site` 目录中的文件，或传统 `/etc/svc/profile/site.xml` 和 `/var/svc/profile/site.xml` 文件中的任何值。
- system-profile（系统配置文件）** 系统配置文件位置 `/etc/svc/profile/generic.xml` 和 `/etc/svc/profile/platform.xml` 中的任何值。

manifest（清单） 系统清单位置 `/lib/svc/manifest` 和 `/var/svc/manifest` 中的任何值。

任何单独的层中不能有属性冲突。`admin` 层中的冲突属性只是覆盖之前的属性。如果任何其他层中的多个文件提供该同一属性，并且该属性未在更高的层中设置，则整个实例标记为冲突，并且不会由 `svc.startd(1M)` 启动，直到冲突定义被去除，或在更高的层设置该属性。请求单个值（包括 `svccfg` 和 `svccprop`）的其他 `libscf` 使用者看到所有适当的值中的随机属性设置。我们不保证返回哪个冲突值。

快照 有关系统信息库中每个实例的历史数据由服务管理工具进行维护。此数据以只读快照形式提供，用于管理检查和回滚。可能会提供以下快照类型集：

- `initial` 由管理员创建的或安装软件包期间生成的实例的初始配置。
- `previous` 执行管理撤消操作时捕获的当前配置。
- `running` 正在运行的实例配置。
- `start` 成功转换为联机状态期间所捕获的配置。

可使用 `svccfg(1M)` 命令与快照进行交互。

专用属性组 某些属性组将标记为 `non-persistent`。这些组不会在快照中进行备份，其内容在系统引导期间会被清除。此类组通常存放活动程序状态，此状态无需在系统重新启动后存在。

配置信息库 每个服务实例的当前状态以及与服务和实例关联的属性都存储在由 `svc.configd(1M)` 管理的系统信息库中。

服务管理工具数据的系统信息库由 `svc.configd(1M)` 管理。

服务束、清单和配置文件 与服务或服务实例关联的信息存储在配置系统信息库中，可作为 XML 文件导出。此类 XML 文件称作服务束，是可移植的，很适合用于备份。服务束分为以下几种类型：

- 清单** 包含与特定服务或服务实例集关联的完整属性集的文件。
- 配置文件** 包含每个实例的 `enabled` 属性（在通用属性组中为布尔型）的一组服务实例和值的文件。

配置文件也可以包含服务和实例中属性的配置值。模板元素不能在配置文件中定义。

配置文件可以使用来自 DTD 的一组宽松元素，如 `service_bundle(4)` 中所述。要使用这些元素，应向 `DOCTYPE` 条目添加以下定义：

```
<!ENTITY % profile "INCLUDE">
<!ENTITY % manifest "IGNORE">
```

可使用 [svccfg\(1M\)](#) 命令在信息库中导入或导出服务束。有关服务束文件格式的说明以及有关制作服务束的准则，请参见 [service_bundle\(4\)](#)。

里程碑

smf 里程碑是一种服务，它聚合了多个服务相关性。通常，里程碑本身并不执行任何有用的操作，但它声明一种其他服务可以依赖的特定系统就绪状态。一个示例就是**名称/服务**里程碑，它仅依赖当前启用的名称服务。

传统启动脚本

`/etc/rc?.d` 目录中的启动程序作为相应运行级别的里程碑的一部分执行：

```
/etc/rcS.d    milestone/single-user:default
```

```
/etc/rc2.d    milestone/multi-user:default
```

```
/etc/rc3.d    milestone/multi-user-server:default
```

每个程序的执行都表示为程序路径所指定的功能简化的服务实例。这些实例保持特殊的 `legacy-run` 状态。

这些实例没有已启用属性（在通用属性组中为 `boolean` 类型），通常无法使用 [svcadm\(1M\)](#) 命令进行处理。系统不会对这些程序执行错误诊断或重新启动。

另请参见

[svcs\(1\)](#)、[inetd\(1M\)](#)、[snmp-notify\(1M\)](#)、[sasinfo\(1M\)](#)、[svcadm\(1M\)](#)、[svccfg\(1M\)](#)、[svc.configd\(1M\)](#)

引用名	smf_bootstrap – 服务管理工具的引导、打包和兼容性行为
描述	<p>服务管理工具在提供服务清单、合并服务清单更改、描述服务配置稳定性、使用服务配置覆盖以及使用服务配置文件等方面建立了相关约定。</p>
引导时的清单装入	<p>在系统引导期间，以及管理员或程序运行时，处理标准目录树 <code>/lib/svc/manifest</code> 和 <code>/var/svc/manifest</code> 的清单：</p> <pre>\$ svcadm restart manifest-import</pre> <p>处理之前未导入，或自上一次导入以来发生更改的清单。使用散列来确定清单是否已更改。</p> <p>当标准位置中的清单第一次导入时，它的属性、实例和服务作为 <code>manifest</code> 层的一部分添加到系统信息库。</p> <p>当标准位置中的清单更新时，将自动导入它们。添加新的服务和实例，升级属性（如果它们发生变化），并在服务、实例和属性被移除时将其删除。</p> <p>在引导期间，将在两个不同的阶段处理清单。</p> <p>服务 <code>svc:/system/early-manifest-import:default</code> 是一项伪服务，负责首次清单处理。此服务仅在 <code>svc.startd(1M)</code> 初始化任何服务之前处理 <code>/lib/svc/manifest</code> 目录树中的清单，从而使 <code>/lib/svc/manifest</code> 中提供的服务始终以最新定义启动。由于此服务是伪服务，因此将忽略 <code>svcadm(1M)</code> 命令，但可以使用 <code>svcs(1)</code> 观察状态及获取日志文件信息。</p> <p><code>svc:/system/manifest-import:default</code> 服务负责第二次清单处理并按相应顺序依次导入 <code>/lib/svc/manifest</code> 和 <code>/var/svc/manifest</code> 目录树中的清单文件。</p> <p>对 <code>/var/svc/manifest</code> 的支持是对引入 <code>system/early-manifest-import:default</code> 之前该目录树中所提供清单的兼容性支持。当服务可能以旧系统信息库配置启动时，由于其更新的清单尚未导入，因此 <code>/var/svc/manifest</code> 中所提供的服务可能会出现与升级相关的问题。同样，由于其清单文件尚未得到处理，因此新添加的服务也可能不可用或者已删除的服务仍在引导期间启动。为避免这些问题，强烈建议开发者将清单移至 <code>/lib/svc/manifest</code>。</p>
配置文件应用	<p><code>early-manifest-import</code> 和 <code>manifest-import</code> 服务还应用配置文件。</p> <p><code>/etc/svc/profile/generic.xml</code> 和 <code>/etc/svc/profile/platform.xml</code> 中的系统提供的配置文件导入到 <code>system-profile</code> 层。</p> <p><code>/etc/svc/profile/site</code> 目录中的站点特定的配置文件和传统站点文件 <code>/etc/svc/profile/site.xml</code> 和 <code>/var/svc/profile/site.xml</code> 导入到 <code>site-profile</code> 层。</p> <p>管理员可以通过运行以下命令请求重新应用这些配置文件：</p> <pre>\$ svcadm restart manifest-import</pre>

	配置文件定义的属性、实例和服务的行为与由清单定义的相同。
打包操作期间的清单处理	<p>软件包中的所有服务清单都应由类 <code>manifest</code> 来标识。安装和移除服务清单的类操作脚本包含在软件包系统中。调用 <code>pkg install</code> 时，会导入服务清单。</p> <p>调用 <code>pkg uninstall</code> 时，会删除清单中禁用的实例。清单中联机或降级的实例将先被禁用，然后再被删除。清单中没有剩余实例的所有服务也将被删除。</p>
稳定性声明	<p>清单中所提供的每个服务组和每个属性组都应根据 attributes(5) 定义声明稳定性级别。通过了解稳定性级别，应用程序开发者可以根据服务或对象的组件是否存在或者它们是否可能跨发行版边界仍能正常使用，确定功能开发的可能性。</p> <p>在 smf(5) 上下文中，稳定性值还可以跨服务的发行版边界（可包含该服务的修补程序）标识属性组中属性的预期更改范围。以下两部分对这一点进行了更详细地讨论。</p>
属性组删除	<p>service_bundle(4) 文档类型定义包含一个删除属性，该属性适用于服务清单中的每个属性组。如果该删除属性设置为 <code>true</code>，它会指示 svccfg(1M) 及其他清单导入工具从系统信息库中删除此属性组。如果该删除属性缺失，或者存在但设置为 <code>false</code>，将保留系统信息库中的属性组。</p> <p>不会删除声明为稳定或发展中的属性组。可跨任何发行版边界删除声明为不稳定的属性组。</p>
另请参见	svcs(1) 、 svcadm(1M) 、 svccfg(1M) 、 svc.startd(1M) 、 libscf(3LIB) 、 service_bundle(4) 、 attributes(5) 、 pkg(1)
附注	smf(5) 的现有版本不支持多个系统信息库。

