



man pages section 1M: Maintenance Commands

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900
U.S.A.

Part No: 816-1055-10
November 2001

Copyright 2001 Sun Microsystems, Inc. 901 San Antonio Road Palo Alto, CA 94303-4900 U.S.A. All rights reserved.

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

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

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

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

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

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

Copyright 2001 Sun Microsystems, Inc. 901 San Antonio Road Palo Alto, CA 94303-4900 U.S.A. Tous droits réservés

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

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

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

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

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



011029@2471



Contents

Preface 9

Introduction 15

Intro(1M) 16

Maintenance Commands 21

accept(1M) 22

add_allocatable(1M) 24

add_drv(1M) 26

adminvi(1M) 30

arp(1M) 32

atohexlabel(1M) 34

audit(1M) 35

auditconfig(1M) 37

auditd(1M) 41

auditreduce(1M) 43

audit_startup(1M) 51

auditstat(1M) 52

audit_warn(1M) 54

automount(1M) 56

automountd(1M) 64

autopush(1M) 65

bootparamd(1M) 67

bsmconv(1M) 68

bsmunconv(1M) 69

chk_encodings(1M) 70
 chroot(1M) 71
 clist(1M) 72
 coreadm(1M) 73
 cron(1M) 77
 devfsadm(1M) 79
 devfsadmd(1M) 81
 device_clean(1M) 83
 devpolicy(1M) 84
 dfmounts(1M) 85
 dfshares(1M) 87
 dispadmin(1M) 89
 dl_booting(1M) 91
 dl_restore(1M) 92
 dminfo(1M) 93
 drvconfig(1M) 95
 du(1M) 98
 eeprom(1M) 101
 format(1M) 107
 fsdb_ufs(1M) 111
 ftpd(1M) 120
 fuser(1M) 130
 getfsattr(1M) 132
 getfsattr_ufs(1M) 133
 halt(1M) 134
 hextoalabel(1M) 135
 ifconfig(1M) 136
 inetd(1M) 150
 in.ftpd(1M) 153
 init(1M) 163
 init.wbem(1M) 168
 in.named(1M) 171
 in.rarpd(1M) 192
 in.rdisc(1M) 194
 in.rexecd(1M) 196
 in.rlogind(1M) 198
 in.routed(1M) 200

in.rshd(1M)	205
install(1M)	208
in.tftpd(1M)	210
ipseccnf(1M)	211
ipseckey(1M)	226
lockd(1M)	235
lpadmin(1M)	236
lpfilter(1M)	248
lpforms(1M)	254
lpmove(1M)	261
lpsched(1M)	263
lpshut(1M)	265
lpsystem(1M)	266
lpusers(1M)	267
mkdevalloc(1M)	269
mkdevdb(1M)	271
mkdevmaps(1M)	273
modload(1M)	275
modunload(1M)	276
mount(1M)	277
mountall(1M)	284
mouted(1M)	286
mount_hfs(1M)	288
mount_nfs(1M)	291
mount_pcfs(1M)	300
mount_tmpfs(1M)	302
mount_ufs(1M)	305
named(1M)	310
ndd(1M)	331
netstat(1M)	333
newsecfs(1M)	339
nfsd(1M)	341
nfsstat(1M)	343
nis_cachemgr(1M)	348
nisclient(1M)	350
nisd(1M)	355
nisd_resolv(1M)	358

nispasswd(1M) 359
 nispopulate(1M) 361
 nissetup(1M) 366
 nscd(1M) 367
 nslookup(1M) 369
 pbind(1M) 378
 pfsh(1M) 381
 pkgchk(1M) 382
 poweroff(1M) 385
 praudit(1M) 386
 prtconf(1M) 388
 psradm(1M) 392
 rarpd(1M) 395
 rdate(1M) 397
 rdisc(1M) 398
 reboot(1M) 400
 reject(1M) 402
 rem_drv(1M) 404
 remove_allocatable(1M) 405
 rexecd(1M) 406
 rlogind(1M) 408
 rmmount(1M) 410
 route(1M) 413
 routed(1M) 420
 rpcbind(1M) 425
 rpc.bootparamd(1M) 427
 rpc.getpeerinfod(1M) 428
 rpcinfo(1M) 429
 rpc.nisd(1M) 434
 rpc.nisd_resolv(1M) 437
 rpc.nispasswd(1M) 438
 rpc.tbootparamd(1M) 440
 rpc.yppasswdd(1M) 441
 rpc.yppupdated(1M) 443
 rshd(1M) 444
 runpd(1M) 447
 rwall(1M) 448

sendmail(1M)	449
setaudit(1M)	465
setfsattr(1M)	466
setuname(1M)	468
share(1M)	469
shareall(1M)	471
share_nfs(1M)	472
showmount(1M)	481
smc(1M)	482
smcron(1M)	487
smexec(1M)	493
smgroup(1M)	499
smhost(1M)	503
smmaillist(1M)	508
smmultiuser(1M)	512
smnetidb(1M)	517
smnettpl(1M)	521
smnetwork(1M)	526
smprofile(1M)	530
smrole(1M)	536
smuser(1M)	545
snoop(1M)	556
spray(1M)	568
statd(1M)	569
su(1M)	570
swap(1M)	573
sysdef(1M)	576
sysh(1M)	578
tbootparam(1M)	579
telinit(1M)	580
tftpd(1M)	585
tnchkdb(1M)	586
tnctl(1M)	588
tnd(1M)	590
tninfo(1M)	592
tokmapctl(1M)	594
tokmapd(1M)	596

traceroute(1M) 598
uadmin(1M) 605
umount(1M) 606
umountall(1M) 613
unshare(1M) 615
unshareall(1M) 616
unshare_nfs(1M) 617
updatehome(1M) 618
writeaudit(1M) 620
ypbind(1M) 623
yppasswdd(1M) 626
ypserv(1M) 628
ypupdated(1M) 631
ypxfr(1M) 632
ypxfr_1perday(1M) 634
ypxfr_1perhour(1M) 636
ypxfr_2perday(1M) 638
ypxfrd(1M) 640

Index 643

Preface

Overview

A man page is provided for both the naive user and the sophisticated user who is familiar with the Trusted Solaris operating environment and is in need of online information. A man page is intended to answer concisely the question “What does it do?” The man pages in general comprise a reference manual. They are not intended to be a tutorial.

Trusted Solaris Reference Manual

In the AnswerBook2™ and online man command forms of the man pages, all man pages are available:

- Trusted Solaris man pages that are unique for the Trusted Solaris environment
- SunOS 5.8 man pages that have been changed in the Trusted Solaris environment
- SunOS 5.8 man pages that remain unchanged.

The printed manual, the *Trusted Solaris 8 Reference Manual* contains:

- Man pages that have been added to the SunOS operating system by the Trusted Solaris environment
- Man pages that originated in SunOS 5.8, but have been modified in the Trusted Solaris environment to handle security requirements.

Users of printed manuals need both manuals in order to have a full set of man pages, since the *SunOS 5.8 Reference Manual* contains the common man pages that are not modified in the Trusted Solaris environment.

Man Page Sections

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

- Section 1 describes, in alphabetical order, commands available with the operating system.
- Section 1M describes, in alphabetical order, commands that are used chiefly for system maintenance and administration purposes.
- Section 2 describes all of the system calls. Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible returned value.
- Section 3 describes functions found in various libraries, other than those functions that directly invoke UNIX system primitives, which are described in Section 2 of this volume.
- Section 4 outlines the formats of various files. The C structure declarations for the file formats are given where applicable.
- Section 5 contains miscellaneous documentation such as character set tables.
- Section 6 contains available games and demos.
- Section 7 describes various special files that refer to specific hardware peripherals, and device drivers. STREAMS software drivers, modules and the STREAMS-generic set of system calls are also described.
- Section 9 provides reference information needed to write device drivers in the kernel operating systems environment. It describes two device driver interface specifications: the Device Driver Interface (DDI) and the Driver/Kernel Interface (DKI).
- Section 9E describes the DDI/DKI, DDI-only, and DKI-only entry-point routines a developer may include in a device driver.
- Section 9F describes the kernel functions available for use by device drivers.
- Section 9S describes the data structures used by drivers to share information between the driver and the kernel.

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

NAME

This section gives the names of the commands or functions documented, followed by a brief description of what they do.

SYNOPSIS

This section shows the syntax of commands or functions. When a command or file does not exist in the standard path, its full pathname is shown. Options and

arguments are alphabetized, with single letter arguments first, and options with arguments next, unless a different argument order is required.

The following special characters are used in this section:

- [] The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.
- . . . Ellipses. Several values may be provided for the previous argument, or the previous argument can be specified multiple times, for example, 'filename...'.
"filename...".
- | Separator. Only one of the arguments separated by this character can be specified at a time.
- { } Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.

PROTOCOL

This section occurs only in subsection 3R to indicate the protocol description file.

DESCRIPTION

This section defines the functionality and behavior of the service. Thus it describes concisely what the command does. It does not discuss OPTIONS or cite EXAMPLES. Interactive commands, subcommands, requests, macros, functions and such, are described under USAGE.

IOCTL

This section appears on pages in Section 7 only. Only the device class which supplies appropriate parameters to the `ioctl` (2) system call is called `ioctl` and generates its own heading. `ioctl` calls for a specific device are listed alphabetically (on the man page for that specific device). `ioctl` calls are used for a particular class of devices all of which have an `io` ending, such as `mtio`(7I)

OPTIONS

This section lists the command options with a concise summary of what each option does. The options are listed literally and in the order they appear in the SYNOPSIS section. Possible arguments to options are discussed under the option, and where appropriate, default values are supplied.

OPERANDS

This section lists the command operands and describes how they affect the actions of the command.

OUTPUT

This section describes the output – standard output, standard error, or output files – generated by the command.

RETURN VALUES

If the man page documents functions that return values, this section lists these values and describes the conditions under which they are returned. If a function can return only constant values, such as 0 or -1, these values are listed in tagged

paragraphs. Otherwise, a single paragraph describes the return values of each function. Functions declared void do not return values, so they are not discussed in RETURN VALUES.

ERRORS

On failure, most functions place an error code in the global variable `errno` indicating why they failed. This section lists alphabetically all error codes a function can generate and describes the conditions that cause each error. When more than one condition can cause the same error, each condition is described in a separate paragraph under the error code.

USAGE

This section lists special rules, features, and commands that require in-depth explanations. The subsections listed here are used to explain built-in functionality:

- Commands
- Modifiers
- Variables
- Expressions
- Input Grammar

EXAMPLES

This section provides examples of usage or of how to use a command or function. Wherever possible a complete example including command-line entry and machine response is shown. Whenever an example is given, the prompt is shown as `example%`, or if the user must be root, `example#`. Examples are followed by explanations, variable substitution rules, or returned values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS, and USAGE sections.

ENVIRONMENT VARIABLES

This section lists any environment variables that the command or function affects, followed by a brief description of the effect.

EXIT STATUS

This section lists the values the command returns to the calling program or shell and the conditions that cause these values to be returned. Usually, zero is returned for successful completion, and values other than zero for various error conditions.

FILES

This section lists all file names referred to by the man page, files of interest, and files created or required by commands. Each is followed by a descriptive summary or explanation.

ATTRIBUTES

This section lists characteristics of commands, utilities, and device drivers by defining the attribute type and its corresponding value. See `attributes(5)` for more information.

SUMMARY OF TRUSTED SOLARIS CHANGES

This section describes changes to a Solaris item by Trusted Solaris software. It is present in man pages that have been modified from Solaris software.

SEE ALSO

This section lists references to other man pages, in-house documentation and outside publications. The references are divided into two sections, so that users of printed manuals can easily locate a man page in its appropriate printed manual.

DIAGNOSTICS

This section lists diagnostic messages with a brief explanation of the condition causing the error.

WARNINGS

This section lists warnings about special conditions which could seriously affect your working conditions. This is not a list of diagnostics.

NOTES

This section lists additional information that does not belong anywhere else on the page. It takes the form of an aside to the user, covering points of special interest. Critical information is never covered here.

BUGS

This section describes known bugs and, wherever possible, suggests workarounds.

Introduction

Intro(1M)

NAME	Intro – introduction to maintenance commands and application programs												
DESCRIPTION	<p>This section describes Trusted Solaris™ commands that are used chiefly for system maintenance and administration. The Trusted Solaris environment includes the following commands:</p> <ul style="list-style-type: none"> ■ Commands that are unique to and originate in the Trusted Solaris environment, such as <code>adminvi(1M)</code>, which enables administrators and other users to edit files while preventing certain <code>vi</code> actions that present a security risk. ■ SunOS 5.8 commands that have been modified to work within the Trusted Solaris security policy, such as <code>mount(1M)</code>. Man pages for modified commands have been rewritten to remove information that is not accurate for how the command behaves within the Trusted Solaris environment. Modified man pages also add descriptions for any new features, options, and arguments. ■ SunOS 5.8 commands that remain unchanged from the Solaris 8 release, such as <code>ln</code>. <p>Because of command restructuring for the Virtual File System architecture, there are several instances of multiple manual pages that begin with the same name. For example, there are multiple <code>mount</code> pages – <code>mount(1M)</code>, <code>mount_hfs(1M)</code>, <code>mount_nfs(1M)</code>, <code>mount_tmpfs(1M)</code>, and <code>mount_ufs(1M)</code>. In each such case the first of the multiple pages describes the syntax and options of the generic command, that is, those options applicable to all FSTypes (file system types). The succeeding pages describe the functionality of the FSType-specific modules of the command. These pages list the command followed by an underscore (<code>_</code>) and the FSType to which they pertain. Note that the administrator should not attempt to call these modules directly. The generic command provides a common interface to all of them. Thus the FSType-specific manual pages should not be viewed as describing distinct commands, but rather as detailing those aspects of a command that are specific to a particular FSType.</p>												
COMMAND SYNTAX	<p>Unless otherwise noted, commands described in this section accept options and other arguments according to the following syntax:</p> <p><i>name</i> [<i>option</i> (s)] [<i>cmdarg</i> (s)] where:</p> <table> <tr> <td><i>name</i></td><td>The name of an executable file.</td></tr> <tr> <td><i>option</i></td><td>– <i>noargletter</i>(s) or, – <i>argletter</i>< ><i>optarg</i></td></tr> <tr> <td></td><td>where < > is optional white space.</td></tr> <tr> <td><i>noargletter</i></td><td>A single letter representing an option without an argument.</td></tr> <tr> <td><i>argletter</i></td><td>A single letter representing an option requiring an argument.</td></tr> <tr> <td><i>optarg</i></td><td>Argument (character string) satisfying preceding <i>argletter</i>.</td></tr> </table>	<i>name</i>	The name of an executable file.	<i>option</i>	– <i>noargletter</i> (s) or, – <i>argletter</i> < > <i>optarg</i>		where < > is optional white space.	<i>noargletter</i>	A single letter representing an option without an argument.	<i>argletter</i>	A single letter representing an option requiring an argument.	<i>optarg</i>	Argument (character string) satisfying preceding <i>argletter</i> .
<i>name</i>	The name of an executable file.												
<i>option</i>	– <i>noargletter</i> (s) or, – <i>argletter</i> < > <i>optarg</i>												
	where < > is optional white space.												
<i>noargletter</i>	A single letter representing an option without an argument.												
<i>argletter</i>	A single letter representing an option requiring an argument.												
<i>optarg</i>	Argument (character string) satisfying preceding <i>argletter</i> .												