引用名	smf_method – 方法的服务管理框架约定						
描述	<p>在服务管理框架smf(5) 中通过 <code>svc.startd(1M)</code> 管理的服务类由符合简单 <code>fork(2)-exec(2)</code> 模型的应用程序组成。<code>svc.startd(1M)</code> 主守护进程及其他重启程序支持 <code>fork(2)-exec(2)</code> 模型（可能带有附加功能）。<code>svc.startd(1M)</code> 守护进程及其他重启程序需要按照本手册页中所述的约定激活、处理或检查服务实例的方法。</p>						
调用格式	<p>方法调用的格式不受约定的控制。在某些情况下，方法调用可能由守护进程或提供服务的其他二进制可执行文件的直接调用组成。对于使用可执行脚本或其他间接可执行文件的情况，约定建议采用以下格式：</p> <pre>/path/to/method_executable abbr_method_name</pre> <p>该建议格式中使用的 <i>abbr_method_name</i> 是支持的方法，如 <code>start</code> 或 <code>stop</code>。在相关重启程序页中提供了重启程序支持的方法集。<code>svc.startd(1M)</code> 守护进程支持 <code>start</code>、<code>stop</code> 和 <code>refresh</code> 方法。</p> <p>除了本页中引用的方法之外，重启程序还可能会定义其他种类的方法。围绕此类扩展的约定将由重启程序定义，可能与此处给出的定义不同。</p>						
环境变量	<p>重启程序向方法提供了四个环境变量，这些变量用于确定调用方法的上下文。</p> <p>SMF_FMRI 为其调用方法的实例的服务故障管理资源标识符 (Fault Management Resource Identifier, FMRI)。</p> <p>SMF_METHOD 所调用方法的完整名称，如 <code>start</code> 或 <code>stop</code>。</p> <p>SMF_RESTARTER 调用方法的重启程序的服务 FMRI</p> <p>SMF_ZONENAME 正在其中运行方法的区域的名称。也可使用 <code>zonename(1)</code> 命令获取此名称。</p> <p>在方法调用任何持久性进程之前，应从环境中删除这些变量。在下述包含文件中，为使用 Bourne 兼容 shell 脚本编写服务方法的服务作者提供了便利 shell 函数 <code>smf_clear_env</code>。</p> <p>方法上下文可能会导致设置其他环境变量（如下文所述）。</p>						
方法定义	<p>方法至少由 <code>method</code> 类型的属性组中的三个属性定义。</p> <p>这些属性包括：</p> <table><tr><td><code>exec (astring)</code></td><td>方法可执行字符串。</td></tr><tr><td><code>timeout_seconds (count)</code></td><td>方法超时之前所经过的秒数。有关更多详细信息，参见超时部分。</td></tr><tr><td><code>type (astring)</code></td><td>方法类型。当前始终设置为 <code>method</code>。</td></tr></table>	<code>exec (astring)</code>	方法可执行字符串。	<code>timeout_seconds (count)</code>	方法超时之前所经过的秒数。有关更多详细信息，参见 超时部分 。	<code>type (astring)</code>	方法类型。当前始终设置为 <code>method</code> 。
<code>exec (astring)</code>	方法可执行字符串。						
<code>timeout_seconds (count)</code>	方法超时之前所经过的秒数。有关更多详细信息，参见 超时部分 。						
<code>type (astring)</code>	方法类型。当前始终设置为 <code>method</code> 。						

要进一步完善方法的执行环境，可定义方法上下文。有关更多信息，参见方法上下文部分。

方法标记

当重启程序 `svc.startd` 在方法的 `exec` 字符串中定义一组标记时，将使用适当的值对这些标记进行分析和扩展。其他重启程序可能不支持方法标记。inet 服务 `inetd(1M)` 的委托重启程序不支持以下方法扩展。

`%%`
 `%`

`%r`
 重启程序的名称，如 `svc.startd`

`%m`
 所调用方法的完整名称，如 `start` 或 `stop`。

`%s`
 服务的名称

`%i`
 实例的名称

`%f`
 实例的 FMRI

`%{prop[: ,]}`
 属性的值。prop 可能是以 / 分隔的属性 FMRI、属性组名称和属性名，也可能是 application 属性组中的属性名。这些值可以后跟 , (逗号) 或 : (冒号)。使用分隔符 (如果提供) 分隔多个值。如果未提供分隔符，将使用空格。出现在字符串值中的以下 shell 元字符将用 \ (反斜杠) 引起来：
 ; & () | ^ < > newline space tab \ " ' ,
 无效扩展会导致方法失败。

有两个显式标记可用来代替方法命令。

`:kill [-signal]`
 将指定的信号 (缺省情况下为 SIGTERM) 发送到主实例合同中的所有进程。始终返回 SMF_EXIT_OK。此标记应当用于替换常见的 `pkill` 调用。

`:is`
 始终返回 SMF_EXIT_OK。此标记应当用于对重启程序是必需的但对特定服务是不必要的方法。

正在退出和退出状态

`start` 方法的必需行为是延迟退出，直到服务实例已经准备好应答请求或以其他方式工作。

以下退出状态代码在 `<libscf.h>` 和 shell 支持文件中定义。

SMF_EXIT_OK	0	方法已退出，已成功执行其操作。
SMF_EXIT_ERR_FATAL	95	方法出现致命性失败，必须通过管理干预才能恢复。
SMF_EXIT_ERR_CONFIG	96	无法恢复的配置错误。返回此退出状态的常见情况是已启用的服务实例缺少必需的配置文件。
SMF_EXIT_ERR_NOSMF	99	方法被错误地在 smf(5) 工具外部调用。依赖于 smf(5) 功能的服务应退出并显示此状态值。
SMF_EXIT_ERR_PERM	100	方法所需的某种形式的权限（如文件访问、特权、授权或其他凭证）在调用时不可用。
SMF_EXIT_ERR_OTHER	non-zero	方法的任何非零退出状态均被视为未知错误。一系列未知错误可能会被重启程序或以重启程序的名义识别为故障。

利用精确的退出代码，使负责的重启程序可以将错误响应分类为间歇性的（值得执行重新启动）或永久性的（要求管理干预）。

超时

每个方法都可以有以秒为单位指定的独立超时。特定超时的选择应基于对检测因不可响应性而导致的方法故障的站点预期。具有复制文件系统或其他故障转移资源的站点可选择延长缺省的方法超时。没有远程资源的站点也可以选择缩短超时。方法超时由 `timeout_seconds` 属性指定。

如果为方法指定 `0 timeout_seconds`，该属性将向重启程序声明服务没有超时。此设置不可取，但可用于确实需要它的服务。

也接受 `-1 timeout_seconds`，但此指定已过时。

Shell 编程支持

定义上述退出状态值的一组环境变量通过文件 `/lib/svc/share/smf_include.sh` 中的便利 shell 函数提供。此文件是 Bourne shell 脚本，可通过源运算符包含在任何 Bourne 兼容 shell 中。

为帮助编写可充当 SMF 方法的脚本以及 `/etc/init.d` 脚本，提供了 `smf_present()` shell 函数。如果 [smf\(5\)](#) 工具不可用，`smf_present()` 会返回非零退出状态。

此类脚本的一个可能的结构如下所示：

```
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
```

此示例显示了所提供的两个便利函数的用法。

方法上下文

服务管理工具提供了一种常见机制，可用于设置在其中执行 `fork(2)-exec(2)` 模型服务的上下文。

所需的方法上下文应由服务开发者提供。所有服务实例都应尽可能地以最低权限级别运行，以限制潜在的安全危害。

方法上下文可包含以下属性：

use_profile

一个布尔值，指定是否应使用配置文件来代替 `user`、`group`、`privileges` 和 `limit_privileges` 属性。

环境

要以一种包含多个 `NAME=value` 字符串的形式插入方法所在环境的环境变量。

profile

RBAC（Role-Based Access Control，基于角色的访问控制）配置文件的名称，该配置文件与方法可执行文件一起用于标识 `exec_attr(4)` 中的条目。

user

数字或文本形式的用户 ID。

group

数字或文本形式的组 ID。

supp_groups

一个可选字符串，按数字或文本形式的 ID 指定补充组成员资格。

privileges

一个可选字符串，指定 `privileges(5)` 中所定义的权限集。

limit_privileges

一个可选字符串，指定 `privileges(5)` 中所定义的限制权限集。

working_directory

从中启动方法的起始目录。可将 `:home` 用作一个标记，指示使用其 `uid` 启动方法的用户起始目录。如果未设置该属性，将使用 `:home`。

corefile_pattern

一个可选字符串，按照 `coreadm(1M)` 指定要用于服务的核心文件模式。多数重启程序都提供了缺省值。设置此属性会覆盖全局核心模式的本地定制。

project

数字或文本形式的项目 ID。可将 `:default` 用作一个标记，为使用其 `uid` 启动方法的用户指示由 `getdefaultproj(3PROJECT)` 标识的项目。

resource_pool

启动方法的资源池的名称。可将 `:default` 用作一个标记，指示在上述 `project` 属性中给出的 `project(4)` 条目中所指定的池。

通过为服务或实例指定 `method_context` 属性组，可为整个服务实例设置方法上下文。通过在方法属性组上提供方法上下文属性，方法可能会覆盖实例方法上下文。

无效方法上下文设置始终会导致方法失败，但导致发出警告的无效环境变量除外。

除了上述定义的上下文之外，许多 `fork(2)`-`exec(2)` 模型重启程序在作为方法调用可执行文件时还使用以下约定：

参数数组

`argv[]` 中参数的设置与 `exec` 字符串的结果 `/bin/sh -c` 一致。

文件描述符

文件描述符 0 为 `/dev/null`。文件描述符 1 和 2 建议为基于服务日志文件。

文件

`/lib/svc/share/smf_include.sh`

退出状态值的定义。

`/usr/include/libscf.h`

退出状态代码的定义。

另请参见

`zonename(1)`、`coreadm(1M)`、`inetd(1M)`、`svccfg(1M)`、`svc.startd(1M)`、`exec(2)`、`fork(2)`、`g`

附注

`smf(5)` 的现有版本不支持多个系统信息库。

当服务配置为以 `root` 身份但以不同于 `limit_privileges` 的权限启动时，所生成的进程可识别权限。这可能会出乎开发者的意料，他们本以为 `seteuid(<非零 UID>)` 能够将权限降为基本或更低的级别。

引用名	smf_restarter – 重启程序的服务管理工具约定	
描述	<p>服务管理工具中的所有服务实例都必须由重启程序来管理。本手册页介绍该框架中所有重启程序共有的配置、功能和报告特性。特定重启程序独有的特性将在重启程序的手册页中介绍。</p> <p>对于每个管理服务，重启程序依赖在服务实例上检索属性来确定配置。重启程序管理一组属性组，以此向显示工具（如 svcs(1)）传送服务的当前处置情况。</p>	
服务配置	所有服务的常见重启程序配置均在 <code>general</code> 属性组中捕获。此组包含以下必需和可选的属性设置。	
	<code>enabled</code>	此属性是必需的。如果设置此属性，实例的重启程序会尝试维护服务的可用性。
	<code>restarter</code>	此属性是可选的，允许指定备用重启程序管理服务实例。如果 <code>restarter</code> 属性为空或缺失，则重启程序缺省为 <code>svc.startd(1M)</code> 。
	<code>single_instance</code>	此属性是可选的。如果设置此属性，在任意时间将只允许一个服务实例转换为联机或降级状态。
服务报告	所有重启程序均使用 <code>restarter</code> 属性组报告状态，该属性组包含以下属性：	
	<code>next_state</code>	存储在这些属性中的实例的当前状态和下一状态（如果当前处于转换中）。有关可能的状态的说明，参见 smf(5) 。
	<code>auxiliary_state</code>	一个不带空格的 <code>astring</code> ，包含描述完整重启程序特定的状态以及重启程序状态属性的准确术语。不能一直设置辅助状态，在从任何状态进行转换期间应始终将其清除。每个重启程序都必须定义它所使用的辅助状态的精确列表。
	<code>state_timestamp</code>	访问当前状态的时间。
	<code>contract</code>	执行服务实例所依据的主进程合同 ID（如果有）。
另请参见	svcs(1) 、 svc.startd(1M) 、 service_bundle(4) 、 smf(5) 、 smf_method(5)	

引用名

smf_security – 服务管理工具的安全行为

描述

服务管理工具 [smf\(5\)](#) 的配置子系统需要修改服务配置的权限。通过 [user_attr\(4\)](#) 和 [prof_attr\(4\)](#) 将下文所述的授权与用户相关联，可以向用户授予适当的权限。参见 [rbac\(5\)](#)。

以下授权用于处理服务和服务实例。

<code>solaris.smf.modify</code>	有权添加、删除或修改服务、服务实例或其属性以及读取受保护的属性值。
---------------------------------	-----------------------------------

属性组授权

[smf\(5\)](#) 配置子系统将属性与每个服务和服务实例相关联。相关属性将进行分组。组可以表示执行方法、凭证信息、应用程序数据或重启程序状态。创建或修改属性组的能力可能会导致 [smf\(5\)](#) 组件执行可能需要系统权限的操作。相应地，框架也需要处理属性组的适当权限。

每个属性组都有与其用途相对应的类型。核心属性组类型有 `method`、`dependency`、`application` 和 `framework`。也可以引入其他属性组类型，但前提是它们必须符合 [smf\(5\)](#) 中的扩展命名约定。不过，以下基本授权仅适用于核心属性组类型：

<code>solaris.smf.modify.method</code>	有权更改值或者创建、删除或修改 <code>method</code> 类型的属性组。
<code>solaris.smf.modify.dependency</code>	有权更改值或者创建、删除或修改 <code>dependency</code> 类型的属性组。
<code>solaris.smf.modify.application</code>	有权更改值、读取受保护的属性值以及创建、删除或修改 <code>application</code> 类型的属性组。
<code>solaris.smf.modify.framework</code>	有权更改值或者创建、删除或修改 <code>framework</code> 类型的属性组。
<code>solaris.smf.modify</code>	有权添加、删除或修改服务、服务实例或其属性以及读取受保护的属性值。

属性组特定的授权可以由包含在该属性组中的属性来指定。

<code>modify_authorization</code>	授权允许在属性组中添加、删除或修改属性以及从受保护的属性组检索属性值。
<code>value_authorization</code>	授权允许更改属性组的任何属性值（ <code>modify_authorization</code> 除外）以及从受保护的属性组检索任何属性值（ <code>modify_authorization</code> 除外）。
<code>read_authorization</code>	授权允许在属性组中检索属性值。如果存在具有此名称的字符串值属性，则会将所在属性组标识为受保护。此属性不适用于非 <code>application</code> 类型的属性组。参见 受保护的属性组 。

上述授权属性仅当具有 `astr` 类型时才能使用。如果实例属性组没有其中任一属性，但实例的服务具有与该属性同名的属性组，则将使用该属性的值。

受保护的属性组 通常，任何用户无需显式授权即可读取系统信息库中的所有属性值。非 `framework` 类型的属性组可用于存储其值需要保护的属性。除非具有适当的授权，否则不能公开这些属性组。通过字符串值的 `read_authorization` 属性存在与否来指示属性组的受保护状态。如果此属性存在，只能按照**属性组授权**中的说明来检索属性组中的所有属性值。

如果管理域所使用的策略禁止备份被视为敏感的数据，应从其备份中排除 SMF 系统信息库数据库。即使存在此类策略，也可备份未受保护的属性值，方法是使用 `svccfg(1M)` 归档命令为不包含受保护属性值的系统信息库创建归档文件。

服务操作授权 对服务实例执行的某些操作可能会导致服务中断或停用。这些操作需要授权才能确保任何拒绝服务操作都是有意的管理操作。此类操作包含执行 `refresh` 或 `restart` 方法的请求或将服务实例置于维护或其他非运行状态的请求。以下授权允许请求此类操作：

`solaris.smf.manage` 有权请求任何服务实例的重新启动、刷新或其他状态修改。

此外，`general/action_authorization` 属性还可以指定允许对该服务实例请求服务操作的其他授权。需要 `solaris.smf.manage` 授权才能修改此属性。

定义的权限配置文件 包含了两个权限配置文件，它们可以提供用于处理典型 `smf(5)` 操作的分组授权。

服务管理 服务管理员可以通过任意方式处理系统信息库中的任何服务。服务管理员与 `solaris.smf.manage` 和 `solaris.smf.modify` 授权相对应。

要使用 `pkg(1)` 命令添加或删除其服务清单文件中包含服务清单的软件包，至少需要服务管理配置文件。

服务操作员 服务操作员能够在系统上启用或禁用任何服务实例以及请求执行其 `restart` 或 `refresh` 方法。服务操作员与 `solaris.smf.manage` 和 `solaris.smf.modify.framework` 授权相对应。

站点可根据需要定义其他权限配置文件。

远程系统信息库修改 远程系统信息库服务器可能会由于附加权限检查而拒绝修改尝试。请参见“注释”。

另请参见 `auths(1)`、`profiles(1)`、`svccfg(1M)`、`prof_attr(4)`、`user_attr(4)`、`rbac(5)`、`smf(5)`、`pkg(1)`

附注 `smf(5)` 的现有版本不支持远程系统信息库。

当服务配置为以 `root` 身份但以不同于 `limit_privileges` 的权限启动时，所生成的进程可识别权限。这可能会出乎开发者的意料，他们本以为 `seteuid(<非零 UID>)` 能够将权限降为基本或更低的级别。

引用名	smf_template – 对服务元数据的服务管理框架支持
描述	<p>模板由服务开发者定义，用于在整体上描述关于服务的元数据或描述服务的各个配置属性，包括便于阅读的说明以及有效配置的定义。</p> <p>通过 SMF 命令为管理员提供访问模板的权限，这些命令用于描述配置值以及根据模板验证配置。</p> <p>工具开发者可以使用模板为服务配置提供更有帮助用户界面。</p>
模板数据	服务元数据在模板中定义为服务清单的一部分。
使用模板数据	<p>可以使用 <code>svcs -lv</code> 和 <code>svccfg describe</code> 命令以便于阅读的格式访问关于属性的元数据。</p> <p>可以使用 <code>svccfg(1M)</code> 的 <code>validate</code> 子命令根据模板数据验证服务实例或清单。可以使用一组 <code>libscf(3LIB)</code> 接口访问模板数据。</p>
模板定义	<p>服务清单是定义模板的唯一接口。</p> <p>服务作者应该为服务特定的属性组以及它们所引入的属性提供模板元数据，包括 <code>common_names</code>、<code>descriptions</code>、<code>choices</code> 和 <code>constraints</code>。在 C 语言环境中，服务作者必须至少为属性组和属性提供说明。服务作者不能为框架附带的属性组（如方法和相关项）提供模板元数据。</p> <p>有关为服务制作模板定义的示例，请参见 示例 部分。</p>
模板编写	<p>对属性组的相关模板数据的所有模板接口搜索依次在实例、服务、服务的重启程序上执行，最后将全局执行。</p> <p>属性组模板由其作者定义，可以应用于特定实例、服务及其所有实例、重启程序的委托或者全局应用。典型服务作者在实例或服务上定义模板。在实例上定义的模板仅应用于该实例，可覆盖在服务上定义的该属性组的模板。在服务上定义的模板应用于该服务的所有实例。</p> <p>重启程序作者可以在其清单中定义模板，这些模板应用于使用其重启程序的任何服务（也称为委托）。SMF 框架作者在 <code>svc:/system/svc/global</code> 的清单中为整个 SMF 框架熟知其含义的属性组定义了模板。</p> <p>如果模板已在全局定义或由重启程序定义，但又由服务或实例重新定义，将被标记为验证错误。服务作者可以只为其服务特定的且 SMF 框架未使用的属性组创建模板，从而避免这些错误。</p> <p>属性组模板也可以按名称或类型设置通配符。只有适用于属性组的最特定的模板定义才受支持。</p>