**Rules for the
Display and
Entering of Labels***cmdarg*

Pathname (or other command argument) *not* beginning with – or,
– by itself indicating the standard input.

When entering labels on the command line in a UNIX shell, follow these rules. For rules for entering labels in graphical user interfaces, see *Rules for the Display and Entering of Labels*. For rules for entering labels in configuration files, see *RULES FOR INCLUDING LABELS IN A CONFIGURATION FILE* in Intro(4).

Enter a sensitivity label (SL) or clearance, in text in the form:

```
{ + } { classification } { { +|- }word } ...
```

Items in curly brackets are optional. A vertical bar (|) represents a choice between two items. Items followed by an ellipsis may be repeated zero or more times. Leading and trailing white space is ignored. Items may be separated by blanks, tabs, commas or slashes (/).

The system always displays labels in uppercase. Users may enter labels in any combination of uppercase and lowercase.

The classification part of the label must be a valid classification name as defined in *label_encodings*(4). Classification names may contain embedded blanks or punctuation, if they are so defined in *label_encodings*. Short and long forms of classification names may be used interchangeably.

The words (*compartments* and *markings*) used in labels must be valid words as defined in *label_encodings*. Words may contain embedded blanks or punctuation if they are so defined in *label_encodings*.

Short and long forms of words may be used interchangeably. Words may be specified in any order; however they are processed left to right, so that where words conflict with each other, the word furthest to the right takes precedence.

You may use plus and minus signs when modifying an existing label to turn on or off the compartments and markings associated with the words.

A CMW label is represented in text in the form:

```
ADMIN_LOW [ SENSITIVITY_LABEL ]
```

Items in curly brackets are optional. Leading and trailing white space is ignored. Items may be separated by blanks, tabs, commas, or slashes (/).

EXAMPLES**EXAMPLE 1** Using quotes in labels

On the command line, enclose any label with more than one word in double quotes because, without quotes, a second word or letter separated by a space is interpreted as a second argument. Enclose labels containing [and] characters in quotes to suppress the shell's use of those characters in filename substitution.

```
$ setlabel "[ts a b]" somefile
$ setlabel "[ts,a,b]" somefile
```

EXAMPLE 1 Using quotes in labels *(Continued)*

```
$ setlabel "[ts/a      b]" somefile
```

EXAMPLE 2 Using case in labels

Use any combination of upper and lowercase letters. You may separate items in a label with blanks, tabs, commas or slashes (/). Do not use any other punctuation.

```
$ setlabel -s SECRET somefile
```

EXAMPLE 3 Using brackets in labels

When entering an SL with a command option that sets the SL, you do not need to use brackets around the SL.

```
$ setlabel -s "TOP SECRET A B" somefile
```

EXAMPLE 4 Setting a label

To set somefile's SL to SECRET A.

```
$ setlabel "[Secret a]" somefile
```

To turn on compartment B in *somefile*'s SL.

```
$ setlabel -s +b somefile
```

To turn off compartment A in *somefile*'s SL.

```
$ setlabel -s -A somefile
```

**TRUSTED
SOLARIS
DIFFERENCES**

The responsibilities and privileges of the superuser have been divided among several administrative roles. When a man page that has not been modified for the Trusted Solaris system states that superuser is required to execute a certain command or option, remember that one or more privileges are required instead. The site's security administrator can perform privilege debugging [see *runpd(1M)*] to find out which privileges are needed and can then decide to give the privilege to the command after assessing whether the command and any users set up to use that command can make use of the privilege in a manner that does not violate the site's security policy.

The ability of the UNIX superuser to bypass access restrictions, to execute restricted commands, and to use some command options not available to other users has been replaced with the *profile mechanism*, which allows the security administrator to assign to various users different sets of commands and to assign different privileges to the commands using *rights profiles*. When a command or one of its options needs a privilege in order to succeed, that privilege is a *required* privilege; if the required privilege is not given to the command in a user's rights profile by the security administrator, the command will not work. Required privileges are indicated on the

man page with the words "must have," as shown in this sentence: "The `ifconfig(1M)` command must have the `sys_net_config` privilege to modify network interfaces."

In other cases, when the command is designed to work within security policy and it fails when certain DAC or MAC checks are not passed, an *override* privilege may be assigned at the security administrator's discretion. On man pages, the names of privileges that may be used to override access restrictions are given in the `ERRORS` section. The override privileges that may be given to bypass DAC or MAC restrictions on files or directories are given below:

The DAC override privileges are `file_dac_read` and `file_dac_write`. If a user does not have DAC access to a file, the security administrator may assign one or both of these privileges to the command, depending on whether read or write access or both are desired. The MAC override privileges are `file_mac_read` and `file_mac_write`. If a user does not have MAC access to a file, the security administrator may assign one or both of these privileges to the command, depending on whether read or write access or both are desired.

Besides being able to assign an override privilege, the security administrator has other options. For example, to avoid the use of privilege the security administrator may specify that the command will execute with another user's ID (usually the root ID 0) or group ID, one that allows access to the file or directory based on its permissions or its ACL.

To find out how privileges are made available to commands and to find out exactly which tasks, commands, and privileges are assigned to each of the roles by means of rights profiles shipped with the default system, see *Trusted Solaris Administrator's Procedures*.

Also, check with your security administrator to find out which roles are configured at your site and if any of the roles have been reconfigured to suit your site's security policy.

SUMMARY OF TRUSTED SOLARIS CHANGES

Besides the usual UNIX DAC checks performed when a process acting on behalf of a user attempts to access a file or directory, *mandatory access* checks also must be passed. For each possible type of access failure, a specific override *privilege* may be assigned to the command at the security administrator's discretion.

The printed *Trusted Solaris 8 4/01 Reference Manual* contains only the Trusted Solaris original and modified (from the Solaris environment) man pages. The online set of man pages viewed by the `man` command accesses all man pages; AnswerBook2™ can access all man pages in the AnswerBook2 collections. For a fuller description, see *Trusted Solaris Manual Page Display* in Intro(1). The `SEE ALSO` man page heading has been subdivided to help users of the printed manual locate a referenced man page.

Note – When a SUMMARY OF TRUSTED SOLARIS CHANGES is provided on a modified man page, it is intended as a convenience to summarize for you the major

Intro(1M)

changes all in one place. Do not rely on the SUMMARY OF TRUSTED SOLARIS CHANGES alone, but also read the entire man page.

ATTRIBUTES

See attributes(5) in the *SunOS 5.8 Reference Manual* for a discussion of the attributes listed in this section.

SEE ALSO

Commands that are listed under the Trusted Solaris 8 4/01 Reference Manual heading in the SEE ALSO section are commands that have been changed or added in the Trusted Solaris environment. Commands that are listed under the SunOS 5.8 Reference Manual heading in the SEE ALSO section are commands that are unchanged in the Trusted Solaris environment. If you are using printed manuals, refer to the *SunOS 5.8 Reference Manual* for Solaris commands that are unchanged in the Trusted Solaris environment.

**Trusted Solaris 8
4/01 Reference
Manual**

runpd(1M)

Trusted Solaris Administration Overview, Trusted Solaris Administrator's Procedures

**SunOS 5.8
Reference Manual
DIAGNOSTICS**

getopt(1), getopt(3C), attributes(5)

Upon termination, each command returns 0 for normal termination and non-zero to indicate troubles such as erroneous parameters, bad or inaccessible data, or other inability to cope with the task at hand. It is called variously "exit code," "exit status," or "return code," and is described only where special conventions are involved.

NOTES

Unfortunately, not all commands adhere to the standard syntax.

Maintenance Commands

accept(1M)

NAME	accept, reject – Accept or reject print requests
SYNOPSIS	accept <i>destination</i> ... reject [-r <i>reason</i>] <i>destination</i> ...
DESCRIPTION	<p>accept allows the queueing of print requests for the named destinations.</p> <p>reject prevents queueing of print requests for the named destinations.</p> <p>Use lpstat -a to check if destinations are accepting or rejecting print requests.</p> <p>accept and request must be run on the print server; they have no meaning on a client system.</p>
OPTIONS	<p>The following options are supported for reject.</p> <p>-r <i>reason</i> Assigns a reason for rejection of print requests for <i>destination</i>. Enclose <i>reason</i> in quotes if it contains blanks. <i>reason</i> is reported by lpstat -a. By default, <i>reason</i> is unknown reason for existing destinations, and new printer for destinations added to the system but not yet accepting requests.</p>
OPERANDS	<p>The following operands are supported.</p> <p><i>destination</i> The name of the destination accepting or rejecting print requests. Destination specifies the name of a printer or class of printers [see lpadmin(1M)]. Specify <i>destination</i> using atomic name. See printers.conf(4) for information regarding the naming conventions for atomic names.</p>
EXIT STATUS	<p>The following exit values are returned:</p> <p>0 Successful completion.</p> <p>non-zero An error occurred.</p>
FILES	/var/spool/lp/* LP print queue.
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWpcu
CSI	Enabled (see NOTES)

SUMMARY OF TRUSTED SOLARIS CHANGES	See the lpstat(1) man page for Trusted Solaris security that affects the accept command.
Trusted Solaris 8 4/01 Reference Manual	enable(1) , lp(1) , lpstat(1) , lpadmin(1M) , lpsched(1M)

accept(1M)

printers.conf (4), attributes(5)

accept and reject only affect queuing on the print server's spooling system. Requests made from a client system remain queued in the client system's queuing mechanism until they are cancelled or accepted by the print server's spooling system.

accept is CSI-enabled except for the *destination* name.

add_allocatable(1M)

NAME	add_allocatable – add entries to allocation databases															
SYNOPSIS	/usr/sbin/add_allocatable [-f] [-s] -n name -t type -d device-list [-l minSL] [-h maxSL] [-a authorization] [-c clean] [-o key=value]															
DESCRIPTION	<p>add_allocatable creates or updates database entries for allocatable devices and certain non-allocatable devices. add_allocatable updates the device_allocate(4) and device_maps(4) databases. The database entries are needed when devices are user-allocatable. The database entries are also needed for the frame buffer and printers because the label ranges for these non-allocatable devices are managed by the device allocation mechanism.</p> <p>add_allocatable can be used in shell scripts, such as installation scripts for driver packages, to automate the administrative work of setting up a new device.</p>															
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:															
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWtsu</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu										
ATTRIBUTE TYPE	ATTRIBUTE VALUE															
Availability	SUNWtsu															
OPTIONS	<table><tr><td>-f</td><td>Force an update of an already-existing entry with the specified information. add_allocatable exits with an error if this option is not specified when an entry with the specified device name already exists.</td></tr><tr><td>-s</td><td>Turn on silent mode. add_allocatable does not print any error or warning messages.</td></tr><tr><td>-n name</td><td>Puts a device name name into the device-name fields of the device’s entries in device_allocate and device_maps.</td></tr><tr><td>-t type</td><td>Puts a device type type into the device-type fields of the device’s entries in device_allocate and device_maps.</td></tr><tr><td>-d device-list</td><td>Puts the device special file names specified in device-list into the device-list field of the device’s entry in device_maps. The list of devices must be separated by white space, and the list must be quoted because white spaces are treated by the shell as argument delimiters.</td></tr><tr><td>-l minSL</td><td>Puts the minimum sensitivity label minSL of the device into the dev-minimum field of the device’s entry in device_allocate. The default sensitivity label ADMIN_LOW is used when this optional argument is not specified.</td></tr><tr><td>-h maxSL</td><td>Puts the maximum sensitivity label maxSL into the dev-maximum field of the device’s entry in device_allocate. The default sensitivity label ADMIN_HIGH is used when this optional argument is not specified.</td></tr></table>		-f	Force an update of an already-existing entry with the specified information. add_allocatable exits with an error if this option is not specified when an entry with the specified device name already exists.	-s	Turn on silent mode. add_allocatable does not print any error or warning messages.	-n name	Puts a device name name into the device-name fields of the device’s entries in device_allocate and device_maps.	-t type	Puts a device type type into the device-type fields of the device’s entries in device_allocate and device_maps.	-d device-list	Puts the device special file names specified in device-list into the device-list field of the device’s entry in device_maps. The list of devices must be separated by white space, and the list must be quoted because white spaces are treated by the shell as argument delimiters.	-l minSL	Puts the minimum sensitivity label minSL of the device into the dev-minimum field of the device’s entry in device_allocate. The default sensitivity label ADMIN_LOW is used when this optional argument is not specified.	-h maxSL	Puts the maximum sensitivity label maxSL into the dev-maximum field of the device’s entry in device_allocate. The default sensitivity label ADMIN_HIGH is used when this optional argument is not specified.
-f	Force an update of an already-existing entry with the specified information. add_allocatable exits with an error if this option is not specified when an entry with the specified device name already exists.															
-s	Turn on silent mode. add_allocatable does not print any error or warning messages.															
-n name	Puts a device name name into the device-name fields of the device’s entries in device_allocate and device_maps.															
-t type	Puts a device type type into the device-type fields of the device’s entries in device_allocate and device_maps.															
-d device-list	Puts the device special file names specified in device-list into the device-list field of the device’s entry in device_maps. The list of devices must be separated by white space, and the list must be quoted because white spaces are treated by the shell as argument delimiters.															
-l minSL	Puts the minimum sensitivity label minSL of the device into the dev-minimum field of the device’s entry in device_allocate. The default sensitivity label ADMIN_LOW is used when this optional argument is not specified.															
-h maxSL	Puts the maximum sensitivity label maxSL into the dev-maximum field of the device’s entry in device_allocate. The default sensitivity label ADMIN_HIGH is used when this optional argument is not specified.															

add_allocatable(1M)

- a *authorization* Puts one or more authorizations or other characters specified in *authorization* into the device authorization field of the device's entry in `device_allocate`. When more than one authorization is specified, the list must be separated by commas and must be quoted. When the device is not allocatable, *authorization* is specified with an asterisk (*) and must be quoted. When the device is allocatable and is allocatable by any user, *authorization* is specified with the at sign (@) and must be quoted. When this optional argument is not specified, the default value '@' is used, and the device is allocatable by any user with no authorization required.
- c *clean* Puts the `device_clean(1M)` program *clean* into the device-clean field of the device's entry in `device_allocate(4)`. The default value `/bin/true` is used when this optional argument is not specified.
- o *key=value* Accepts a string of colon-separated *key=value* pairs for incorporation into `device_allocate(4)`.

ERRORS When successful, `add_allocate` returns an exit status of 0 (true). `add_allocate` returns a nonzero exit status in the event of an error. The exit codes are as follows:

- 1 Invocation syntax error
- 2 Unknown system error
- 3 A `device_allocate` entry already exists. This error occurs only when the `-f` option is not specified.
- 4 Permission denied. User does not have DAC or MAC access to database.

FILES `/etc/security/device_allocate`
Mandatory access control file for devices
`/etc/security/device_maps`
List of physical devices associated with a device name and type

**Trusted Solaris 8
4/01 Reference
Manual**
**SunOS 5.8
Reference Manual**

`allocate(1)`, `device_allocate(4)`, `device_clean(1M)`, `device_maps(4)`,
`remove_allocatable(1M)`
`attributes(5)`

add_drv(1M)

NAME	add_drv – Add a new device driver to the system	
SYNOPSIS	add_drv [-b <i>basedir</i>] [-c <i>class_name</i>] [-i ' <i>identify_name...</i> '] [-m ' <i>permission</i> ', ' <i>...</i> '] [-n] [-f] [-v] <i>device_driver</i>	
DESCRIPTION	<p>The add_drv command is used to inform the system about newly installed device drivers.</p> <p>Each device on the system has a name associated with it. This name is represented by the <code>name</code> property for the device. Similarly, the device may also have a list of driver names associated with it. This list is represented by the <code>compatible</code> property for the device.</p> <p>The system determines which devices will be managed by the driver being added by examining the contents of the <code>name</code> property and the <code>compatible</code> property (if it exists) on each device. If the value in the <code>name</code> property does not match the driver being added, each entry in the <code>compatible</code> property is tried, in order, until either a match occurs or there are no more entries in the <code>compatible</code> property.</p> <p>In some cases, adding a new driver may require a reconfiguration boot. See the NOTES section.</p>	
OPTIONS	-b <i>basedir</i>	Installs the driver on the system with a root directory of <i>basedir</i> rather than installing on the system executing add_drv. This option is typically used in package post-installation scripts when the package is not being installed on the system executing the pkgadd command. The system using <i>basedir</i> as its root directory must reboot to complete the driver installation.
	-c <i>class_name</i>	The driver being added to the system exports the class <i>class_name</i> .
	-i ' <i>identify_name</i> '	A white-space separated list of aliases for the driver <i>device_driver</i> .
	-m ' <i>permission</i> '	Specify the file system permissions for device nodes created by the system on behalf of <i>device_driver</i> .
	-n	Do not try to load and attach <i>device_driver</i> , just modify the system configuration files for the <i>device_driver</i> .
	-f	Normally if a reconfiguration boot is required to complete the configuration of the driver into the system, add_drv will not add the driver. The force flag forces add_drv to add the driver even if a reconfiguration boot is required. See the -v flag.
	-v	The verbose flag causes add_drv to provide additional information regarding the success or failure of a driver's configuration into the system. See the EXAMPLES section.

EXAMPLES**EXAMPLE 1** Adding The SUNW, Example Driver to the System

The following example adds the SUNW, example driver to the system, with an alias name of SUNW, alias. It assumes the driver has already been copied to /usr/kernel/drv.

```
example# add_drv -m '* 0666 bin bin', 'a 0644 root sys' \
-i 'SUNW,alias' SUNW,example
```

Every minor node created by the system for the SUNW, example driver will have the permission 0666, and be owned by user bin in the group bin, except for the minor device a, which will be owned by root, group sys, and have a permission of 0644.

EXAMPLE 2 Adding The Driver To The Client /export/root/sun1

The following example adds the driver to the client /export/root/sun1. The driver is installed and loaded when the client machine, sun1, is rebooted. This second example produces the same result as the first, except the changes are on the diskless client, sun1, and the client must be rebooted for the driver to be installed.

```
example# add_drv -m '* 0666 bin bin', 'a 0644 root sys' \
-i 'SUNW,alias' -b /export/root/sun1 \
SUNW,example
```

EXAMPLE 3 Adding A Driver For A Device That Is Already Managed By An Existing Driver

The following example illustrates the case where a new driver is added for a device that is already managed by an existing driver. Consider a device that is currently managed by the driver dumb_framebuffer. The name and compatible properties for this device are as follows:

```
name="display"
compatible="whizzy_framebuffer", "dumb_framebuffer"
```

If add_drv is used to add the whizzy_framebuffer driver, the following will result.

```
example# add_drv whizzy_framebuffer
Error: Could not install driver (whizzy_framebuffer)
Device managed by another driver.
```

If the -v flag is specified, the following will result.

```
example# add_drv -v whizzy_framebuffer
Error: Could not install driver (whizzy_framebuffer)
Device managed by another driver.
Driver installation failed because the following
entries in /devices would be affected:
```

```
/devices/iommu@f,e0000000/sbus@f,e0001000/display[:*]
(Device currently managed by driver "dumb_framebuffer")
```

The following entries in /dev would be affected:

```
/dev/fbs/dumb_framebuffer0
```

add_drv(1M)

EXAMPLE 3 Adding A Driver For A Device That Is Already Managed By An Existing Driver (Continued)

If the -v and -f flags are specified, the driver will be added resulting in the following.

```
example# add_drv -vf whizzy_framebuffer
A reconfiguration boot must be performed to complete the
installation of this driver.
```

The following entries in /devices will be affected:

```
/devices/iommu@f,e0000000/sbus@f,e0001000/display[:*]
(Device currently managed by driver "dumb_framebuffer"
```

The following entries in /dev will be affected:

```
/dev/fbs/dumb_framebuffer0
```

The above example is currently only relevant to devices exporting a generic device name.

EXIT STATUS

add_drv returns 0 on success and 1 on failure.

FILES

/kernel/drv	Boot device drivers
/usr/kernel/drv	Other drivers that could potentially be shared between platforms
/platform/'uname -i'/kernel/drv	Platform-dependent drivers
/etc/driver_aliases	Driver aliases file
/etc/driver_classes	Driver classes file
/etc/minor_perm	Minor node permissions
/etc/name_to_major	Major number binding

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

To succeed, this command needs the sys_devices privilege. This command is intended to be invoked at ADMIN_LOW with effective user ID 0; when invoked by other users, the command needs the file_dac_write privilege

Trusted Solaris 8 4/01 Reference Manual

devfsadm(1M), drvconfig(1M), rem_drv(1M)

boot(1M), devlinks(1M), disks(1M), kernel(1M), modinfo(1M), ports(1M),
tapes(1M), driver.conf(4), system(4), attributes(5),
ddi_create_minor_node(9F)

Writing Device Drivers

- NOTES** Aliases may require quoting (with double-quotes) if they contain numbers.
- It is possible to add a driver for a device already being managed by a different driver, where the driver being added appears in the device's `compatible` list before the current driver. In such cases, a reconfiguration boot is required (see `boot(1M)` and `kernel(1M)`). After the reconfiguration boot, device nodes in `/devices`, entries in `/dev`, and references to these files may no longer be valid (see the `-v` flag). If a reconfiguration boot would be required to complete the driver installation, `add_drv` will fail unless the `-f` option is specified. See Example 3 in the `EXAMPLES` section.
- BUGS** `add_drv` will accept a full pathname for *device_driver*. However, the kernel does not use the full pathname; it only uses the final component and searches the internal driver search path for the driver. This can lead to the kernel loading a different driver than expected.
- For this reason, it is *not* recommended that you use `add_drv` with a full pathname. See `kernel(1M)` for more information on the driver search path.

adminvi(1M)

NAME	adminvi – Edit text with restrictions
SYNOPSIS	adminvi <i>filename...</i>
DESCRIPTION	The admin text editor is a modified version of <i>vi</i> that provides a restricted text-editing environment. <i>adminvi</i> provides all the capabilities of <i>vi</i> except that <i>adminvi</i> does not allow the user to execute shell commands or to write any files other than the files specified on the command line.
OPTIONS	<p>Refer to the <i>vi</i>(1) man page for a complete list of options. <i>adminvi</i> modifies the following options:</p> <ul style="list-style-type: none"> -x Heuristic file encryption is not allowed. -C Forced file encryption is not allowed. -L Listing the names of files saved as the result of an editor or system crash is not allowed. -r <i>filename</i> Recovering files saved as the result of an editor or system crash is not allowed. <i>filename</i> A filename must be specified.
USAGE	<p>Refer to the <i>vi</i>(1) man page for a complete usage description.</p> <p><i>adminvi</i> modifies <i>vi</i> commands to prevent use of the <i>!</i> operator and shell metacharacters in filenames given to commands such as <i>:r</i> and <i>:so</i>.</p>
Commands	<p>The actions of these commands are changed:</p> <ul style="list-style-type: none"> <i>:!</i> The command to execute a shell command is not allowed. <i>:C</i> The forced-encryption command is not allowed. <i>:cd, :chdir</i> The change directory command is not allowed. <i>:crypt, :X</i> The heuristic-encryption command is not allowed. <i>:e</i> If the command to change the file being edited specifies a filename other than the filenames that were given on the <i>adminvi</i> command line, the file is edited in read-only mode. <i>:pre</i> The command to preserve the edit buffers is not allowed. <i>:rec</i> The command to recover preserved edit buffers is not allowed. <i>:sh</i> The command to run a shell is not allowed. <i>:w</i> This command accepts only the filenames that were given on the <i>adminvi</i> command line.
ATTRIBUTES	See <i>attributes</i> (5) for descriptions of the following attributes:

adminvi(1M)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWtsu

**SunOS 5.8
Reference Manual
NOTES**

vi(1), attributes(5)

These interfaces are uncommitted. Although they are not expected to change between minor releases of the Trusted Solaris environment, they may.

arp(1M)

NAME	arp – Address resolution display and control
SYNOPSIS	<pre>arp <i>hostname</i> arp -a arp -d <i>hostname</i> arp -f <i>filename</i> arp -s <i>hostname ether_address</i> [<i>temp</i>] [<i>pub</i>] [<i>trail</i>]</pre>
DESCRIPTION	<p>The <code>arp</code> program displays and modifies the Internet-to-Ethernet address translation tables used by the address resolution protocol (see <code>arp(7P)</code>).</p> <p>With no flags, the program displays the current ARP entry for <i>hostname</i>. The host may be specified by name or by number, using Internet dot notation.</p>
OPTIONS	<p>-a Display all the current ARP entries. The definition for the flags in the table are:</p> <p> P Publish; includes IP address for the machine and the addresses that have explicitly been added by the -s option. ARP will respond to ARP requests for this address.</p> <p> S Static; not learned for the ARP protocol.</p> <p> U Unresolved; waiting for ARP response.</p> <p> M Mapping; only used for the multicast entry for '224.0.0.0'.</p> <p>-d Delete an entry for the host called <i>hostname</i>. To succeed, the process must inherit the <code>sys_net_config</code> privilege.</p> <p>-f Read the file named <i>filename</i> and set multiple entries in the ARP tables. See option -s for argument definitions. To succeed, the process must inherit the <code>sys_net_config</code> privilege. Entries in the file should be of the form:</p> <p><i>hostname ether_address</i> [<i>temp</i>] [<i>pub</i>] [<i>trail</i>]</p> <p>-s Create an ARP entry for the host called <i>hostname</i> with the Ethernet address <i>ether_address</i>. The Ethernet address is given as six hexadecimal bytes separated by colons. The entry will be permanent unless the word <i>temp</i> is given in the command. If the word <i>pub</i> is given, the entry will be published. For instance, this system will respond to ARP requests for <i>hostname</i> even though the <i>hostname</i> is not its own. The word <i>trail</i> indicates that trailer encapsulations may be sent to this host. <code>arp -s</code> can be used for a limited form of proxy ARP when a host on one of the directly attached networks is not physically present on the subnet. Another machine can then be configured to respond to ARP requests using <code>arp -s</code>. This is useful in certain SLIP or PPP configurations. To succeed, the process must inherit the <code>sys_net_config</code> privilege.</p>

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES
To run, options -d, -f , and -s need to inherit the sys_net_config privilege.
ifconfig(1M)
arp(7P), attributes(5)

Trusted Solaris 4/01 Reference Manual
SunOS 5.6

atohexlabel(1M)

NAME	atohexlabel – convert a character-coded label to its hexadecimal equivalent				
SYNOPSIS	<pre> /usr/sbin/atohexlabel [character_coded_CMW_label] /usr/sbin/atohexlabel -c [character_coded_clearance] /usr/sbin/atohexlabel -s [character_coded_sensitivity_label] </pre>				
DESCRIPTION	atohexlabel converts a character-coded label of the type specified into its standard, formatted hexadecimal equivalent and writes the result to the standard output file. If no character-coded label is specified, one is read from standard input.				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWtsu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWtsu				
OPTIONS	<p>-c Identifies the character-coded label as a clearance.</p> <p>-s Identifies the character-coded label as a sensitivity label.</p>				
EXIT STATUS	<p>The following exit values are returned:</p> <p>0 On success.</p> <p>1 On failure, and writes diagnostics to the standard error file.</p>				
FILES	<p>/etc/security/tsol/label_encodings</p> <p>The label encodings file contains the classification names, words, constraints, and values for the defined labels of this system.</p>				
Trusted Solaris 8 4/01 Reference Manual	<p>label_encodings(4)</p> <p><i>Trusted Solaris administrator's document set</i></p>				
SunOS 5.8 Reference Manual DIAGNOSTICS	<p>attributes(5)</p> <p>label translation unavailable</p> <p>The label services are currently unavailable either because the label daemon is not running or because the label_encodings file is incorrect or unavailable.</p> <p><i>label</i> is not translatable by this process</p> <p>This process is not allowed to translate <i>label</i>.</p> <p>error in <i>label</i> at position <i>n</i></p> <p><i>label</i> is not a valid label. An error is noted in position <i>n</i> of the string.</p>				

NAME	audit – Control the behavior of the audit daemon					
SYNOPSIS	audit -n -s -t					
DESCRIPTION	<p>The audit command is the general administrator’s interface to maintaining the audit trail. The audit daemon may be notified to read the contents of the audit_control(4) file and re-initialize the current audit directory to the first directory listed in the audit_control file or to open a new audit file in the current audit directory specified in the audit_control file as last read by the audit daemon. The audit daemon may also be signaled to close the audit trail and disable auditing. The audit command must inherit the sys_audit privilege.</p>					
OPTIONS	<p>-n Signal audit daemon to close the current audit file and open a new audit file in the current audit directory.</p> <p>-s Signal audit daemon to read audit control file. The audit daemon stores the information internally.</p> <p>-t Signal audit daemon to close the current audit trail file, disable auditing and die.</p>					
DIAGNOSTICS	<p>The audit command will exit with 0 upon success and a positive integer upon failure.</p>					
FILES	/etc/security/audit_user	File containing user information for system audit daemon.				
	/etc/security/audit_control	File containing information for system audit daemon.				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWcsu</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE					
Availability	SUNWcsu					
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>This functionality is active only if auditing is enabled. By default, auditing is enabled in the Trusted Solaris environment. By default, the machine halts when audit files run out of disk space. The Trusted Solaris environment adds programming interfaces, audit classes, and audit events.</p> <p>This command must be run at ADMIN_HIGH, and must inherit the sys_audit privilege.</p>					
Trusted Solaris 8 4/01 Reference Manual	<p>praudit(1M), audit(2), audit_control(4), audit_user(4)</p> <p>Trusted Solaris Audit Administration Manual</p>					
SunOS 5.8 Reference Manual	<p>attributes(5)</p>					

audit(1M)

NOTES	This command does not modify a process's preselection mask. It only affects which audit directories are used for audit data storage and to specify the minimum size free.
--------------	---