模板详细信息

服务和实例模板

`template` 元素定义模板块的开始。以下所有更深入的定义可包含在模板块中。`template` 元素可包含在 `service` 或 `instance` 元素中。如果它包含在 `service` 元素中，则应用于服务及该服务的所有实例。如果它包含在 `instance` 元素中，则仅应用于服务的该实例。

我们建议您尽可能地为整个服务定义模板数据。

```
<service ... >
  <template>
  </template>
</service>
```

服务和实例通用名称

整个服务或实例可以定义一个通用名称，用来描述服务/实例的用途。

```
<template>
  <common_name>
    <loctext xml:lang='C'>console login</loctext>
  </common_name>
</template>
```

`common_name` 是一个自由格式的字符串，但专门用作 GUI 或 CLI 中的标签。

定义通用名称时，请遵循以下准则：

- 简短。通常一两个字词就很合适。将名称限定为 40 个字符以内。
- 清晰。服务名、属性组名或属性名对用户可能没什么帮助，但 `common_name` 有助于阐明实体的用途。
- 无标点符号。`common_name` 不是句子或段落。它不应包含分句或短语。只有在为了满足商标需求时才应使用标点符号。
- 大写字母只能用于首字母缩略词或正确的名称。对于英语以外的语言环境，对句子分段使用适当的大写。

服务和实例说明

`description` 元素包含属性组的较长说明，适用于状态行或工具提示：

```
<template>
  <description>
    <loctext xml:lang='C'>Provide the text login prompt on console.
    </loctext>
  </description>
</template>
```

description 准则

- 使用正确的语法。*description* 是旨在供用户阅读的句子。
- 简短。几个句子通常最合适。

文档

可明确地定义此服务的文档，这样在服务出现问题或者服务的使用者需要更多相关信息时，就可以轻松地查找文档。

属性组

`pg_pattern` 元素包含属性组的定义：

```
<template>
  <pg_pattern name="pgname" type="pgtype" target="this" required="true">
  </pg_pattern>
</template>
```

name 是属性组的名称，*type* 是属性组的类型。

target 指定此定义的目标。"this" 是指定义服务或实例。"instance" 只能在服务的模板块中使用，这意味着该定义应用于此服务的所有实例。"delegate" 只能在重启程序的模板块中使用，并且应用于委托给该重启程序的所有实例。"all" 只能由主重启程序使用，是指系统上的所有服务。目标的缺省值为 "this"。

required 指示此属性组是否为必需。*required* 的缺省值为 `false`。如果 *required* 为 `true`，则必须指定 *name* 和 *type*。

name 和/或 *type* 可忽略。如果这两个属性中的任何一个被忽略，则将其视为通配符。例如，如果在 `pg_pattern` 定义中忽略 *name* 属性，则 `pg_pattern` 将适用于具有指定类型的所有特性组。

属性组名称

`common_name` 元素包含属性组的便于阅读的本地化名称：

```
<pg_pattern ...>
  <common_name>
    <loctext xml:lang='C'>start method</loctext>
  </common_name>
</pg_pattern>
```

`common_name` 是一个自由格式的字符串，但专门用作 GUI 或 CLI 中的标签。

参见上文“服务实例和通用名称”中的 `common_name` 准则。

属性组说明

`description` 元素包含属性组的较长说明，适用于状态行或工具提示：

```
<pg_pattern ...>
  <description>
    <loctext xml:lang='C'>A required method which starts the service.
  </loctext>
</description>
</pg_pattern>
```

参见上文“服务和实例说明”中用于指定 *description* 的准则。

属性

prop_pattern 元素包含特定属性的定义：

```
<pg_pattern ...>
  <prop_pattern name="propname" type="proptype" required="true">
</prop_pattern>
</pg_pattern>
```

name 是属性的名称，*type* 是属性的类型。

required 指示属性是否为必需。*required* 的缺省值为 *false*。

name 始终为必需。*type* 只有在 *required* 为 *false* 时才可选。

属性名

common_name 元素包含属性的便于阅读的本地化名称：

common_name 是一个自由格式的字符串字段，但专门用作 GUI 或 CLI 中的标签。

```
<prop_pattern ...>
  <common_name>
    <loctext xml:lang='C'>retry interval</loctext>
  </common_name>
</prop_pattern>
```

参见上文“服务实例和通用名称”中的 *common_name* 准则。

属性单位

units 元素包含数字属性的便于阅读的本地化单位：

```
<prop_pattern ...>
  <units>
    <loctext xml:lang='C'>seconds</loctext>
  </units>
</prop_pattern>
```

units 准则

- 简短。尽量只使用单个字词或标签。复数形式通常最合适。

- 无标点符号。*units* 不是句子或段落。它不应包含分句或短语。只有在为了满足商标需求时才应使用标点符号。

属性说明

description 元素包含属性的较长说明，适用于状态行或工具提示：

```
<prop_pattern ...>
  <description> <loctext xml:lang='C'>
    The number of seconds to wait before retry.
  </loctext> </description>
</prop_pattern>
```

参见上文“服务和实例说明”中用于指定 *description* 的准则。

属性可见性

visibility 元素指定更高级别软件中的简化视图是否要显示此属性。

```
<prop_pattern ...>
  <visibility value="hidden | readonly | readwrite"/>
</prop_pattern>
```

有些属性是内部实现的详细信息，不应显示为配置设置。其他属性可能仅为只读。此属性用于指定这些限制。值 *hidden* 指示不应显示属性，*readonly* 指示属性不能修改，*readwrite* 指示属性可修改。

此属性不是一种安全机制，它只是专门帮助用户避免因粗心大意而造成破坏以及从 CLI 输出或 GUI 显示中删除不必要的杂乱信息。在许多命令和 UI 中的完全公开模式下，隐藏属性是可见的。

属性格式

cardinality 和 *internal_separators* 元素用于限制属性的结构：

```
<prop_pattern ...>
  <cardinality min="1" max="1"/>
  <internal_separators>,</internal_separators>
</prop_pattern>
```

cardinality 指示可接受的属性值数量。*min* 是最小数量，*max* 是最大数量。这两者均可选。如果这两者均未指定，则 *<cardinality/>* 等同于缺省数量的值、零个值或零个以上的值。

internal_separators 指定在打包了多个实际值的属性值中使用的分隔符。

值约束

constraints 元素指定属性可接受哪些值：

```

<prop_pattern ...>
<constraints>
  <value name="blue" />
  <range min="1" max="7"/>
  <include_values type="values"/>
</constraints>
</prop_pattern>

```

value 元素包含可能的属性值。*range* 包含整数范围。

value 和 *range* 可以按任何组合形式使用，限制其使用将会禁止许多有效的说明。如果未指定值约束，则属性可采用任何值。

include_values 包含由值块指定的所有值（参见**值说明**部分）。

值选择

选择块指示 UI 应向用户提供哪些值：

```

<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>

```

与用于约束时一样，*range* 和 *value* 也分别包含范围和各个值。

include_values 包含由约束块或值块（参见下一部分）指定的所有值。

值说明

与属性名一样，属性可采用的值也可能具有难以理解的表示形式。*values* 元素包含特定属性值的便于阅读的本地化说明：

```

<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 是一个自由格式的字符串，但专门用作 GUI 或 CLI 中的标签。

参见上文“服务实例和通用名称”中的 *common_name* 准则。

示例

假定要定义基本模板数据的基本服务如下所示：

```
<?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 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>
```

该服务可定义一些基本模板数据，从而为在 <service> 标记内使用此服务的管理员提供帮助。最有用的就是记录服务本身的用途以及服务特定的配置。