NAME	auditconfig – configure auditing
SYNOPSIS	auditconfig <i>option...</i>
DESCRIPTION	<p>auditconfig provides a command line interface to get and set kernel audit parameters. A process must have the PRIV_SYS_AUDIT, PRIV_PROC_AUDIT_TCB, or PRIV_PROC_AUDIT_APPL privilege in its set of effective privileges to use the -getcond, -getclass, -getpinfo, and -getpolicy options. A process must have the PRIV_SYS_AUDIT privilege in its set of effective privileges to use the -setcond, -setclass, -chkconf, -conf, -setpmask, -setumask, -setsmask, -getfsize, -setfsize, and -setpolicy options.</p>
OPTIONS	<p>-chkconf Check the configuration of kernel audit event to class mappings. If the runtime class mask of a kernel audit event does not match the configured class mask, a mismatch is reported.</p> <p>-conf Configure kernel audit event to class mappings. Runtime class mappings are changed to match those in the audit event to class database file.</p> <p>-getfsize Return the maximum audit file size in bytes and the current size of the audit file in bytes.</p> <p>-setfsize <i>size</i> Set the maximum size of an audit file to <i>size</i> bytes. When the size limit is reached, the audit file is closed and another is started.</p> <p>-getcond Display the kernel audit condition. The condition displayed is the literal string <i>auditing</i> meaning auditing is enabled and turned on (the kernel audit module is constructing and queuing audit records) or <i>noaudit</i> meaning auditing is enabled but turned off (the kernel audit module is not constructing and queuing audit records), or <i>disabled</i> meaning that the audit module has not been enabled. See <i>auditon(2)</i> and <i>auditd(1M)</i> for further information.</p> <p>-setcond[<i>auditing noaudit</i>] Set the kernel audit condition to the <i>condition</i> specified where <i>condition</i> is the literal string <i>auditing</i> indicating auditing should be enabled or <i>noaudit</i> indicating auditing should be disabled.</p> <p>-getclass <i>event</i> Display the preselection mask associated with the specified kernel audit event. <i>event</i> is the kernel event number or event name.</p> <p>-setclass <i>event audit_flag[audit_flag ...]</i> Map the kernel event <i>event</i> to the classes specified by <i>audit_flags</i>. <i>event</i> is an event number or name. An <i>audit_flag</i> is a two-character string representing an audit class. See <i>audit_control(4)</i> for further information.</p>

auditconfig(1M)

- lsevent
Display the currently configured (runtime) kernel and user level audit event information.
- getpinfo *pid*
Display the audit ID, preselection mask, terminal ID and audit session ID for the specified process.
- setpmask *pid flags*
Set the preselection mask of the specified process. *flags* is the text representation of the flags similar to that in `audit_control(4)`.
- setsmask *asid flags*
Set the preselection mask of all processes with the specified audit session ID.
- setumask *audit flags*
Set the preselection mask of all processes with the specified audit ID.
- lspolicy
Display the kernel audit policies with a description of each policy.
- getpolicy
Display the kernel audit policy.
- setpolicy[+|-]*policy_flag[,policy_flag ...]*
Set the kernel audit policy. A policy *policy_flag* is literal strings that denotes an audit policy. A prefix of + adds the policies specified to the current audit policies. A prefix of - removes the policies specified from the current audit policies. The following are the valid policy flag strings (`auditconfig -lspolicy` also lists the current valid audit policy flag strings):
 - acl
Include in the audit data an ACL attribute for each object accessed. Note that regardless of policy, if there is no ACL associated with an object, an attribute will not be generated. This information is not included by default.
 - ahlt
Halt the machine if an asynchronous audit event occurs that cannot be delivered because the audit queue has reached the high-water mark or because there are insufficient resources to construct an audit record. By default, records are dropped and a count is kept of the number of dropped records.
 - arge
Include the `execv(2)` system call environment arguments to the audit record. This information is not included by default.
 - argv
Include the `execv(2)` system call parameter arguments to the audit record. This information is not included by default.

	cnt	Do not suspend processes when audit resources are exhausted. Instead, drop audit records and keep a count of the number of records dropped. By default, process are suspended until audit resources become available.
	group	Include the supplementary group token in audit records. By default, the group token is not included.
	slabel	Include slabels in audit records. This information is included by default.
	passwd	Include as part of the audit record any bad authentication data encountered during a login operation. The default action is not to include the password in the audit record.
	path	Add secondary path tokens to audit record. These are typically the pathnames of dynamically linked shared libraries or command interpreters for shell scripts. By default, they are not included.
	trail	Include the trailer token in every audit record. By default, the trailer token is not included.
	seq	Include the sequence token as part of every audit record. By default, the sequence token is not included. The sequence token attaches a sequence number to every audit record.
	windata_down	Include in an audit record any downgraded data moved between windows. By default, this information is not included.
	windata_up	Include in an audit record any upgraded data moved between windows. By default, this information is not included.
EXAMPLES	EXAMPLE 1 A sample auditconfig program <pre> # # map kernel audit event number 10 to the "fr" audit class # % auditconfig -setclass 10 fr # # turn on inclusion of exec arguments in exec audit records # % auditconfig -setpolicy +argv </pre>	
EXIT STATUS	0	Successful completion.
	1	An error occurred.

auditconfig(1M)

FILES /etc/security/audit_event Audit event definition and class mappings.
 /etc/security/audit_class Audit class definitions.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES This functionality is active only if auditing is enabled. By default, auditing is enabled in the Trusted Solaris environment. By default, the machine halts when audit files run out of disk space. The Trusted Solaris environment adds programming interfaces, audit classes, and audit events.

These policy flags have been added to the Trusted Solaris auditing module: acl, ahl, slabel, passwd, windata_down, and windata_up.

A process must have the PRIV_SYS_AUDIT, PRIV_PROC_AUDIT_TCB, or PRIV_PROC_AUDIT_APPL privilege in its set of effective privileges to use the -getcond, -getclass, -getpinfo, and -getpolicy options. A process must have the PRIV_SYS_AUDIT privilege in its set of effective privileges to use the -setcond, -setclass, -chkconf, -conf, -setpmask, -setumask, -setsmask, -getfsize, -setfsize, and -setpolicy options.

Trusted Solaris 8 4/01 Reference Manual auditd(1M), praudit(1M), auditon(2), execv(2), audit_class(4),
 audit_control(4), audit_event(4)

Trusted Solaris Audit Administration Manual

SunOS 5.8 Reference Manual attributes(5)

NAME	auditd – Audit daemon
SYNOPSIS	<code>/usr/sbin/auditd</code>
DESCRIPTION	<p>The audit daemon controls the generation and location of audit trail files. If auditing is desired, auditd reads the <code>audit_control(4)</code> file to get a list of directories into which audit files can be written and the percentage limit for how much space to reserve on each filesystem before changing to the next directory.</p> <p>If auditd receives the signal <code>SIGUSR1</code>, the current audit file is closed and another is opened. If <code>SIGHUP</code> is received, the current audit trail is closed, the <code>audit_control</code> file reread, and a new trail is opened. If <code>SIGTERM</code> is received, the audit trail is closed and auditing is terminated. The program <code>audit(1M)</code> sends these signals and is recommended for this purpose.</p> <p>Each time the audit daemon opens a new audit trail file, it updates the file <code>audit_data(4)</code> to include the correct name.</p>
Auditing Conditions	<p>The audit daemon invokes the program <code>audit_warn(1M)</code> under the following conditions with the indicated options:</p> <p><code>audit_warn soft pathname</code> The file system upon which <i>pathname</i> resides has exceeded the minimum free space limit defined in <code>audit_control(4)</code>. A new audit trail has been opened on another file system.</p> <p><code>audit_warn allsoft</code> All available file systems have been filled beyond the minimum free space limit. A new audit trail has been opened anyway.</p> <p><code>audit_warn hard pathname</code> The file system upon which <i>pathname</i> resides has filled or for some reason become unavailable. A new audit trail has been opened on another file system.</p> <p><code>audit_warn allhard count</code> All available file systems have been filled or for some reason become unavailable. The audit daemon will repeat this call to <code>audit_warn</code> every twenty seconds until space becomes available. <i>count</i> is the number of times that <code>audit_warn</code> has been called since the problem arose.</p> <p><code>audit_warn ebusy</code> There is already an audit daemon running.</p> <p><code>audit_warn tmpfile</code> The file <code>/etc/security/audit/audit_tmp</code> exists, indicating a fatal error.</p> <p><code>audit_warn nostart</code> The internal system audit condition is <code>AUC_FCHDONE</code>. Auditing cannot be started without rebooting the system.</p>

auditd(1M)

audit_warn auditoff

The internal system audit condition has been changed to not be AUC_AUDITING by someone other than the audit daemon. This causes the audit daemon to exit.

audit_warn postsigterm

An error occurred during the orderly shutdown of the auditing system.

audit_warn getacdir

There is a problem getting the directory list from /etc/security/audit/audit_control.

The audit daemon will hang in a sleep loop until this file is fixed.

FILES

/etc/security/audit/audit_control

File containing information for system audit daemon.

/etc/security/audit/audit_data

File containing current information on audit daemon.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

This functionality is active only if auditing is enabled. By default, auditing is enabled in the Trusted Solaris environment. By default, the machine halts when audit files run out of disk space. The Trusted Solaris environment adds programming interfaces, audit classes, and audit events.

auditd reads the audit_control(4) file rather than the audit_control(4) file and updates the audit_data(4) file rather than the audit_data(4) file.

Trusted Solaris 8 4/01 Reference Manual

audit(1M), audit_warn(1M), praudit(1M), auditon(2), auditsvc(2), audit.log(4), audit_control(4), audit_data(4)

Trusted Solaris Audit Administration Manual

SunOS 5.8 Reference Manual

attributes(5)

NAME	auditreduce – merge and select audit records from audit trail files
SYNOPSIS	auditreduce [<i>options</i>] [<i>audit-trail-file...</i>]
DESCRIPTION	<p>auditreduce allows you to select or merge records from audit trail files. Audit files may be from one or more machines.</p> <p>The merge function merges together audit records from one or more input audit trail files into a single output file. The records in an audit trail file are assumed to be sorted in chronological order (oldest first) and this order is maintained by auditreduce in the output file.</p> <p>Unless instructed otherwise, auditreduce will merge the entire audit trail, which consists of all the audit trail files in the directory structure <i>audit_root_dir</i>/*/files (see <i>audit_control</i>(4) for details of the structure of the audit root). Unless stated with the -R or -S option, <i>audit_root_dir</i> defaults to <i>/etc/security/audit</i>. By using the file selection options it is possible to select some subset of these files, or files from another directory, or files named explicitly on the command line.</p> <p>The select function allows audit records to be selected on the basis of numerous criteria relating to the record's content (see <i>audit.log</i>(4) for details of record content). A record must meet all of the <i>record-selection-option</i> criteria to be selected.</p>
Audit Trail Filename Format	<p>Any audit trail file not named on the command line must conform to the audit trail filename format. Files produced by the audit system already have this format. Output file names produced by auditreduce are in this format. It is:</p> <p><i>start-time . end-time . suffix</i></p> <p>where <i>start-time</i> is the 14-character timestamp of when the file was opened, <i>end-time</i> is the 14-character timestamp of when the file was closed, and <i>suffix</i> is the name of the machine which generated the audit trail file, or some other meaningful suffix (e.g., <i>all</i>, if the file contains a combined group of records from many machines). The <i>end-time</i> may be the literal string <i>not_terminated</i>, to indicate that the file is still being written to by the audit system. Timestamps are of the form <i>yyyymmddhhmmss</i> (year, month, day, hour, minute, second). The timestamps are in Greenwich Mean Time (GMT).</p>
File Selection Options	<p>The file selection options indicate which files are to be processed and certain types of special treatment.</p> <p>-A</p> <p>All of the records from the input files will be selected regardless of their timestamp. This option effectively disables the -a, -b, and -d options. This is useful in preventing the loss of records if the -D option is used to delete the input files after they are processed. Note, however, that if a record is <i>not</i> selected due to another option, then -A will not override that.</p>

auditreduce(1M)

-C

Only process complete files. Files whose filename *end-time* timestamp is *not_terminated* are not processed (such a file is currently being written to by the audit system). This is useful in preventing the loss of records if -D is used to delete the input files after they are processed. It does not apply to files specified on the command line.

-D *suffix*

Delete input files after they are processed. The files are only deleted if the entire run is successful. If *auditreduce* detects an error while reading a file, then that file is not deleted. If -D is specified, -A, -C and -O are also implied. *suffix* is given to the -O option. This helps prevent the loss of audit records by ensuring that all of the records are written, only complete files are processed, and the records are written to a file before being deleted. Note that if both -D and -O are specified in the command line, the order of specification is significant. The *suffix* associated with the latter specification is in effect.

-M *machine*

Allows selection of records from files with *machine* as the filename suffix. If -M is not specified, all files are processed regardless of suffix. -M can also be used to allow selection of records from files that contain combined records from many machines and have a common suffix (such as *all*).

-N

Select objects in *new mode*. This flag is off by default, thus retaining backward compatibility. In the existing, *old mode*, specifying the -e, -f, -g, -r, or -u flags would select not only actions taken with those IDs, but also certain objects owned by those IDs. When running in *new mode*, only actions are selected. In order to select objects, the -o option must be used.

-O *suffix*

Direct output stream to a file in the current *audit_root_dir* with the indicated suffix. *suffix* may alternatively contain a full pathname, in which case the last component is taken as the suffix, ahead of which the timestamps will be placed, ahead of which the remainder of the pathname will be placed. If the -O option is not specified, the output is sent to the standard output. When *auditreduce* places timestamps in the filename, it uses the times of the first and last records in the merge as the *start-time* and *end-time*.

-Q

Quiet. Suppress notification about errors with input files.

-R *pathname*

Specify the pathname of an alternate audit root directory *audit_root_dir* to be *pathname*. Therefore, rather than using */etc/security/audit/*/files* by default, *pathname/*/files* will be examined instead.

-S *server*

This option causes *auditreduce* to read audit trail files from a specific location (server directory). *server* is normally interpreted as the name of a subdirectory of

Record Selection Options

the audit root, therefore `auditreduce` will look in `audit_root_dir/server/files` for the audit trail files. But if `server` contains any `'/'` characters, it is the name of a specific directory not necessarily contained in the audit root. In this case, `server/files` will be consulted. This option allows archived files to be manipulated easily, without requiring that they be physically located in a directory structure like that of `/etc/security/audit`.

-V

Verbose. Display the name of each file as it is opened, and how many records total were written to the output stream.

The record selection options listed below are used to indicate which records are written to the output file produced by `auditreduce`.

Multiple arguments of the same type are not permitted.

-a *date-time*

Select records that occurred at or after *date-time*. The *date-time* argument is described under Option Arguments, below. *date-time* is in local time. The **-a** and **-b** options can be used together to form a range.

-b *date-time*

Select records that occurred before *date-time*.

-c *audit-classes*

Select records by audit class. Records with events that are mapped to the audit classes specified by *audit-classes* are selected. Audit class names are defined in `audit_class(4)`. The *audit-classes* can be a comma separated list of audit *flags* like those described in `audit_control(4)`. Using the audit *flags*, one can select records based upon success and failure criteria.

-d *date-time*

Select records that occurred on a specific day (a 24-hour period beginning at 00:00:00 of the day specified and ending at 23:59:59). The day specified is in local time. The time portion of the argument, if supplied, is ignored. Any records with timestamps during that day are selected. If any hours, minutes, or seconds are given in *time*, they are ignored. **-d** can not be used with **-a** or **-b**.

-e *effective-user*

Select records with the specified *effective-user*.

-f *effective-group*

Select records with the specified *effective-group*.

-g *real-group*

Select records with the specified *real-group*.

-j *subject-ID*

Select records with the specified *subject-ID* where *subject-ID* is a process ID.

auditreduce(1M)

-m *event*

Select records with the indicated *event*. The *event* is the literal string or the *event* number.

-o *object_type=objectID_value*

Select records by object type. A match occurs when the record contains the information describing the specified *object_type* and the object ID equals the value specified by *objectID_value*. The allowable object types and values are as follows:

<i>file=pathname</i>	Select records containing file system objects with the specified <i>pathname</i> , where <i>pathname</i> is a comma separated list of regular expressions. If a regular expression is preceded by a tilde (~), files matching the expression are excluded from the output. For example, the option <i>file="/usr/openwin,/usr,/etc"</i> would select all files in /usr or /etc except those in /usr/openwin. The order of the regular expressions is important because auditreduce processes them from left to right, and stops when a file is known to be either selected or excluded. Thus the option <i>file=/usr,/etc,~/usr/openwin</i> would select all files in /usr and all files in /etc. Files in /usr/openwin are not excluded because the regular expression /usr is matched first. Care should be given in surrounding the <i>pathname</i> with quotes so as to prevent the shell from expanding any tildas.
<i>filegroup=group</i>	Select records containing file system objects with <i>group</i> as the owning group.
<i>fileowner=user</i>	Select records containing file system objects with <i>user</i> as the owning user.
<i>msgqid=ID</i>	Select records containing message queue objects with the specified <i>ID</i> where <i>ID</i> is a message queue ID.
<i>msgqgroup=group</i>	Select records containing message queue objects with <i>group</i> as the owning or creating group.
<i>msgqowner=user</i>	Select records containing message queue objects with <i>user</i> as the owning or creating user.
<i>pid=ID</i>	Select records containing process objects with the specified <i>ID</i> where <i>ID</i> is a process ID. Process are objects when they are receivers of signals.
<i>procgrou=group</i>	Select records containing process objects with <i>group</i> as the real or effective group.
<i>procowner=user</i>	Select records containing process objects with <i>user</i> as the real or effective user.
<i>semid=ID</i>	Select records containing semaphore objects with the specified <i>ID</i> where <i>ID</i> is a semaphore ID.

auditreduce(1M)

`semgroup=group` Select records containing semaphore objects with *group* as the owning or creating group.

`semowner=user` Select records containing semaphore objects with *user* as the owning or creating user.

`shmid=ID` Select records containing shared memory objects with the specified *ID* where *ID* is a shared memory ID.

`shmgroup=group` Select records containing shared memory objects with *group* as the owning or creating group.

`shmowner=user` Select records containing shared memory objects with *user* as the owning or creating user.

`sock=port_number | machine`

Select records containing socket objects with the specified *port_number* or the specified *machine* where *machine* is a machine name as defined in `hosts(4)`.

`-r real-user`

Select records with the specified *real-user*.

`-s sensitivity-label`

Select records with the specified *sensitivity-label*, which may be a range as explained under Option Arguments, *sensitivity-label*.

`-u audit-user`

Select records with the specified *audit-user*. When one or more *filename* arguments appear on the command line, only the named files are processed. Files specified in this way need not conform to the audit trail filename format. However, `-M`, `-S`, and `-R` may not be used when processing named files. If the *filename* is “-” then the input is taken from the standard input.

Option Arguments

audit-trail-file

An audit trail file as defined in `audit.log(4)`. An audit trail file not named on the command line must conform to the audit trail file name format. Audit trail files produced as output of `auditreduce` are in this format as well. The format is:

`start-time . end-time . suffix`

start-time is the 14 character time stamp denoting when the file was opened.

end-time is the 14 character time stamp denoting when the file was closed. *end-time* may also be the literal string `not_terminated`, indicating the file is still be written to by the audit daemon or the file was not closed properly (a system crash or abrupt halt occurred). *suffix* is the name of the machine that generated the audit trail file (or some other meaningful suffix; e.g. `all` would be a good suffix if the audit trail file contains a combined group of records from many machines).

date-time

The *date-time* argument to `-a`, `-b`, and `-d` can be of two forms: An absolute *date-time* takes the form:

auditreduce(1M)

yyyymmdd [*hh* [*mm* [*ss*]]]

where *yyyy* specifies a year (with 1970 as the earliest value), *mm* is the month (01-12), *dd* is the day (01-31), *hh* is the hour (00-23), *mm* is the minute (00-59), and *ss* is the second (00-59). The default is 00 for *hh*, *mm* and *ss*.

An offset can be specified as: *+n d|h|m|s* where *n* is a number of units, and the tags *d*, *h*, *m*, and *s* stand for days, hours, minutes and seconds, respectively. An offset is relative to the starting time. Thus, this form can only be used with the *-b* option.

event

The literal string or ordinal event number as found in `audit_event(4)`. If *event* is not found in the `audit_event` file it is considered invalid.

group

The literal string or ordinal group ID number as found in `group(4)`. If *group* is not found in the `group` file it is considered invalid. *group* may be negative.

pathname

A regular expression describing a pathname.

sensitivity-label

The literal string representation of an sensitivity label or a range of two valid sensitivity labels. To specify a range, use *[x] : [y]* where *x* and *y* are valid sensitivity labels. Only those records that are fully bounded by *x* and *y* will be selected. If *x* or *y* is omitted, the default uses `ADMIN_LOW` or `ADMIN_HIGH` respectively.

user

The literal username or ordinal user ID number as found in `passwd(4)`. If the username is not found in the `passwd` file it is considered invalid. *user* may be negative.

EXAMPLES

EXAMPLE 1 The auditreduce command.

`praudit(1M)` is available to display audit records in a human-readable form.

This will display the entire audit trail in a human-readable form:

```
% auditreduce | praudit
```

If all the audit trail files are being combined into one large file, then deleting the original files could be desirable to prevent the records from appearing twice:

```
% auditreduce -V -d /etc/security/audit/combined/all
```

This will print what user `milner` did on April 13, 1988. The output will be displayed in a human-readable form to the standard output:

```
% auditreduce -d 19880413 -u milner | praudit
```


EXAMPLE 1 The auditreduce command. (Continued)

The above example may produce a large volume of data if `milner` has been busy. Perhaps looking at only login and logout times would be simpler. The `-c` option will select records from a specified class:

```
% auditreduce -d 19880413 -u milner -c lo | praudit
```

To see `milner`'s login/logout activity for April 13, 14, and 15 the following is used. The results are saved to a file in the current working directory. Note that the name of the output file will have `milnerlo` as the *suffix*, with the appropriate timestamp prefixes. Note that the long form of the name is used for the `-c` option:

```
% auditreduce -a 19880413 -b +3d -u milner -c login_logout -o milnerlo
```

To follow `milner`'s movement about the file system on April 13, 14, and 15 the `chdir` record types could be viewed. Note that in order to get the same time range as the above example we needed to specify the `-b` time as the day after our range. This is because 19880416 defaults to midnight of that day, and records before that fall on 0415, the end-day of the range.

```
% auditreduce -a 19880413 -b 19880416 -u milner -m AUE_CHDIR | praudit
```

In this example the audit records are being collected in summary form (the login/logout records only). The records are being written to a summary file in a different directory than the normal audit root to prevent the selected records from existing twice in the audit root.

```
% auditreduce -d 19880330 -c lo -o /etc/security/audit_summary/logins
```

If activity for user ID 9944 has been observed, but that user is not known to the system administrator, then the following example will search the entire audit trail for any records generated by that user. `auditreduce` will query the system as to the current validity of ID 9944, and print a warning message if it is not currently active:

```
% auditreduce -o /etc/security/audit_suspect/user9944 -u 9944
```

FILES /etc/security/audit/server/files/*
Location of audit trails, when stored.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

The functionality described in this man page is available only if auditing is enabled. By default, auditing is enabled in the Trusted Solaris environment.

auditreduce(1M)

	<p>The Trusted Solaris environment has added the <i>-s sensitivity-label</i> record selection option to this command.</p>
Trusted Solaris 8 4/01 Reference Manual	<p>praudit(1M), audit.log(4), audit_class(4), audit_control(4)</p> <p><i>Trusted Solaris Audit Administration Manual</i></p>
SunOS 5.8 Reference Manual DIAGNOSTICS	<p>group(4), hosts(4), passwd(4), attributes(5)</p> <p>auditreduce will print out error messages if there are command line errors and then exit. If there are fatal errors during the run auditreduce will print an explanatory message and exit. In this case the output file may be in an inconsistent state (no trailer or partially written record) and auditreduce will print a warning message before exiting. Successful invocation returns 0 and unsuccessful invocation returns 1.</p> <p>Since auditreduce may be processing a large number of input files, it is possible that the machine-wide limit on open files will be exceeded. If this happens, auditreduce will print a message to that effect, give information on how many file there are, and exit.</p> <p>If auditreduce prints a record's timestamp in a diagnostic message, that time is in local time. However, when filenames are displayed, their timestamps are in GMT.</p>
BUGS	<p>Conjunction, disjunction, negation, and grouping of record selection options should be allowed.</p>

NAME	audit_startup – Audit subsystem initialization script
SYNOPSIS	/etc/security/audit_startup
DESCRIPTION	The audit_startup script is used to initialize the audit subsystem before the audit daemon is started. This script is configurable by the security administrator, and currently consists of a series of auditconfig(1M) commands to set the system default policy, and download the initial event to class mapping.
SUMMARY OF TRUSTED SOLARIS CHANGES	By default, the audit module is enabled in the Trusted Solaris environment. By default, the machine halts when audit files run out of disk space. The Trusted Solaris environment adds programming interfaces, audit classes, and audit events.
Trusted Solaris 8 4/01 Reference Manual	auditconfig(1M), auditd(1M)
Solaris 8 4/01 Reference Manual	attributes(5)

auditstat(1M)

NAME	auditstat – Display kernel audit statistics	
SYNOPSIS	auditstat [-c <i>count</i>] [-h <i>numlines</i>] [-i <i>interval</i>] [-n] [-v]	
DESCRIPTION	auditstat displays kernel audit statistics. To succeed, it must inherit the sys_audit privilege. The fields displayed are as follows:	
	aud	The total number of audit records processed by the audit(2) system call.
	ctl	This field is obsolete.
	drop	The total number of audit records that have been dropped. Records are dropped according to the kernel audit policy. See auditon(2), AUDIT_CNT policy for details.
	enq	The total number of audit records put on the kernel audit queue.
	gen	The total number of audit records that have been constructed (not the number written).
	kern	The total number of audit records produced by user processes (as a result of system calls).
	mem	The total number of Kbytes of memory currently in use by the kernel audit module.
	nona	The total number of non-attributable audit records that have been constructed. These are audit records that are not attributable to any particular user.
	rblk	The total number of times that auditsvc(2) has blocked waiting to process audit data.
	tot	The total number of Kbytes of audit data written to the audit trail.
	wblk	The total number of times that user processes blocked on the audit queue at the high water mark.
	wrtn	The total number of audit records written. The difference between enq and wrtn is the number of outstanding audit records on the audit queue that have not been written.
	OPTIONS	-c <i>count</i>
-h <i>numlines</i>		Display a header for every <i>numlines</i> of statistics printed. The default is to display the header every 20 lines. If <i>numlines</i> is equal to zero, the header is never displayed.
-i <i>interval</i>		Display the statistics every <i>interval</i> where <i>interval</i> is the number of seconds to sleep between each collection.
-n		Display the number of kernel audit events currently configured.
-v		Display the version number of the kernel audit module software.

EXIT STATUS

0	Successful completion.
1	An error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

The functionality described in this man page is available only if auditing is enabled. By default, auditing is enabled in the Trusted Solaris environment. To succeed, this command must have the `sys_audit` privilege.

Trusted Solaris 8 4/01 Reference Manual

auditconfig(1M), praudit(1M), audit(2), auditon(2), auditsvc(2)

attributes(5)

audit_warn(1M)

NAME	audit_warn – Audit daemon warning script	
SYNOPSIS	/etc/security/audit_warn [<i>option</i> [<i>arguments</i>]]	
DESCRIPTION	<p>The audit_warn script processes warning or error messages from the audit daemon. When a problem is encountered, the audit daemon, auditd(1M) calls audit_warn with the appropriate arguments. The <i>option</i> argument specifies the error type.</p> <p>The system administrator can specify a list of mail recipients to be notified when an audit_warn situation arises by defining a mail alias called audit_warn in aliases(4). The users that make up the audit_warn alias are typically the administrative roles.</p>	
OPTIONS	allhard <i>count</i>	Indicates that the hard limit for all filesystems has been exceeded <i>count</i> times. The default action for this option is to send mail to the audit_warn alias only if the <i>count</i> is 1, and to write a message to the machine console every time. It is recommended that mail <i>not</i> be sent every time as this could result in a the saturation of the file system that contains the mail spool directory.
	allsoft	Indicates that the soft limit for all filesystems has been exceeded. The default action for this option is to send mail to the audit_warn alias and to write a message to the machine console.
	auditoff	Indicates that someone other than the audit daemon changed the system audit state to something other than AUC_AUDITING. The audit daemon will have exited in this case. The default action for this option is to send mail to the audit_warn alias and to write a message to the machine console.
	ebusy	Indicates that the audit daemon is already running. The default action for this option is to send mail to the audit_warn alias and to write a message to the machine console.
	getacdir <i>count</i>	Indicates that there is a problem getting the directory list from audit_control(4). The audit daemon will hang in a sleep loop until the file is fixed. The default action for this option is to send mail to the audit_warn alias only if <i>count</i> is 1, and to write a message to the machine console every time. It is recommended that mail <i>not</i> be sent every time as this could result in a the saturation of the file system that contains the mail spool directory.
	hard <i>filename</i>	Indicates that the hard limit for the file has been exceeded. The default action for this option is to send mail to the audit_warn alias and to write a message to the machine console.
	nostart	Indicates that auditing could not be started. The default action for this option is to send mail to the audit_warn alias and to write a message to the machine console. Some administrators may prefer to modify audit_warn to reboot the system when this error occurs.

audit_warn(1M)

postsigterm	Indicates that an error occurred during the orderly shutdown of the audit daemon. The default action for this option is to send mail to the audit_warn alias and to write a message to the machine console.
soft <i>filename</i>	Indicates that the soft limit for <i>filename</i> has been exceeded. The default action for this option is to send mail to the audit_warn alias and to write a message to the machine console.
tmpfile	Indicates that the temporary audit file already exists indicating a fatal error. The default action for this option is to send mail to the audit_warn alias and to write a message to the machine console.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsr

SUMMARY OF TRUSTED SOLARIS CHANGES The functionality described in this man page is available only if auditing is enabled. By default, auditing is enabled in the Trusted Solaris environment. By default, the machine halts when audit files run out of disk space. The Trusted Solaris environment adds programming interfaces, audit classes, and audit events.