```
<template>
  <common_name> <loctext xml:lang='C'>
    all-purpose demonstration
  </loctext> </common_name>
  <documentation>
    <manpage title='foo' 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='astrin'
      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='astrin'
      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>
</pg_pattern>
</template>
```

文件

/usr/share/lib/xml/dtd/service_bundle.dtd.1

另请参见

[svcs\(1\)](#)、[svccfg\(1M\)](#)、[libscf\(3LIB\)](#)、[service_bundle\(4\)](#)、[smf\(5\)](#)

引用名 solaris10 – Solaris 10 标记区域

描述 solaris10 标记使用 [brands\(5\)](#) 中介绍的标记区域框架，使得 Solaris 10 二进制应用程序无需经过修改即可在具有最新 Solaris 操作系统内核的计算机上运行。

solaris10 标记只支持在区域中运行 Solaris 10 10/09（Solaris 10 更新 8）或更高发行版。

solaris10 标记包括将 Solaris 10 系统映像安装到非全局区域所必需的工具。它还支持将 Solaris 10 本地区域迁移到 solaris10 标记区域所必需的工具。该标记支持在运行最新 Solaris 操作系统的 SPARC 或 x86 计算机上执行 32 位和 64 位 Solaris 10 应用程序。

配置与管理 solaris10 标记支持完全根非全局区域模型。所有必需的 Solaris 10 软件及任何其他软件包都安装在区域的专用文件系统中。

区域必须位于各自的 [zfs\(1M\)](#) 数据集上，并且只支持 ZFS。在安装或连接区域时，将自动创建 ZFS 数据集。如果无法创建 ZFS 数据集，区域也无法安装或连接。

[zonecfg\(1M\)](#) 实用程序可用于配置 solaris10 标记区域。在创建区域时可以使用 SUNWsolaris10 模板，也可以手动设置配置。一旦安装了标记区域，便不能更改或删除该区域的标记。[zoneadm\(1M\)](#) 实用程序可用于报告区域的标记类型和管理区域。[zlogin\(1\)](#) 实用程序可用于登录到区域。

当前对于委托 ZFS 数据集配置的支持是试验性的，尚未经过测试。对于在超虚拟化 xVM 域中运行这些区域的支持是试验性的，并且在区域中运行的 64 位 x86 应用程序存在已知问题。不能在标记区域中配置 `/dev/sound` 设备。此外，在全局区域中用于检查 solaris10 标记区域中执行的进程时，[mdb\(1\)](#) 和 [dtrace\(1M\)](#) 无法完全正常运行。

solaris10 标记安装程序支持通过已安装的 Solaris 10 系统的映像安装区域。这可以是完整的 [flash_archive\(4\)](#)、[cpio\(1\)](#) 或 [pax\(1\)](#) xustar 归档文件。[cpio](#) 归档文件可以使用 [gzip\(1\)](#) 或 [bzip2\(1\)](#) 进行压缩。此映像也可以是级别为 0 的 [ufsdump\(1M\)](#)，或者是到 Solaris 10 系统根目录树顶级的路径。不能从标准的 Solaris 10 分发介质中安装区域。

为了将本地区域从 Solaris 10 系统迁移到最新的 Solaris 操作系统内核，`attach` 子命令支持从已安装的 Solaris 10 本地区域的归档文件中安装区域。与安装程序一样，这可以是 `zonepath` 的 [cpio\(1\)](#) 或 [pax\(1\)](#) xustar 归档文件。[cpio](#) 归档文件可以使用 [gzip\(1\)](#) 或 [bzip2\(1\)](#) 进行压缩。映像还可以是到 Solaris 10 区域的 `zonepath` 目录树顶级的路径。除了从 Solaris 10 本地区域迁移外，相同的迁移选项也可以在将 solaris10 标记区域从一台主机迁移到另一台主机时使用。从 Solaris 10 进行迁移时，可以将区域配置为 `sparsEROOT` 区域。这种情况下，在生成归档文件前，应该在主机上已配置好区域。这样可以确保在归档文件中包含继承的目录。

子命令 支持 [zoneadm\(1M\)](#) 标记特定子命令的下列参数：

`attach [-a archive] [-d path] [-c sysidecfg]`

将指定的 Solaris 10 本地区域映像连接到标记区域。如果 `-a` 或 `-d` 都未指定，则假定已使用区域文件正确安装区域的 `zonepath`。

- a *archive* 到已安装的 Solaris 10 本地区域或 **solaris10** 标记区域的 **zonepath** 的 **cpio(1)**、**pax(1)** xustar 或 zfs 归档文件的路径。cpio 和 zfs 归档文件可以使用 **gzip** 或 **bzip2** 进行压缩。
- c *sysidcfg* 指定 *sysidcfg* 文件将导致连接后在区域上发生 **sys-unconfig**。然后 *sysidcfg* 文件应用到区域。
- d *path* 到已安装的 Solaris 10 本地区域或 **solaris10** 标记区域的 **zonepath** 的 **zonepath** 目录的路径。

clone [-c *sysidcfg*]

安装某个区域时，可直接复制现有的已安装区域。该子命令是一种备选的区域安装方法。

- c *sysidcfg* 提供取消配置克隆的区域后将应用的 *sysidcfg* 文件。

install [-a *archive*] [-d *path*] [-p] [-s] [-u] [-v] \

[-c *sysidcfg*]

将指定的 Solaris 10 系统映像安装到区域中。要求使用 -u 或 -p 选项以及 -a 或 -d 选项。

-a *archive*

到已安装的 Solaris 10 系统的 **flash_archive(4)**、**cpio(1)**、**pax(1)** xustar、zfs 归档文件或级别 **0** **ufsdump(1M)** 的路径。cpio 和 zfs 归档文件可以使用 **gzip** 或 **bzip2** 进行压缩。

-c *sysidcfg*

提供安装后将应用的 *sysidcfg* 文件。

-d *path*

到已安装的 Solaris 10 系统的根目录的路径。

-p

在安装区域后保留系统配置。

-s

无提示安装。

-u

安装区域后对区域运行 **sys-unconfig**。

-v

安装过程的详细输出。

应用程序支持

solaris10 区域只支持用户级的 Solaris10 应用程序。不能使用 **solaris10** 区域中的 Solaris 10 设备驱动程序或 Solaris 10 内核模块。不过，取决于内核模块，您可能能够将最新的 Solaris 内核模块版本用于 Solaris 10 用户级应用程序。

属性类型	属性值
可用性	system/zones/brand/brand-solaris10
接口稳定性	Committed（已确定）

cpio(1)、mdb(1)、pax(1)、zlogin(1)、dtrace(1M)、ufsdump(1M)、zfs(1M)、zoneadm(1M)、zonecfg(1M)

引用名	solaris, ipkg – Solaris 标记区域
描述	solaris 标记使用标记区域框架（如 brands(5) 中所述）运行与安装在全局区域中的同一软件一起安装的区域。使用 solaris 标记时，系统软件必须始终与全局区域保持同步。区域中的系统软件包使用映像包管理系统进行管理。请参见 pkg(5) 。
配置与管理	solaris 标记支持完全根非全局区域模型。所有必需的系统软件及任何其他软件包都安装在区域的专用文件系统中。区域只能位于其自己的 zfs(1M) 数据集中并且仅支持 ZFS。在安装或连接区域时，将自动创建 ZFS 数据集。如果无法创建 ZFS 数据集，区域也无法安装或连接。
子命令	<p>以下 solaris 标记特定的子命令选项受 zoneadm(1M) 支持。</p> <pre>attach [-a archive] [-d path] [-z ZBE] [-u] [-c config_profile.xml dir]</pre> <p>将指定的 solaris 标记区域映像连接到区域。如果 -a 和 -d 均未指定，则假定区域的区域路径已经与区域文件一起正确安装。zoneadm 会在计算机上检查区域要连接到的包级别。如果区域通过全局区域所依赖的包（具有不同的版本号）不同于源计算机上的相关包，zoneadm 将报告这些冲突并且不执行连接。</p> <p>如果目标系统只有比源系统更新的相关软件包（修订版本号较高），则可以使用 -u 或 -u 选项更新相关软件包以匹配位于新系统上的软件包修订版。</p> <p>连接区域时，可能会存在多个区域引导环境 (zone boot environment, ZBE)，attach 子命令必须确定要连接其中哪一个。选择条件如下所示，并使用第一个匹配项。</p> <ul style="list-style-type: none"> ■ 如果使用 -z 选项来指定 ZBE，则选择它。 ■ 如果只有一个 ZBE，则选择它。 ■ 如果只有一个 ZBE 与该全局区域引导环境关联，则选择它。 ■ 如果只有一个活动的 ZBE 与该全局区域引导环境关联，则选择它。 ■ 如果只有一个活动的 ZBE，则选择它。 <p>如果选定的 ZBE 与其他全局区域引导环境关联，则 ZBE 将被克隆，并且连接选定 ZBE 的克隆。有关引导环境的更多信息，请参见 beadm(1M)。</p> <p>-a archive 到已安装 Oracle Solaris 标记区域的 cpio(1) 或 pax(1) xustar 或 zfs(1M) 归档文件的路径。</p> <p>cpio 和 pax 归档文件必须采用区域路径或区域根路径，并且不得包含绝对路径（路径不得以 / 开头）。ZFS 归档文件可以是 zonepath dataset、zone rpool dataset 或单独区域引导环境。</p> <p>如果 ZFS 归档文件包含多个 ZBE，则连接可能失败，并显示一条消息，指示必须使用 -z 选项连接特定的 ZBE。在这种情况下，仍然提取归档文件中的 ZBE，并且该区域将无法使用 -a 和 -d 选项，直到使用 -z 选项连接提取的一个 ZBE。如果不希望连接提取的一个 ZBE，而是删除 ZBE，则使用 zoneadm -z <zone> mark incomplete，后跟 zoneadm -z <zone> uninstall。</p> <p>cpio 和 ZFS 归档文件可以使用 gzip 或 bzip2 进行压缩。</p>

-c *config_profile.xml* | *dir*

提供从系统信息库安装后将应用的配置文件或配置文件的目录。

所有配置文件必须具有 .xml 扩展名。

-d *path*

区域路径目录的路径，该目录属于 ipkg 标记区域的区域路径。

-u

更新区域中的最少数量的软件包，以允许区域的软件包与在全局区域中安装的软件包兼容。

-U

将区域中的所有软件包更新为与安装在全局区域中的软件包兼容的最新版本。

-z *ZBE*

连接指定的现有区域引导环境。如果指定的区域引导环境与其他全局区域关联，则指定的 ZBE 将被克隆，并连接 ZBE 的克隆。

clone [-c *config_profile.xml* | *dir*]

-c *config_profile.xml* | *dir* 提供从系统信息库安装后将应用的配置文件或配置文件的目录。

所有配置文件必须具有 .xml 扩展名。

install [-m *manifest.xml*] [-c *config_profile.xml* | *dir*]

install [-a *archive* | -d *path* | -z *ZBE*] [-p] [-s] [-u] [-v] [-c *config_profile.xml* | *dir*]

solaris 标记安装程序支持从软件信息库或从运行同一发行版的已安装系统映像安装区域。此映像可以是 **cpio(1)**、**pax(1)** xustar 或 ZFS 归档文件。cpio 或 ZFS 归档文件可以使用 gzip 或 bzip2 进行压缩。该映像也可以是到系统根目录树顶级的路径，或者是预存在的区域路径。

如果 -a、-d 或 -z 选项都未指定，则从系统信息库安装区域。要在区域中安装其他软件包，可以复制并编辑缺省区域清单

/usr/share/auto_install/manifest/zone_default.xml 以包含所需的软件包。应使用 -m 选项指定安装该修改的清单。

要从系统映像安装区域，需要使用 -a、-d、-z 选项之一。在此例中还需要使用 -u 或 -p 选项。

-a *archive*

到已安装系统的 **cpio(1)** 或 **pax(1)** xustar 或 ZFS 归档文件的路径。

如果 ZFS 归档文件包含多个引导环境，则安装了活动的引导环境。如果安装无法确定哪个引导环境是活动的引导环境，安装程序将提供提取的引导环境列表，并建议一个使用 -z 选项的连接命令来连接特定的引导环境。

	cpio 和 ZFS 归档文件可以使用 gzip 或 bzip2 进行压缩。
<code>-c config_profile.xml dir</code>	提供从系统信息库安装后将应用的配置文件或配置文件的目录。
	所有配置文件必须具有 .xml 扩展名。
<code>-d path</code>	区域路径目录的路径，该目录属于 ipkg 标记区域的区域路径。
<code>-m manifest.xml</code>	将指定给自动化安装程序的清单文件。
<code>-p</code>	从归档文件或路径安装区域后，保留系统配置。
<code>-s</code>	无提示安装
<code>-u</code>	安装后取消配置系统。
<code>-v</code>	安装过程的详细输出。
<code>-z ZBE</code>	连接指定的现有区域引导环境。如果指定的区域引导环境与其他全局区域关联，则指定的 ZBE 将被克隆，并连接 ZBE 的克隆。

示例

示例 1 创建 ZFS 归档文件进行安装

以下示例显示如何创建归档文件，以便进行物理到虚拟 (P2V) 迁移。这将在未配置、安装或运行非全局区域的系统的全局区域中执行。它假设根池名为 `rpool`。

首先，创建整个根池的快照。