Trusted Solaris 8 4/01 Reference Manual audit(1M), auditd(1M), audit.log(4), audit_control(4)

Trusted Solaris Audit Administration Manual

SunOS 5.8 Reference Manual aliases(4), attributes(5)

automount(1M)

NAME	automount – Install automatic mount points
SYNOPSIS	<code>/usr/sbin/automount [-t <i>duration</i>] [-v]</code>
DESCRIPTION	<p>The automount utility installs <code>autofs</code> mount points and associates an automount map with each mount point. The <code>autofs</code> file system monitors attempts to access directories within it and notifies the <code>automountd(1M)</code> daemon. The daemon uses the map to locate a file system, which it then mounts at the point of reference within the <code>autofs</code> file system. A map can be assigned to an <code>autofs</code> mount using an entry in the <code>/etc/auto_master</code> map or a direct map.</p> <p>If the file system is not accessed within an appropriate interval (10 minutes by default), the <code>automountd</code> daemon unmounts the file system.</p> <p>The file <code>/etc/auto_master</code> determines the locations of all <code>autofs</code> mount points. By default, this file contains four entries:</p> <pre># Master map for automounter # +auto_master /net -hosts -nosuid /home auto_home /xfn -xfn</pre> <p>The <code>+auto_master</code> entry is a reference to an external NIS or NIS+ master map. If one exists, then its entries are read as if they occurred in place of the <code>+auto_master</code> entry. The remaining entries in the master file specify a directory on which an <code>autofs</code> mount will be made followed by the automounter map to be associated with it. Optional mount options may be supplied as an optional third field in the each entry. These options are used for any entries in the map that do not specify mount options explicitly. Security attributes may also follow the automounter map name. These consist of a semicolon separated list of security attributes to be associated with the map. See <code>mount(1M)</code> for a description of these security attributes. As with mount options, security attributes in <code>/etc/auto_master</code> are used for any entries in the map that do not specify security attributes explicitly. The security attribute list must be preceded by a <code>-o</code> flag to distinguish it from mount options. The <code>automount</code> command is usually run without arguments. It compares the entries <code>/etc/auto_master</code> with the current list of <code>autofs</code> mounts in <code>/etc/mnttab</code> and adds, removes or updates <code>autofs</code> mounts to bring the <code>/etc/mnttab</code> up to date with the <code>/etc/auto_master</code>. At boot time it installs all <code>autofs</code> mounts from the master map. Subsequently, it may be run to install <code>autofs</code> mounts for new entries in the master map or the direct map, or to perform unmounts for entries that have been removed from these maps.</p> <p>Note – <code>automount -S allowed=all</code> can be written as the ordinary mount option <code>-o allowed=all</code>. This form will be discarded by base Solaris automounts, hence both base Solaris and Trusted Solaris can share name service databases containing such security attributes.</p>
OPTIONS	The following options are supported:

automount(1M)

- t *duration* Specifies a *duration*, in seconds, that a file system is to remain mounted when not in use. The default is 10 minutes.
- v Verbose mode. Notifies of `autofs` mounts, unmounts, or other non-essential information.

Map Entry Format

A simple map entry (mapping) takes the form:

```
key [ -mount-options ] [ -oattribute-list ] location . . .
```

where *key* is the full pathname of the directory to mount when used in a direct map, or the simple name of a subdirectory in an indirect map. *mount-options* is a comma-separated list of mount options, and *location* specifies a file system from which the directory may be mounted. In the case of a simple NFS mount, the options that can be used are as specified in `mount_nfs(1M)`, and *location* takes the form:

```
host: pathname
```

host is the name of the host from which to mount the file system, and *pathname* is the absolute pathname of the directory to mount.

Options to other file systems are documented on the other `mount_*` reference manual pages, for example, `mount_cacheefs(1M)`.

Replicated File Systems

Multiple *location* fields can be specified for replicated NFS file systems, in which case `automount` and the kernel will each try to use that information to increase availability. If the read-only flag is set in the map entry, `automount` mounts a list of locations that the kernel may use, sorted by several criteria. When a server does not respond, the kernel will switch to an alternate server. The sort ordering of `automount` is used to determine how the next server is chosen. If the read-only flag is not set, `automount` will mount the best single location, chosen by the same sort ordering, and new servers will only be chosen when an unmount has been possible, and a remount is done. Servers on the same local subnet are given the strongest preference, and servers on the local net are given the second strongest preference. Among servers equally far away, response times will determine the order if no weighting factors (see below) are used.

If the list includes server locations using both the NFS Version 2 Protocol and the NFS Version 3 Protocol, `automount` will choose only a subset of the server locations on the list, so that all entries will be the same protocol. It will choose servers with the NFS Version 3 Protocol so long as an NFS Version 2 Protocol server on a local subnet will not be ignored. See the *NFS Administration Guide* for additional details.

If each *location* in the list shares the same *pathname* then a single *location* may be used with a comma-separated list of hostnames:

```
hostname,hostname . . . : pathname
```

automount(1M)

Requests for a server may be weighted, with the weighting factor appended to the server name as an integer in parentheses. Servers without a weighting are assumed to have a value of zero (most likely to be selected). Progressively higher values decrease the chance of being selected. In the example,

```
man -ro alpha,bravo,charlie(1),delta(4) : /usr/man
```

hosts alpha and bravo have the highest priority; host delta has the lowest.

Server proximity takes priority in the selection process. In the example above, if the server delta is on the same network segment as the client, but the others are on different network segments, then delta will be selected; the weighting value is ignored. The weighting has effect only when selecting between servers with the same network proximity.

In cases where each server has a different export point, the weighting can still be applied. For example:

```
man -ro alpha : /usr/man   bravo,charlie(1) : /usr/share/man \
delta(3) : /export/man
```

A mapping can be continued across input lines by escaping the NEWLINE with a backslash (\). Comments begin with a number sign (#) and end at the subsequent NEWLINE.

Map Key Substitution

The ampersand (&) character is expanded to the value of the key field for the entry in which it occurs. In this case:

```
jane sparcserver : /home/&
```

the & expands to jane.

Wildcard Key

The asterisk (*) character, when supplied as the key field, is recognized as the catch-all entry. Such an entry will match any key not previously matched. For instance, if the following entry appeared in the indirect map for /config:

```
*          & : /export/config/&
```

this would allow automatic mounts in /config of any remote file system whose location could be specified as:

```
hostname : /export/config/hostname
```

Variable Substitution

Client-specific variables can be used within an automount map. For instance, if \$HOST appeared within a map, automount would expand it to its current value for the client's host name. Supported variables are:

ARCH	The application architecture is derived from the output of <code>uname -m</code>	The architecture name. For example, "sun4" on a sun4u machine.
CPU	The output of <code>uname -p</code>	The processor type. For example, "sparc"
HOST	The output of <code>uname -n</code>	The host name. For example, "biggles"
OSNAME	The output of <code>uname -s</code>	The OS name. For example, "SunOS"
OSREL	The output of <code>uname -r</code>	The OS release name. For example "5.7"
OSVERS	The output of <code>uname -v</code>	The OS version. For example, "beta1.0"
NATISA	The output of <code>isainfo -n</code>	The native instruction set architecture for the system. For example, "sparcv9"

If a reference needs to be protected from affixed characters, you can surround the variable name with curly braces (`{ }`).

Multiple Mounts

A multiple mount entry takes the form:

```
key [-mount-options] [ [mountpoint] [-mount-options] location . . . ] . . .
```

The initial `/[mountpoint]` is optional for the first mount and mandatory for all subsequent mounts. The optional `mountpoint` is taken as a pathname relative to the directory named by `key`. If `mountpoint` is omitted in the first occurrence, a `mountpoint` of `/` (root) is implied.

Given an entry in the indirect map for `/src`

```
BETA -RO \  
/SVR1,SVR2 : /EXPORT/SRC/BETA\  
/1.0SVR1,SVR2 : /EXPORT/SRC/BETA/1.0\  
/1.0/MANSVR1,SVR2 : /EXPORT/SRC/BETA/1.0/MAN
```

All offsets must exist on the server under `beta`. `automount` will automatically mount `/src/beta`, `/src/beta/1.0`, and `/src/beta/1.0/man`, as needed, from either `svr1` or `svr2`, whichever host is nearest and responds first.

Other File System Types

The automounter assumes NFS mounts as a default file system type. Other file system types can be described using the `fstype` mount option. Other mount options specific

automount(1M)

to this file system type can be combined with the `fstype` option. The location field must contain information specific to the file system type. If the location field begins with a slash, a colon character must be prepended, for instance, to mount a CD file system:

```
cdrom -fstype=hsfs,ro : /dev/sr0
```

or to perform an `autofs` mount:

```
src -fstype=autofs auto_src
```

Note: Use this procedure only if you are not using Volume Manager.

Mounts using CacheFS are most useful when applied to an entire map as map defaults. The following entry in the master map describes cached home directory mounts. It assumes the default location of the cache directory, `/cache`.

```
/home auto_home -fstype=cachefs,backfstype=nfs
```

See the NOTES section for information on option inheritance.

Indirect Maps

An indirect map allows you to specify mappings for the subdirectories you wish to mount under the `directory` indicated on the command line. In an indirect map, each key consists of a simple name that refers to one or more file systems that are to be mounted as needed.

Direct Maps

Entries in a direct map are associated directly with `autofs` mount points. Each key is the full pathname of an `autofs` mount point. The direct map as a whole is not associated with any single directory.

Included Maps

The contents of another map can be included within a map with an entry of the form

```
+mapname
```

If *mapname* begins with a slash, it is assumed to be the pathname of a local file. Otherwise, the location of the map is determined by the policy of the name service switch according to the entry for the automounter in `/etc/nsswitch.conf`, such as

```
automount: files nis
```

If the name service is `files`, then the name is assumed to be that of a local file in `/etc`. If the key being searched for is not found in the included map, the search continues with the next entry.

Special Maps

There are three special maps available: `-hosts`, `-xfs`, and `-null`. The `-hosts` map is used with the `/net` directory and assumes that the map key is the hostname of an NFS server. The `automountd` daemon dynamically constructs a map entry from the server's list of exported file systems. For instance, a reference to

	<p>/net/hermes/usr would initiate an automatic mount of all exported file systems from hermes that are mountable by the client. References to a directory under /net/hermes will refer to the corresponding directory relative to hermes root.</p> <p>The -xfn map is used to mount the initial context of the Federated Naming Service (FNS) namespace under the /xfn directory. For more information on FNS, see <code>fns(5)</code>, <code>fns_initial_context(5)</code>, <code>fns_policies(5)</code>, and the Federated Naming Service Guide.</p> <p>The -null map, when indicated on the command line, cancels a previous map for the directory indicated. This is most useful in the /etc/auto_master for cancelling entries that would otherwise be inherited from the +auto_master include entry. To be effective, the -null entries must be inserted before the included map entry.</p>				
Executable Maps	<p>Local maps that have the execute bit set in their file permissions will be executed by the automounter and provided with a key to be looked up as an argument. The map generation program is executed with no inherited privileges.</p> <p>The executable map is expected to return the content of an automounter map entry on its stdout or no output if the entry cannot be determined. A direct map cannot be made executable.</p>				
Configuration and the auto_master Map	<p>When initiated without arguments, automount consults the master map for a list of autofs mount points and their maps. It mounts any autofs mounts that are not already mounted, and unmounts autofs mounts that have been removed from the master map or direct map.</p> <p>The master map is assumed to be called auto_master and its location is determined by the name service switch policy. Normally the master map is located initially as a local file /etc/auto_master.</p>				
Browsing	<p>The Solaris 7 release supports browsability of indirect maps. This allows all of the potential mount points to be visible, whether or not they are mounted. The -nobrowse option can be added to any indirect autofs map to disable browsing. For example:</p> <pre> /net -hosts -nosuid,nobrowse /home auto_home </pre> <p>In this case, any <i>hostnames</i> would only be visible in /net after they are mounted, but all potential mount points would be visible under /home. The -browse option enables browsability of autofs file systems. This is the default for all indirect maps.</p>				
EXIT STATUS	<p>The following exit values are returned:</p> <table> <tr> <td>0</td><td>Successful completion.</td></tr> <tr> <td>1</td><td>An error occurred.</td></tr> </table>	0	Successful completion.	1	An error occurred.
0	Successful completion.				
1	An error occurred.				
FILES	<table> <tr> <td>/etc/auto_master</td><td>master automount map.</td></tr> </table>	/etc/auto_master	master automount map.		
/etc/auto_master	master automount map.				

automount(1M)

	<p>/etc/auto_home map to support automounted home directories.</p> <p>/etc/nsswitch.conf the name service switch configuration file.</p>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWcsu</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>Security attributes can be specified in <code>auto_master</code> and in <code>autofs</code> map entries with the <code>-o</code> option. If security attributes are not specified in either <code>auto_master</code> or an <code>autofs</code> map entry, but an entry for the mount point is in <code>/etc/security/tsol/vfstab_adjunct</code>, then security attributes in the <code>vfstab_adjunct</code> file are used.</p> <p><code>automount</code> must be started with a process sensitivity label of <code>ADMIN_LOW</code>, and a clearance of <code>ADMIN_HIGH</code>. It must have the <code>PAF_TRUSTED_PATH</code> process attribute, and must inherit the following privileges: <code>file_mac_read</code>, <code>file_mac_write</code>, <code>file_dac_read</code>, <code>file_dac_write</code>, and <code>sys_mount</code>. For executable maps, the map generation program is executed with no inherited privileges.</p>				
Trusted Solaris 8 4/01 Reference Manual	<p><code>uname(1)</code>, <code>automountd(1M)</code>, <code>mount(1M)</code>, <code>mount_nfs(1M)</code>, <code>vfstab_adjunct(4)</code></p> <p><code>isainfo(1)</code>, <code>ls(1)</code>, <code>mount_cachefs(1M)</code>, <code>attributes(5)</code>, <code>fns(5)</code>, <code>fns_initial_context(5)</code>, <code>fns_policies(5)</code></p> <p><i>NFS Administration Guide</i></p>				
NOTES	<p><code>autofs</code> mount points must not be hierarchically related. <code>automount</code> does not allow an <code>autofs</code> mount point to be created within another <code>autofs</code> mount.</p> <p>Since each direct map entry results in a new <code>autofs</code> mount such maps should be kept short.</p> <p>Entries in both direct and indirect maps can be modified at any time. The new information is used when <code>automountd</code> next uses the map entry to do a mount.</p> <p>New entries added to a master map or direct map will not be useful until the <code>automount</code> command is run to install them as new <code>autofs</code> mount points. New entries added to an indirect map may be used immediately.</p> <p>As of the Solaris 7 release, a listing (see <code>ls(1)</code>) of the <code>autofs</code> directory associated with an indirect map shows all potential mountable entries. The attributes associated with the potential mountable entries are temporary. The real file system attributes will only be shown once the file system has been mounted.</p> <p>Default mount options can be assigned to an entire map when specified as an optional third field in the master map. These options apply only to map entries that have no</p>				

automount(1M)

mount options. Note that map entities with options override the default options, as at this time, the options do not concatenate. The concatenation feature is planned for a future release.

When operating on a map that invokes an NFS mount, the default number of retries for the automounter is 0, that is, a single mount attempt, with no retries. Note that this is significantly different from the default (10000) for the `mount_nfs(1M)` utility.

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same.

automountd(1M)

NAME automountd – autofs mount/unmount daemon

SYNOPSIS **automountd** [-Tvn] [-D *name=value*]

DESCRIPTION automountd is an RPC server that answers filesystem mount and unmount requests from the autofs filesystem. It uses local files or name service maps to locate filesystems to be mounted. These maps are described with the automount(1M) command.

The automountd daemon is automatically invoked in run level 2.

OPTIONS

- T Trace. Expand each RPC call and display it on the standard output.
- v Verbose. Log status messages to the console.
- n Turn off browsing for all autofs mount points. This option overrides the -browse autofs map option on the local host.
- D *name=value* Assign *value* to the indicated automount map substitution variable. These assignments cannot be used to substitute variables in the master map auto_master.

USAGE See largefile(5) for the description of the behavior of automountd when encountering files greater than or equal to 2 Gbyte (2³¹ bytes).

FILES /etc/auto_master Master map for automounter

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES automountd must be started with a process sensitivity label of ADMIN_LOW and a clearance of ADMIN_HIGH. It must have the PAF_TRUSTED_PATH process attribute, and must inherit the following privileges: file_dac_execute, file_dac_read, file_dac_search, file_dac_write, file_mac_read, file_mac_search, file_mac_write, file_owner, file_upgrade_sl, net_mac_read, net_privaddr, net_upgrade_sl, proc_audit_tcb, proc_setsl, sys_mount, and sys_trans_label.

**Trusted Solaris 8
4/01 Reference
Manual** automount(1M)
attributes(5)

NAME	autopush – Configures lists of automatically pushed STREAMS modules
SYNOPSIS	<pre> autopush -f <i>filename</i> autopush -g -M <i>major</i> -m <i>minor</i> autopush -r -M <i>major</i> -m <i>minor</i> </pre>
DESCRIPTION	The autopush command configures the list of modules to be automatically pushed onto the stream when a device is opened. It can also be used to remove a previous setting or get information on a setting.
OPTIONS	<p>The following options are supported:</p> <p>-f <i>filename</i> Sets up the autopush configuration for each driver according to the information stored in <i>filename</i>. An autopush file consists of lines of four or more fields, separated by spaces as shown below:</p> <pre> <i>major minor last-minor module1 module2 ... modulen</i> </pre> <p>The first field is a string that specifies the <i>major</i> device name, as listed in the <code>/kernel/drv</code> directory. The next two fields are integers that specify the <i>minor</i> device number and <i>last-minor</i> device number. The fields following represent the names of modules. If <i>minor</i> is -1, then all minor devices of a major driver specified by <i>major</i> are configured, and the value for <i>last-minor</i> is ignored. If <i>last-minor</i> is 0, then only a single minor device is configured. To configure a range of minor devices for a particular major, <i>minor</i> must be less than <i>last-minor</i>.</p> <p>The remaining fields list the names of modules to be automatically pushed onto the stream when opened, along with the position of an optional anchor. The maximum number of modules that can be pushed is eight. The modules are pushed in the order they are specified. The optional special character sequence [anchor] indicates that a STREAMS anchor should be placed on the stream at the module previously specified in the list; it is an error to specify more than one anchor or to have an anchor first in the list. The <code>sys_devices</code> privilege is required for this command to succeed.</p> <p>A nonzero exit status indicates that one or more of the lines in the specified file failed to complete successfully.</p> <p>-g Gets the current configuration setting of a particular <i>major</i> and <i>minor</i> device number specified with the -M and -m options respectively and displays the autopush modules associated with it. It will also return the starting minor device number if the request corresponds to a setting of a range (as described with the -f option).</p>

autopush(1M)

-M *major* Specifies the major device number.

-m *minor* Specifies the minor device number.

-r Removes the previous configuration setting of the particular *major* and *minor* device number specified with the **-M** and **-m** options respectively. If the values of *major* and *minor* correspond to a previously established setting of a range of minor devices, where *minor* matches the first minor device number in the range, the configuration would be removed for the entire range.

EXIT STATUS The following exit values are returned:

0 Successful completion.

non-zero An error occurred.

EXAMPLES **EXAMPLE 1** Using the autopush command

The following example gets the current configuration settings for the *major* and *minor* device numbers as indicated and displays the autopush modules associated with them for the character-special device `/dev/term/a`:

```
example# autopush -g -M 29 -m 0
Major      Minor      Lastminor    Modules
  29         0         1      ldterm ttcompat
```

SUMMARY OF TRUSTED SOLARIS CHANGES ATTRIBUTES The `sys_devices` privilege is required for this command to succeed.

`/etc/iu.ap` Autopush configuration file.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

`sad(7D)`
`bdconfig(1M)`, `ttymon(1M)`, `attributes(5)`, `streamio(7I)`, `ldterm(7M)`,
`ttcompat(7M)`

STREAMS Programming Guide

NAME	rpc.bootparamd, bootparamd – Boot parameter server				
SYNOPSIS	/usr/sbin/rpc.bootparamd [-d]				
DESCRIPTION	<p>rpc.bootparamd is a server process that provides information from a bootparams database to diskless clients at boot time. See bootparams(4).</p> <p>The source for the bootparams database is determined by the nsswitch.conf(4) file (on the machine running the rpc.bootparamd process).</p> <p>The rpc.bootparamd program can be invoked either by inetd(1M) or directly from the command line.</p>				
OPTIONS	-d Display debugging information.				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>rpc.bootparamd requires the trust path attribute with a UID of 0, and the sensitivity label ADMIN_LOW.</p> <p>/etc/bootparams Boot parameter database.</p> <p>/etc/nsswitch.conf Configuration file for the name-service switch.</p>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
Trusted Solaris 8 4/01 Reference Manual	<p>inetd(1M), nsswitch.conf(4)</p> <p>bootparams(4), attributes(5)</p>				
NOTES	<p>A diskless client requires service from at least one rpc.bootparamd process running on a server that is on the same IP subnetwork as the diskless client.</p> <p>Some routines that compare hostnames use case-sensitive string comparisons; some do not. If an incoming request fails, verify that the case of the hostname in the file to be parsed matches the case of the hostname called for, and attempt the request again.</p>				

bsmconv(1M)

NAME	bsmconv, bsmunconv – enable or disable the Basic Security Module (BSM)				
SYNOPSIS	<p>/etc/security/bsmconv [<i>rootdir...</i>]</p> <p>/etc/security/bsmunconv [<i>rootdir...</i>]</p>				
DESCRIPTION	<p>The bsmconv and bsmunconv scripts are used to enable or disable the BSM features, auditing and device protection.</p> <p>In the Trusted Solaris environment, the bsmconv and bsmunconv scripts are <i>not</i> used to enable or disable auditing, or to protect devices.</p>				
SUMMARY OF TRUSTED SOLARIS CHANGES	Enabling and disabling auditing in the Trusted Solaris environment does <i>not</i> use the bsmconv and bsmunconv scripts. These scripts do not exist in the Trusted Solaris environment. See <i>Trusted Solaris Audit Administration</i> for the procedure to disable and enable auditing. Devices are always protected in the Trusted Solaris environment.				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWcsr</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsr
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsr				
Trusted Solaris 8 4/01 Reference Manual	<p>allocate(1), auditd(1M), audit_startup(1M), audit.log(4), audit_control(4), device_allocate(4)</p> <p><i>Trusted Solaris Audit Administration</i></p>				
SunOS 5.8 Reference Manual	attributes(5)				

bsmunconv(1M)

NAME	bsmconv, bsmunconv – enable or disable the Basic Security Module (BSM)				
SYNOPSIS	<pre>/etc/security/bsmconv [rootdir...] /etc/security/bsmunconv [rootdir...]</pre>				
DESCRIPTION	<p>The bsmconv and bsmunconv scripts are used to enable or disable the BSM features, auditing and device protection.</p> <p>In the Trusted Solaris environment, the bsmconv and bsmunconv scripts are <i>not</i> used to enable or disable auditing, or to protect devices.</p>				
SUMMARY OF TRUSTED SOLARIS CHANGES	Enabling and disabling auditing in the Trusted Solaris environment does <i>not</i> use the bsmconv and bsmunconv scripts. These scripts do not exist in the Trusted Solaris environment. See <i>Trusted Solaris Audit Administration</i> for the procedure to disable and enable auditing. Devices are always protected in the Trusted Solaris environment.				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWcsr</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsr
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsr				
Trusted Solaris 8 4/01 Reference Manual	<p>allocate(1), auditd(1M), audit_startup(1M), audit.log(4), audit_control(4), device_allocate(4)</p> <p><i>Trusted Solaris Audit Administration</i></p>				
SunOS 5.8 Reference Manual	<p>attributes(5)</p>				

chk_encodings(1M)

NAME chk_encodings – Check the label encodings file syntax

SYNOPSIS `/usr/sbin/chk_encodings [-a] [-c maxclass] [pathname]`

DESCRIPTION `chk_encodings` checks the syntax of the label-encodings file specified by *pathname*; with the `-a` option, `chk_encodings` also prints a semantic analysis of the label-encodings file specified by *pathname*. If *pathname* is not specified, `chk_encodings` checks and analyzes `/etc/security/tsol/label_encodings`.

If label-encodings file analysis was requested, whatever analysis can be provided is written to the standard output file even if errors were found.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWtsu

OPTIONS

- `-a` Provide a semantic analysis of the label encodings file.
- `-c maxclass` Accept a maximum classification value of *maxclass* (default 255) in the label encodings file CLASSIFICATIONS section.

ERRORS When successful, `chk_encodings` returns an exit status of 0 (true) and writes to the standard output file a confirmation that no errors were found in *pathname*. Otherwise, `chk_encodings` returns an exit status of nonzero (false) and writes an error diagnostic to the standard output file.

FILES `/etc/security/tsol/label_encodings`
The label encodings file contains the classification names, words, constraints, and values for the defined labels of this system.

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

`label_encodings(4)`
`attributes(5)`

NAME	chroot – Change root directory for a command				
SYNOPSIS	/usr/sbin/chroot <i>newroot</i> <i>command</i>				
DESCRIPTION	<p>The chroot utility causes <i>command</i> to be executed relative to <i>newroot</i>. The meaning of any initial slashes (<i>/</i>) in the pathnames is changed to <i>newroot</i> for <i>command</i> and any of its child processes. Upon execution, the initial working directory is <i>newroot</i>.</p> <p>Notice that redirecting the output of <i>command</i> to a file,</p> <pre>example# chroot <i>newroot</i> <i>command</i> ><i>x</i></pre> <p>will create the file <i>x</i> relative to the original root of <i>command</i>, not the new one.</p> <p>The new root pathname is always relative to the current root. Even if a chroot is currently in effect, the <i>newroot</i> argument is relative to the current root of the running process.</p> <p>The <code>proc_chroot</code> privilege is required to run this command.</p>				
RETURN VALUES	The exit status of chroot is the return value of <i>command</i> .				
EXAMPLES	<p>EXAMPLE 1 Using the chroot utility.</p> <p>The chroot utility provides an easy way to extract tar files (see tar(1)) written with absolute filenames to a different location:</p> <pre>example# cp /usr/sbin/static/tar /tmp example# dd if=/dev/nrst0 chroot /tmp tar xvf -</pre> <p>Note that tar is statically linked, so it is not necessary to copy any shared libraries to the <i>newroot</i> filesystem.</p>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
SUMMARY OF TRUSTED SQUARES	To succeed, this command needs the <code>proc_chroot</code> privilege.				
References	<p>chroot(2)</p> <p>cd(1), ttyname(3C), attributes(5)</p>				
NOTES	<p>Exercise extreme caution when referencing device files in the new root file system.</p> <p>References by routines such as ttyname(3C) to <code>stdin</code>, <code>stdout</code>, and <code>stderr</code> will find that the device associated with the file descriptor is unknown after chroot is run.</p>				

SUMMARY OF
TRUSTED
SQUARES
Trusted Squares
4/01 Reference
Manual
SUNWcsu
Reference Manual
NOTES

clist(1M)

NAME	clist – Display a list of commands in a profile shell
SYNOPSIS	clist [-hpniu] (obsolete)
DESCRIPTION	The profile shell clist command is replaced in Trusted Solaris 8 and later releases with the profiles(1) and smprofile(1M) commands.

NAME	coreadm – core file administration																
SYNOPSIS	<p>coreadm [-g <i>pattern</i>] [-i <i>pattern</i>] [-d <i>option...</i>] [-e <i>option...</i>]</p> <p>coreadm [-p <i>pattern</i>] [<i>pid...</i>]</p> <p>coreadm -u</p>																
DESCRIPTION	<p>The coreadm command is used to specify the name and location of core files produced by abnormally-terminating processes. See core(4).</p> <p>The first form shown in the synopsis can be executed only by an administrator and is used to configure system-wide core file options, including a global core file name pattern and a per-process core file name pattern for the init(1M) process. All such settings are saved in coreadm's configuration file <code>/etc/coreadm.conf</code> for setting on reboot. See init(1M).</p> <p>The second form can be executed by non-privileged users and is used to specify the file name pattern to be used by the operating system when generating a per-process core file.</p> <p>The third form can be executed only by an administrator and is used to update all system-wide core file options based on the contents of <code>/etc/coreadm.conf</code>. Normally this option is used only on reboot by the startup script <code>/etc/init.d/coreadm</code>.</p> <p>A core file name pattern is a normal file system path name with embedded variables, specified with a leading % character, that are expanded from values in effect when a core file is generated by the operating system. The possible variables are:</p> <table> <tr><td>%p</td><td>process-ID</td></tr> <tr><td>%u</td><td>effective user-ID</td></tr> <tr><td>%g</td><td>effective group-ID</td></tr> <tr><td>%f</td><td>executable file name</td></tr> <tr><td>%n</td><td>system node name (<code>uname -n</code>)</td></tr> <tr><td>%m</td><td>machine name (<code>uname -m</code>)</td></tr> <tr><td>%t</td><td>decimal value of <code>time(2)</code></td></tr> <tr><td>%%</td><td>literal %</td></tr> </table> <p>For example, the core file name pattern:</p> <pre>/var/core/core.%f.%p</pre> <p>would result, for command <code>foo</code> with process-ID 1234, in the core file name:</p> <pre>/var/core/core.foo.1234</pre> <p>The coreadm command with no arguments reports the current system configuration, for example:</p>	%p	process-ID	%u	effective user-ID	%g	effective group-ID	%f	executable file name	%n	system node name (<code>uname -n</code>)	%m	machine name (<code>uname -m</code>)	%t	decimal value of <code>time(2)</code>	%%	literal %
%p	process-ID																
%u	effective user-ID																
%g	effective group-ID																
%f	executable file name																
%n	system node name (<code>uname -n</code>)																
%m	machine name (<code>uname -m</code>)																
%t	decimal value of <code>time(2)</code>																
%%	literal %																

coreadm(1M)

```
$ coreadm
  global core file pattern: /var/core/core.%f.%p
  init core file pattern: core
    global core dumps: enabled
    per-process core dumps: enabled
  global setid core dumps: enabled
per-process setid core dumps: disabled
  global core dump logging: disabled
```

The `coreadm` command with only a list of process-IDs reports each process's per-process core file name pattern, for example:

```
$ coreadm 278 5678
278:   core.%f.%p
5678:  /home/george/cores/%f.%p.%t
```

Only the owner of a process can interrogate a process in this manner, unless the command is run with the `proc_owner` privilege. To succeed, `coreadm` must be invoked at a label that dominates the process label, or with the `proc_mac_read` privilege.

When a process is dumping core, the operating system will generate two possible core files, the global core file and the per-process core file. Both files, one or the other, or no file will be generated, based on the system options in effect at the time.