```
# zfs snapshot -r rpool@p2v
```

然后，销毁与交换设备和转储设备关联的快照，因为在目标系统上不需要它们。

```
# zfs destroy rpool/swap@p2v
# zfs destroy rpool/dump@p2v
```

最后，生成使用 gzip 压缩的 ZFS 复制流归档文件。在本示例中，它存储在远程 NFS 服务器上。

```
# zfs send -R rpool@p2v | gzip > /net/somehost/p2v/p2v.zfs.gz
```

示例 2 使用 ZFS 归档文件安装区域

以下示例使用示例 1 中的归档文件安装区域。

```
# zoneadm -z p2vzone install -a /net/somehost/p2v/p2v.zfs.gz -p
```

示例3 创建 ZFS 归档文件进行连接

以下示例显示如何创建归档文件，以便进行虚拟到虚拟 (v2v) 迁移。它假设该区域的 `zonepath` 为 `/zones/v2vzone`。

首先，确定 `zonepath` 数据集的名称。

```
# dataset=$(zfs list -H -o name /zones/v2vzone)
```

然后，创建区域数据集的快照。

```
# zfs snapshot -r $dataset@v2v
```

最后，生成使用 `bzip2` 压缩的 ZFS 自包含递归流。

```
# zfs send -rc $dataset@v2v | bzip2 > /net/somehost/v2v/v2v.zfs.bz2
```

示例4 使用 ZFS 归档文件连接区域

以下示例使用 ZFS 归档文件连接区域。

```
# zoneadm -z v2vzone attach -a /net/somehost/v2v/v2v.zfs.bz2
```

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/zones
接口稳定性	Uncommitted（未确定）

另请参见

[cpio\(1\)](#)、[pax\(1\)](#)、[beadm\(1M\)](#)、[sysconfig\(1M\)](#)、[zfs\(1M\)](#)、[zlogin\(1\)](#)、[zonename\(1\)](#)、[zoneadm\(1M\)](#)、[pkg\(5\)](#)，在 IPS 整合中提供

引用名

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

描述

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
POSIX.1-2008	POSIX.1-2004 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

X/Open CAE		
Specification	Description	Release
XNS4	sockets and XTI interfaces	Solaris 2.6
SUSv2	superset of SUS extended to support POSIX.1b-1993, POSIX.1c-1996, and ISO/IEC 9899 (C Standard) Amendment 1	Solaris 7
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 12.2 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 12.2 C Compiler is installed, Oracle Solaris 11 supports ISO/IEC 9899:1990 Amendment 1:1995: C Integrity.

When Oracle Solaris Studio 12.2 C Compiler is installed, Oracle Solaris 11 supports ISO/IEC 9899:1999 Programming Languages – C.

When Oracle Solaris Studio 12.2 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 `exec1p()` 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 <code>_XOPEN_SOURCE</code> . If <code>_XOPEN_SOURCE</code> 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 12.2 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 12.2 C Compiler and the Oracle Solaris Studio 12.2 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 [1f64\(5\)](#) manual page, since these names are used by the implementation to name the alternative entry points.

When using the Oracle Solaris Studio 12.2 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	cc -Xt -xc99=none	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) -lnet
```

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) -lnet
```

SUSv3

```
c99                _XOPEN_SOURCE=600
```

另请参见

[csh\(1\)](#), [ksh\(1\)](#), [sh\(1\)](#), [exec\(2\)](#), [sysconf\(3C\)](#), [system\(3C\)](#), [environ\(5\)](#), [lf64\(5\)](#)

引用名	sticky – mark files for special treatment
描述	<p>The <i>sticky bit</i> (file mode bit <code>01000</code>, 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 <code>/tmp</code>, which must be publicly writable but should deny users permission to arbitrarily delete or rename the files of others.</p> <p>If the sticky bit is set on a regular file and no execute bits are set, the system's page cache will not be used to hold the file's data. This bit is normally set on swap files of diskless clients so that accesses to these files do not flush more valuable data from the system's cache. Moreover, by default such files are treated as swap files, whose inode modification times may not necessarily be correctly recorded on permanent storage.</p> <p>Any user may create a sticky directory. See <code>chmod</code> for details about modifying file modes.</p>
另请参见	chmod(1) , chmod(2) , chown(2) , mkdir(2) , rename(2) , unlink(2)
已知问题	The mkdir(2) function will not create a directory with the sticky bit set.

引用名	tecla, teclarc – User interface provided by the tecla library.														
描述	<p>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 <code>tcsh</code> shell will find the default key bindings very familiar. Users of the <code>bash</code> 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 <code>vi(1)</code> 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.</p>														
Key Sequence Notation	<p>In the rest of this man page, and also in all tecla configuration files, key sequences are expressed as follows.</p> <table><tr><td><code>^A</code> or <code>C-a</code></td><td>This is a 'CONTROL-A', entered by pressing the CONTROL key at the same time as the 'A' key.</td></tr><tr><td><code>\\E</code> or <code>M-</code></td><td>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 <code>M-p</code> 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'.</td></tr><tr><td><code>up</code></td><td>This refers to the up-arrow key.</td></tr><tr><td><code>down</code></td><td>This refers to the down-arrow key.</td></tr><tr><td><code>left</code></td><td>This refers to the left-arrow key.</td></tr><tr><td><code>right</code></td><td>This refers to the right-arrow key.</td></tr><tr><td><code>a</code></td><td>This is just a normal 'A' key.</td></tr></table>	<code>^A</code> or <code>C-a</code>	This is a 'CONTROL-A', entered by pressing the CONTROL key at the same time as the 'A' key.	<code>\\E</code> or <code>M-</code>	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 <code>M-p</code> 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'.	<code>up</code>	This refers to the up-arrow key.	<code>down</code>	This refers to the down-arrow key.	<code>left</code>	This refers to the left-arrow key.	<code>right</code>	This refers to the right-arrow key.	<code>a</code>	This is just a normal 'A' key.
<code>^A</code> or <code>C-a</code>	This is a 'CONTROL-A', entered by pressing the CONTROL key at the same time as the 'A' key.														
<code>\\E</code> or <code>M-</code>	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 <code>M-p</code> 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'.														
<code>up</code>	This refers to the up-arrow key.														
<code>down</code>	This refers to the down-arrow key.														
<code>left</code>	This refers to the left-arrow key.														
<code>right</code>	This refers to the right-arrow key.														
<code>a</code>	This is just a normal 'A' key.														
The Tecla Configuration File	<p>By default, tecla looks for a file called <code>.teclarc</code> in your home directory (ie. <code>~/ .teclarc</code>). 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:</p> <pre>edit-mode vi</pre> <p>This will re-configure the default bindings for vi-mode. The complete set of arguments that this command accepts are:</p> <table><tr><td><code>vi</code></td><td>Install key bindings like those of the vi editor.</td></tr><tr><td><code>emacs</code></td><td>Install key bindings like those of the emacs editor. This is the default.</td></tr></table>	<code>vi</code>	Install key bindings like those of the vi editor.	<code>emacs</code>	Install key bindings like those of the emacs editor. This is the default.										
<code>vi</code>	Install key bindings like those of the vi editor.														
<code>emacs</code>	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> .
<code>forward-delete-word</code>	Delete the word which follows the cursor.
<code>backward-delete-word</code>	Delete the word which precedes the cursor.
<code>upcase-word</code>	Convert all of the characters of the word which follows the cursor, to upper case.
<code>downcase-word</code>	Convert all of the characters of the word which follows the cursor, to lower case.
<code>capitalize-word</code>	Capitalize the word which follows the cursor.
<code>change-case</code>	If the next character is upper case, toggle it to lower case and vice versa.
<code>redisplay</code>	Redisplay the line.
<code>clear-screen</code>	Clear the terminal, then redisplay the current line.
<code>transpose-chars</code>	Swap the character under the cursor with the character just before the cursor.
<code>set-mark</code>	Set a mark at the position of the cursor.
<code>exchange-point-and-mark</code>	Move the cursor to the last mark that was set, and move the mark to where the cursor used to be.
<code>kill-region</code>	Delete the characters that lie between the last mark that was set, and the cursor.
<code>copy-region-as-kill</code>	Copy the text between the mark and the cursor to the cut buffer, without deleting the original text.
<code>yank</code>	Insert the text that was last deleted, just before the current position of the cursor.
<code>append-yank</code>	Paste the current contents of the cut buffer, after the cursor.
<code>up-history</code>	Recall the next oldest line that was entered. Note that in <code>vi</code> mode you are left in command mode.
<code>down-history</code>	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.
<code>history-search-backward</code>	Recall the next oldest 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-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>> 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
^\	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-<</code>	beginning-of-history
<code>M-></code>	end-of-history
<code>\</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-^\</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>\</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\	forward-copy-char
M-y\\	forward-copy-char (META-y-space)
M-ye	forward-copy-word
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-\\	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.

文件

- /usr/lib/libtecla.so The tecla library
- /usr/include/libtecla.h The tecla header file
- ~/.teclarc The personal tecla customization file

属性

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	library/libtecla
Interface Stability	Committed

另请参见

[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\(5\)](#)

引用名

term – conventional names for terminals

描述

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
-w	Wide mode (more than 80 columns)	att4425-w
-am	With auto. margins (usually default)	vt100-am
-nam	Without automatic margins	vt100-nam
-n	Number of lines on the screen	aaa-60
-na	No arrow keys (leave them in local)	c100-na
-np	Number of pages of memory	c100-4p
-rv	Reverse video	att4415-rv

To avoid conflicts with the naming conventions used in describing the different modes of a terminal (for example, -w), it is recommended that a terminal's root name not contain hyphens. Further, it is good practice to make all terminal names used in the [terminfo\(4\)](#) database unique. Terminal entries that are present only for inclusion in other entries via the `use=` facilities should have a '+' in their name, as in `4415+n1`.

Here are some of the known terminal names: (For a complete list, enter the command `ls -C /usr/share/lib/terminfo/?`).

2621,hp2621	Hewlett-Packard 2621 series
2631	Hewlett-Packard 2631 line printer
2631-c	Hewlett-Packard 2631 line printer, compressed mode
2631-e	Hewlett-Packard 2631 line printer, expanded mode
2640,hp2640	Hewlett-Packard 2640 series
2645,hp2645	Hewlett-Packard 2645 series
3270	IBM Model 3270
33, tty33	AT&T Teletype Model 33 KSR
35, tty35	AT&T Teletype Model 35 KSR
37, tty37	AT&T Teletype Model 37 KSR
4000a	Trendata 4000a
4014,tek4014	TEKTRONIX 4014
40, tty40	AT&T Teletype Dataspeed 40/2
43, tty43	AT&T Teletype Model 43 KSR
4410,5410	AT&T 4410/5410 in 80-column mode, version 2
4410-nfk,5410-nfk	AT&T 4410/5410 without function keys, version 1
4410-nsl,5410-nsl	AT&T 4410/5410 without pln defined
4410-w,5410-w	AT&T 4410/5410 in 132-column mode
4410v1,5410v1	AT&T 4410/5410 in 80-column mode, version 1
4410v1-w,5410v1-w	AT&T 4410/5410 in 132-column mode, version 1
4415,5420	AT&T 4415/5420 in 80-column mode
4415-nl,5420-nl	AT&T 4415/5420 without changing labels
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*.

文件 `/usr/share/lib/terminfo/??/*` compiled terminal description database

另请参见 [sh\(1\)](#), [stty\(1\)](#), [tabs\(1\)](#), [tput\(1\)](#), [vi\(1\)](#), [infocmp\(1M\)](#), [curses\(3CURSES\)](#), [profile\(4\)](#), [terminfo\(4\)](#), [environ\(5\)](#)

引用名 threads, pthreads – POSIX pthreads and Solaris threads concepts

用法概要

```
POSIX      cc -mt [ flag... ] file... [ -lrt library... ]
           #include <pthread.h>
```

```
Solaris    cc -mt [ flag... ] file... [ library... ]
           #include <sched.h>
           #include <thread.h>
```

描述 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
	pthread_create()	thr_create()
	pthread_attr_init()	–

	<hr/>	
	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()
	<hr/>	
Functions Related to Exit	POSIX	Solaris
	pthread_exit()	thr_exit()
	pthread_join()	thr_join()
	pthread_detach()	—
	<hr/>	
Functions Related to Thread Specific Data	POSIX	Solaris
	pthread_key_create()	thr_keycreate()
	pthread_setspecific()	thr_setspecific()
	pthread_getspecific()	thr_getspecific()
	pthread_key_delete()	—
	<hr/>	

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()	

pthread_mutexattr_init()	—
pthread_mutexattr_setpshared()	—
pthread_mutexattr_getpshared()	—
pthread_mutexattr_setprotocol()	—
pthread_mutexattr_getprotocol()	—
pthread_mutexattr_setprioceiling()	—
pthread_mutexattr_getprioceiling()	—
pthread_mutexattr_settype()	—
pthread_mutexattr_gettype()	—
pthread_mutexattr_setrobust()	—
pthread_mutexattr_getrobust()	—
pthread_mutexattr_destroy()	—
pthread_mutex_setprioceiling()	—
pthread_mutex_getprioceiling()	—
pthread_mutex_lock()	mutex_lock()
pthread_mutex_trylock()	mutex_trylock()
pthread_mutex_unlock()	mutex_unlock()
pthread_mutex_destroy()	mutex_destroy()

Functions Related to
Condition Variables

POSIX	Solaris
pthread_cond_init()	cond_init()
pthread_condattr_init()	—
pthread_condattr_setpshared()	—
pthread_condattr_getpshared()	—
pthread_condattr_destroy()	—
pthread_cond_wait()	cond_wait()
pthread_cond_timedwait()	cond_timedwait()
pthread_cond_signal()	cond_signal()
pthread_cond_broadcast()	cond_broadcast()

	pthread_cond_destroy()	cond_destroy()
Functions Related to Reader/Writer Locking	POSIX	Solaris
	pthread_rwlock_init()	rwlock_init()
	pthread_rwlock_rdlock()	rw_rdlock()
	pthread_rwlock_tryrdlock()	rw_tryrdlock()
	pthread_rwlock_wrlock()	rw_wrlock()
	pthread_rwlock_trywrlock()	rw_trywrlock()
	pthread_rwlock_unlock()	rw_unlock()
	pthread_rwlock_destroy()	rwlock_destroy()
	pthread_rwlockattr_init()	—
	pthread_rwlockattr_destroy()	—
	pthread_rwlockattr_getpshared()	—
	pthread_rwlockattr_setpshared()	—
Functions Related to Semaphores	POSIX	Solaris
	sem_init()	sema_init()
	sem_open()	—
	sem_close()	—
	sem_wait()	sema_wait()
	sem_trywait()	sema_trywait()
	sem_post()	sema_post()
	sem_getvalue()	—
	sem_unlink()	—
	sem_destroy()	sema_destroy()
Functions Related to fork() Clean Up	POSIX	Solaris
	pthread_atfork()	—

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 (<i>optimized frequent-read occasional-write mutex</i>), 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 <code>fork(2)</code> , it implicitly calls <code>fork1(2)</code> , which replicates only the calling thread. Should there be any outstanding mutexes throughout the process, the application should call <code>pthread_atfork(3C)</code> to wait for and acquire those mutexes prior to calling <code>fork()</code> .

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_PRIOCNTRL} 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_PRIOCNTRL} 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 [prctl\(2\)](#).

SCHED_FX Threads are scheduled according to the Fixed-Priority Class (FX) policy as described in [prctl\(2\)](#).

SolarisThreads	Only scheduling policy supported is SCHED_OTHER, which is timesharing, based on the TS scheduling class.
----------------	----------------------------------------------------------------------------------------------------------

错误 In a multithreaded application, EINTR can be returned from blocking system calls when another thread calls `forkall(2)`.

用法

-mt compiler option	The -mt compiler option compiles and links for multithreaded code. It compiles source files with <code>-D_REENTRANT</code> and augments the set of support libraries properly.
---------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

属性 See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe, Fork 1-Safe

另请参见 `crle(1)`, `fork(2)`, `priocntl(2)`, `libpthread(3LIB)`, `librt(3LIB)`, `libthread(3LIB)`, `pthread_atfork(3C)`, `pthread_create(3C)`, `attributes(5)`, `privileges(5)`, `standards(5)`

《链接程序和库指南》

引用名	trusted_extensions – Solaris Trusted Extensions
描述	<p>Solaris Trusted Extensions 软件是 Solaris 操作系统 (Solaris OS) 的特定配置。Solaris Trusted Extensions（下称 Trusted Extensions）为本地对象和进程、桌面和窗口系统、区域和文件系统以及网络通信提供标签。这些标签用于实现多层安全性 (MLS) 策略，以便根据标签关系来限制信息流。与基于所有权的自主访问控制 (DAC) 相反，Trusted Extensions 实施的 MLS 策略是强制访问控制 (MAC) 的一个示例。</p> <p>缺省情况下，Trusted Extensions 软件是禁用的。该软件可以通过 FMRI 标识的 labeld(1M) 服务启用和禁用（但不能配置）：</p> <pre>svc:/system/labeld:default</pre> <p>有关使用前必须执行的 Trusted Extensions 软件配置的信息，请参见下面列出的管理员指南。在启用或禁用 labeld 后，必须重新引导系统才能激活或取消激活 Trusted Extensions 软件。</p>
另请参见	<p>labeld(1M), label_encodings(4), labels(5)</p> <p>《Trusted Extensions 配置和管理》</p> <p>《Trusted Extensions 用户指南》</p>

引用名 vgrindefs – vgrind's language definition data base

用法概要 /usr/lib/vgrindefs

描述 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.

示例

示例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'.

文件	<code>/usr/lib/vgrindefs</code>	file containing vgrind descriptions
另请参见	<code>ex(1)</code> , <code>lex(1)</code> , <code>troff(1)</code> , <code>vgrind(1)</code>	

引用名	wbem – Web-Based Enterprise Management	
描述	<p>Web-Based Enterprise Management (WBEM) is a set of management and Internet-related technologies intended to unify the management of enterprise computing environments. Developed by the Distributed Management Task Force (DMTF), WBEM enables organizations to deliver an integrated set of standards-based management tools that support and promote World Wide Web technology. The DMTF has developed a set of standards that make up WBEM. This set of standards includes:</p>	
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>	
	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.
	<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>	
WBEM Tools and Services	<p>Tools and services that enable developers to create and Services management applications and instrumentation that manage heterogeneous computer environments include:</p> <ul style="list-style-type: none"> ■ Solaris WBEM Services ■ Solaris WBEM Software Development Kit (SDK) 	
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"> ■ CIM Object Manager, CIM Repository, and MOF Compiler 	

- 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.

属性

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

另请参见

[appletviewer\(1\)](#), [cimworkshop\(1M\)](#), [init.wbem\(1M\)](#), [mofcomp\(1M\)](#), [mofreg\(1M\)](#), [snmpdx\(1M\)](#), [wbemadmin\(1M\)](#), [wbemconfig\(1M\)](#), [wbemlogviewer\(1M\)](#), [attributes\(5\)](#)

引用名 zones – Solaris 应用程序容器

描述 Solaris 中的区域工具提供了用于运行应用程序的隔离环境。在区域中运行的进程可以免受系统中其他活动的监视或干扰。对其他进程、网络接口、文件系统、设备和进程间通信设备的访问受到限制，防止不同区域中的进程之间相互交互。

在区域中可用的权限受到限制，以防止操作造成系统范围的影响。请参见 [privileges\(5\)](#)。

可以使用 [zoneadm\(1M\)](#) 和 [zonecfg\(1M\)](#) 实用程序配置和管理区域。可以为区域指定配置详细信息，在区域中安装文件系统内容（包括软件包）以及管理区域的运行时状态。可以使用 [zlogin\(1\)](#) 在活动区域中运行命令。您不需要通过基于网络的登录服务器（例如 [in.rlogind\(1M\)](#) 或 [sshd\(1M\)](#)）进行登录，即可执行此操作。

区域的自动引导可以通过 FMRI 标识的区域服务来启用和禁用：

```
svc:/system/zones:default
```

请参见 [zoneadm\(1M\)](#)。请注意，区域有一个 `autoboot` 属性，该属性可以设置为 `true`（始终自动引导）。不过，如果禁用了区域服务，将不发生自动引导，而无论指定区域的 `autoboot` 属性设置如何。请参见 [zonecfg\(1M\)](#)。

每个活动区域使用一个字母数字名称和一个数字 ID 来标识。字母数字名称是使用 [zonecfg\(1M\)](#) 实用程序配置的。数字 ID 是在引导区域时自动分配的。[zonename\(1\)](#) 实用程序可以报告当前区域名称，而 [zoneadm\(1M\)](#) 实用程序可用于报告已配置区域的名称和 ID。

区域可以处于多种状态之一：

CONFIGURED	指示区域配置已完全指定并且已提交到稳定的存储设备。
INCOMPLETE	指示区域正在安装或卸载过程中，或者已在此类转换期间中断。
INSTALLED	指示区域配置已在系统上实例化：软件包已安装在区域的根路径下。
READY	指示区域的“虚拟平台”已建立。例如，文件系统已挂载，设备已配置，但尚未启动任何与区域关联的进程。
RUNNING	指示正在运行与区域应用程序环境关联的用户进程。
SHUTTING_DOWN DOWN	指示正在停止区域。如果区域无法退出应用程序环境状态（例如已挂载的文件系统），或者无法销毁虚拟平台的某个部分，则区域会在其中一个状态中挂起。在这些情况下，需要操作者进行干预。

进程访问限制 在区域（不包括全局区域）中运行的进程对其他进程拥有受限的访问权限。只有相同区域中的进程才能通过 `/proc`（请参见 [proc\(4\)](#)）或者通过接受进程 ID 的系统调用接

口（例如 `kill(2)` 和 `prctl(2)`）来查看。尝试访问存在于其他区域（包括全局区域）中的进程时将失败，并且出现与当指定进程不存在时发布的错误代码相同的错误代码。

权限限制	在非全局区域中运行的进程被限制为拥有一部分权限，从而防止某个区域能够执行可能会影响其他区域的操作。这组权限限制了区域中的特权用户（例如超级用户或 <code>root</code> 用户）的能力。可以使用 <code>ppriv(1)</code> 实用程序显示某个区域中可用的权限列表。有关权限的详细信息，请参见 <code>privileges(5)</code> 。
设备限制	<p>对区域中可用的一组设备进行了限制，以防止某个区域中的进程干扰在其他区域中运行的进程。例如，某个区域中的进程应该无法使用 <code>/dev/kmem</code> 修改内核内存或修改根磁盘的内容。因此，缺省情况下，只有一些被认为可安全在区域中使用的 <code>pseudo</code> 设备可用。在特定区域内，可使用 <code>zonecfg(1M)</code> 实用程序使其他设备变得可用。</p> <p>设备和权限限制对于可在非全局区域中运行的实用程序产生重大影响。例如，<code>eeprom(1M)</code>、<code>prtdiag(1M)</code> 和 <code>prtconf(1M)</code> 实用程序无法在区域中使用，因为它们依赖于通常不可用的设备。</p>
标记	在最初创建区域时，会向该区域分配一个标记。标记区域中的软件与在全局区域中发现的软件不匹配。软件可以包含按照不同的方式配置或布局的 Solaris 软件，也可以包含非 Solaris 软件。这种特定的软件集合称为“标记”（请参见 <code>brands(5)</code> ）。在安装后，除非首先卸载区域，否则无法更改区域的标记。
文件系统	<p>每个区域都有自己的文件系统分层结构部分，根目录称为区域 <code>root</code>。区域中的进程只能访问分层结构的该部分中的文件，即位于区域根路径下的文件。这样可以防止一个区域中的进程损坏或检查与另一个区域关联的文件系统数据。<code>chroot(1M)</code> 实用程序可以在区域中使用，但只能将进程限制为区域中可访问的根路径。</p> <p>为了保留文件系统空间，可以使用 <code>lofs(7FS)</code> 文件系统的只读选项，将文件系统部分挂载到一个或多个区域中。这样可以在多个区域中共享相同的文件系统数据，同时保留区域提供的安全保证。</p> <p>在区域中建立的 NFS 和 <code>autofs</code> 挂载是该区域的本地挂载；它们不能从其他区域（包括全局区域）中访问。在停止或重新引导区域时，将删除挂载。</p> <p>委托给某个区域的 ZFS 数据集可以在该区域中进行管理。在委托数据集中可以创建子数据集。在委托数据集中创建的数据集是自我委托的。可以销毁除顶级委托数据集之外的委托数据集。大多数属性都可以在委托数据集上设置，但并不是所有属性都可以。有关详细信息，请参见 <code>zfs(1M)</code>。</p> <p>每个区域都有一个顶级委托数据集，该数据集又包含 <code>ROOT</code>，并且可能包含其他数据集，例如 <code>.../export</code> 和 <code>.../export/home</code>。存在于 <code>ROOT</code> 数据集下方的数据集构成了区域的引导环境。只应使用 <code>zoneadm(1M)</code> 或 <code>beadm(1M)</code> 命令来创建或销毁引导环境数据集。</p>

联网 区域对于 TCP、UDP 和 SCTP 应用程序使用自己的端口号范围，通常有一个或多个单独的 IP 地址（不过，某些 Trusted Extensions 配置在区域之间共享 IP 地址）。

对于 IP 层（IP 路由、ARP、IPsec、IP 过滤器等），区域可以与全局区域共享配置和状态（共享 IP 区域），也可以具有独特的 IP 层配置和状态（独占 IP 区域）。

如果区域将连接到相同的数据链接，即与全局区域相同的 IP 子网，该区域适合使用共享 IP 实例。

如果区域必须在网络上的 IP 层中隔离（例如连接到与全局区域和其他非全局区域不同的 VLAN 或不同的 LAN），出于隔离原因，区域应该具有其专用的 IP。

共享 IP 区域无法在网络中执行某些操作（例如更改 IP 地址或者发送欺骗性 IP 或 Ethernet 包），而与连接到相同网络接口的单独主机在网络中执行的功能相比，专用 IP 区域具有或多或少的相同功能。特别是，这种区域中的超级用户可以更改 IP 地址和发送欺骗性 ARP 包。

在 `zonecfg(1M)` 中，会向共享 IP 的区域分配一个或多个网络接口名称和 IP 地址。网络接口名称还必须在全局区域中配置。

在 `zonecfg(1M)` 中，会向专用 IP 的区域分配一个或多个网络接口名称。网络接口名称必须专门分配给该区域，即名称不能分配给其他某个正在运行的区域，也无法由全局区域使用。

以 DHCP 客户机、IPsec 和 IP 过滤器形式提供的完整 IP 级别功能在专用 IP 的区域中可用，而在共享 IP 的区域中不可用。

主机标识符 区域能够模拟 32 位主机标识符，该标识符可以通过 `zonecfg(1M)` 针对系统整合目的进行配置。如果区域模拟主机标识符，在区域上下文中执行的 `hostid(1)` 和 `sysdef(1M)` 等命令以及 `sysinfo(2)` 和 `gethostid(3C)` 等 C 接口将显示或返回区域的模拟主机标识符，而不是主机计算机的标识符。

属性 有关下列属性的描述，请参见 `attributes(5)`。

属性类型	属性值
可用性	system/core-os

另请参见 `hostid(1)`、`zlogin(1)`、`zonename(1)`、`beadm(1M)`、`in.rlogind(1M)`、`sshd(1M)`、`sysdef(1M)`、`zfs(1M)`