When generated, a global core file will be created mode 600 and will be owned by root. Non-privileged users cannot examine such files.

Ordinary per-process core files are created mode 600 under the credentials of the process. The owner of the process can examine such files.

A process that is or ever has been `setuid`, `setgid`, or privileged since its last `exec(2)`, including a process that began life with privilege and gave up that privilege by way of `setuid(2)` or `setppriv(2)`, presents security issues with respect to dumping core, as it may contain sensitive information in its address space to which the current non-privileged owner of the process should not have access. If `setid` core files are enabled, they will be created mode 600 and will be owned by root.

OPTIONS

The following options are supported:

- | | |
|--------------------------------|--|
| <code>-g <i>pattern</i></code> | Set the global core file name pattern to <i>pattern</i> . The pattern must start with a / and can contain any of the special % variables described in the DESCRIPTION. |
| | To succeed with this option, the command requires the <code>sys_config</code> privilege, and must be invoked with an effective UID of 0 at the label <code>admin_low</code> . |
| <code>-i <i>pattern</i></code> | Set the per-process core file name <i>pattern</i> for <code>init(1M)</code> to <i>pattern</i> . This is the same as <code>coreadm -p <i>pattern</i> 1</code> except that the setting will be persistent across reboot. |

	To succeed with this option, the command requires the <code>sys_config</code> privilege, and must be invoked with an effective UID of 0 at the label <code>admin_low</code> .										
<code>-e option...</code>	<p>Enable the specified core file option. Specify <i>option</i> as one of the following:</p> <table> <tr> <td><code>global</code></td><td>Allow core dumps using global core pattern</td></tr> <tr> <td><code>process</code></td><td>Allow core dumps using per-process core pattern</td></tr> <tr> <td><code>global-setid</code></td><td>Allow set-id or privileged core dumps using global core pattern</td></tr> <tr> <td><code>proc-setid</code></td><td>Allow set-id or privileged core dumps using per-process core pattern</td></tr> <tr> <td><code>log</code></td><td>Generate a <code>syslog(3C)</code> message when generation of a global core file is attempted.</td></tr> </table> <p>Multiple <code>-e</code> and <code>-d</code> options can be specified on the command line. To succeed with these options, the command requires the <code>sys_config</code> privilege, and must be invoked with an effective UID of 0 at the label <code>admin_low</code>.</p>	<code>global</code>	Allow core dumps using global core pattern	<code>process</code>	Allow core dumps using per-process core pattern	<code>global-setid</code>	Allow set-id or privileged core dumps using global core pattern	<code>proc-setid</code>	Allow set-id or privileged core dumps using per-process core pattern	<code>log</code>	Generate a <code>syslog(3C)</code> message when generation of a global core file is attempted.
<code>global</code>	Allow core dumps using global core pattern										
<code>process</code>	Allow core dumps using per-process core pattern										
<code>global-setid</code>	Allow set-id or privileged core dumps using global core pattern										
<code>proc-setid</code>	Allow set-id or privileged core dumps using per-process core pattern										
<code>log</code>	Generate a <code>syslog(3C)</code> message when generation of a global core file is attempted.										
<code>-d option...</code>	<p>Disable the specified core file option. See the <code>-e option</code> for descriptions of possible options.</p> <p>Multiple <code>-e</code> and <code>-d</code> options can be specified on the command line. To succeed with these options, the command requires the <code>sys_config</code> privilege, and must be invoked with an effective UID of 0 at the label <code>admin_low</code>.</p>										
<code>-p pattern</code>	<p>Set the per-process core file name pattern to <i>pattern</i> for each of the specified process-IDs. The pattern can contain any of the special % variables described in the <code>DESCRIPTION</code> and need not begin with /. If it does not begin with /, it will be evaluated relative to the current directory in effect when the process generates a core file.</p> <p>A user can apply the <code>-p</code> option only to processes owned by that user, unless the command is run with <code>proc_owner</code> privilege. To succeed, <code>coreadm -p</code> must be invoked at a label that is dominated by the process label, or with the <code>proc_mac_write</code> privilege. The per-process core file name pattern will be inherited by future child processes of the affected processes. See <code>fork(2)</code>.</p>										
<code>-u</code>	<p>Update system-wide core file options from the contents of the configuration file <code>/etc/coreadm.conf</code>. If the configuration file is missing or contains invalid values, default values are substituted. Following the update, the configuration file is resynchronized with the system core file configuration. To succeed with this option, the command requires the <code>sys_config</code> privilege, and must be</p>										

coreadm(1M)

invoked with an effective UID of 0 at the label `admin_low`.

OPERANDS The following operands are supported:

pid process-ID

EXIT STATUS The following exit values are returned:

0 Successful completion.

1 A fatal error occurred while either obtaining or modifying the system core file configuration.

2 Invalid command line options were specified.

EXAMPLES **EXAMPLE 1** Setting the core file name pattern

When executed from a user's `$HOME/.profile` or `$HOME/.login`, the following command sets the core file name pattern for all processes run during the login session:

```
example$ coreadm -p core.%f.%p $$
```

`$$` is the process-id of the currently running shell. The per-process core file name pattern is inherited by all child processes.

EXAMPLE 2 Dumping user's files into a subdirectory

The following command dumps all of the user's core dumps into the `corefiles` subdirectory of the home directory, discriminated by the system node name. This is useful for users who use many different machines but have a shared home directory.

```
example$ coreadm -p $HOME/corefiles/%n.%f.%p $$
```

FILES `/etc/init.d/coreadm`

`/etc/coreadm.conf`

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES Setting or examining the core file pattern for a process requires MAC access and process ownership. Setting system-wide options requires an effective UID of 0, a label of `admin_low`, and the `sys_config` privilege.

Trusted Solaris 8 4/01 Reference Manual `init(1M)`, `exec(2)`, `fork(2)`, `setppriv(2)`, `setuid(2)`
`gcore(1)`, `time(2)`, `syslog(3C)`, `core(4)`, `attributes(5)`

NAME	cron – Clock daemon
SYNOPSIS	<code>/usr/sbin/cron</code>
DESCRIPTION	<p>The <code>cron</code> command starts a process that executes commands at specified dates and times. Regularly scheduled commands can be specified according to instructions found in <code>crontab</code> files in the directory <code>/var/spool/cron/crontabs</code>. Users can submit their own <code>crontab</code> file using the <code>crontab(1)</code> command. Commands which are to be executed only once may be submitted using the <code>at(1)</code> command.</p> <p><code>cron</code> only examines <code>crontab</code> or <code>at</code> command files during its own process initialization phase and when the <code>crontab</code> or <code>at</code> command is run. This reduces the overhead of checking for new or changed files at regularly scheduled intervals.</p> <p>Since <code>cron</code> never exits, it should be executed only once. This is done routinely through <code>/etc/rc2.d/S75cron</code> at system boot time. The file <code>/etc/cron.d/FIFO</code> is used (among other things) as a lock file to prevent the execution of more than one instance of <code>cron</code>.</p> <p><code>cron</code> captures the output of the job's <code>stdout</code> and <code>stderr</code> streams, and, if it is non-empty, mails the output to the user. If the job does not produce output, no mail is sent to the user (unless the job is an <code>at(1)</code> job and the <code>-m</code> option was specified when the job was submitted).</p>
Setting cron Defaults	<p>To keep a log of all actions taken by <code>cron</code>, <code>CRONLOG=YES</code> (by default) must be specified in the <code>/etc/default/cron</code> file. If <code>CRONLOG=NO</code> is specified, no logging is done. Keeping the log is a user configurable option since <code>cron</code> usually creates huge log files.</p> <p>The <code>PATH</code> for user <code>cron</code> jobs can be set using <code>PATH=</code> in <code>/etc/default/cron</code>. The <code>PATH</code> for root <code>cron</code> jobs can be set using <code>SUPATH=</code> in <code>/etc/default/cron</code>. The security implications of setting <code>PATH</code> and <code>SUPATH</code> should be carefully considered.</p> <p>Example <code>/etc/default/cron</code> file:</p> <pre>CRONLOG=YES PATH=/usr/bin:/usr/ucb:</pre> <p>This example enables logging and sets the default <code>PATH</code> used by non-root jobs to <code>/usr/bin:/usr/ucb:</code>. Root jobs will continue to use <code>/usr/sbin:/usr/bin</code>.</p> <p><code>/etc/cron.d/logchecker</code> is a script that checks to see if the log file has exceeded the system <code>ulimit</code>. If so, the log file is moved to <code>/var/cron/olog</code>.</p>
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>The job directories <code>/var/spool/cron/crontabs</code> and <code>/var/spool/cron/atjobs</code> are multilevel directories (MLDs). The MLD job directory provides for the separation of job files at different sensitivity labels. Hence, there can be multiple <code>crontab</code> files for a single user within the <code>crontabs</code> directory, but each <code>crontab</code> file is at a different sensitivity label. In addition, a user can have multiple <code>atjob</code> files at different sensitivity labels.</p>

cron(1M)

Each crontab file in the crontabs MLD and each atjob file in the atjobs MLD has an ancillary file containing information used by cron to set up a job. The crontab ancillary files are named `username.ad`, and the atjobs ancillary files are named `jobname.ad`.

The clock daemon must be started with the root userid, must have the `PAF_TRUSTED_PATH` process attribute, and it must inherit the following privileges: `file_mac_write`, `net_mac_read`, `proc_setid`, `proc_setsl`, `proc_setclr`, `sys_audit`, `proc_audit_tcb`, `file_dac_read`, and `file_owner`.

If the clock daemon has the `PAF_PRIV_DEBUG` process attribute, it passes the attribute on to the job to be executed. Because the daemon never has the `PAF_TOKMAPPER`, `PAF_DISKLESS_BOOT`, and `PAF_SELAGNT` process attributes, these attributes will not be passed on to the job to be executed.

The clock daemon creates the `/var/cron/log` file at the `ADMIN_HIGH` sensitivity label.

In the default Trusted Solaris environment, there are two pairs of crontab and its ancillary file for the root userid: one pair at the `ADMIN_HIGH` sensitivity label, and the other pair at the `ADMIN_LOW` sensitivity label.

FILES	<code>/etc/cron.d</code>	main cron directory
	<code>/etc/cron.d/FIFO</code>	used as a lock file
	<code>/etc/default/cron</code>	contains cron default settings
	<code>/var/cron/log</code>	cron history information
	<code>/var/spool/cron</code>	spool area
	<code>/etc/cron.d/logchecker</code>	moves log file to <code>/var/cron/olog</code> if log file exceeds system ulimit.
	<code>/etc/cron.d/queuedefs</code>	queue description file for at, batch, and cron.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8
4/01 Reference
Manual
Reference Manual
DIAGNOSTICS

at(1), crontab(1)

sh(1), queuedefs(4), attributes(5)

A history of all actions taken by cron is stored in `/var/cron/log` and (possibly) `/var/cron/olog`.

NAME	devfsadm, devfsadmd – administration command for /dev and /devices												
SYNOPSIS	<pre> /usr/sbin/devfsadm [-C] [-c <i>device_class</i>] [-i <i>driver_name</i>] [-n] [-r <i>root_dir</i>] [-s] [-t <i>table_file</i>] [-v] /usr/lib/devfsadm/devfsadmd </pre>												
DESCRIPTION	<p>The devfsadm command maintains the /dev and /devices namespaces. It replaces the previous suite of devfs administration tools including drvconfig(1M), disks(1M), tapes(1M), ports(1M), audlinks(1M), and devlinks(1M).</p> <p>The default operation is to attempt to load every driver in the system and attach to all possible device instances. devfsadm then creates device special files in /devices and logical links in /dev.</p> <p>devfsadmd(1M) is the daemon version of devfsadm(1M). The daemon is started by the /etc/rc* scripts during system startup and is responsible for handling both reconfiguration boot processing and updating /dev and /devices in response to dynamic reconfiguration event notifications from the kernel.</p> <p>For compatibility purposes, drvconfig(1M), disks(1M), tapes(1M), ports(1M), audlinks(1M), and devlinks(1M) are implemented as links to devfsadm.</p> <p>In addition to managing /dev and /devices, devfsadm also maintains the path_to_inst(4) database.</p>												
OPTIONS	<p>The following options are supported:</p> <table> <tr> <td>-C</td><td>Cleanup mode. Prompts devfsadm to invoke cleanup routines that are not normally invoked to remove dangling logical links. If -c is also used, devfsadm only cleans up for the listed devices' classes.</td></tr> <tr> <td>-c <i>device_class</i></td><td>Restrict operations to devices of class <i>device_class</i>. The following values are defined for <i>device_class</i>: disk, tape, port, audio, and pseudo. This option may be specified more than once to specify multiple device classes.</td></tr> <tr> <td>-i <i>driver_name</i></td><td>Configure only the devices for the named driver, <i>driver_name</i>.</td></tr> <tr> <td>-n</td><td>Do not attempt to load drivers or add new nodes to the kernel device tree.</td></tr> <tr> <td>-s</td><td>Suppress any changes to /dev or /devices. This is useful with the -v option for debugging.</td></tr> <tr> <td>-t <i>table_file</i></td><td>Read an alternate devlink.tab file. devfsadm normally reads /etc/devlink.tab.</td></tr> </table>	-C	Cleanup mode. Prompts devfsadm to invoke cleanup routines that are not normally invoked to remove dangling logical links. If -c is also used, devfsadm only cleans up for the listed devices' classes.	-c <i>device_class</i>	Restrict operations to devices of class <i>device_class</i> . The following values are defined for <i>device_class</i> : disk, tape, port, audio, and pseudo. This option may be specified more than once to specify multiple device classes.	-i <i>driver_name</i>	Configure only the devices for the named driver, <i>driver_name</i> .	-n	Do not attempt to load drivers or add new nodes to the kernel device tree.	-s	Suppress any changes to /dev or /devices. This is useful with the -v option for debugging.	-t <i>table_file</i>	Read an alternate devlink.tab file. devfsadm normally reads /etc/devlink.tab.
-C	Cleanup mode. Prompts devfsadm to invoke cleanup routines that are not normally invoked to remove dangling logical links. If -c is also used, devfsadm only cleans up for the listed devices' classes.												
-c <i>device_class</i>	Restrict operations to devices of class <i>device_class</i> . The following values are defined for <i>device_class</i> : disk, tape, port, audio, and pseudo. This option may be specified more than once to specify multiple device classes.												
-i <i>driver_name</i>	Configure only the devices for the named driver, <i>driver_name</i> .												
-n	Do not attempt to load drivers or add new nodes to the kernel device tree.												
-s	Suppress any changes to /dev or /devices. This is useful with the -v option for debugging.												
-t <i>table_file</i>	Read an alternate devlink.tab file. devfsadm normally reads /etc/devlink.tab.												

devfsadm(1M)

-r *root_dir* Presume that the /dev and /devices directory trees are found under *root_dir*, not directly under root (/). No other use or assumptions are made about *root_dir*.

-v Print changes to /dev and /devices in verbose mode.

EXIT STATUS The following exit values are returned:

0 Successful completion.

1 An error occurred.

FILES

/devices	device nodes directory
/dev	logical symbolic links to /devices
/usr/lib/devfsadm/devfsadmd	devfsadm daemon
/etc/init.d/devfsadm	daemon start/stop script
/etc/rcS.d/S50devfsadm	link to init.d script
/etc/rc0.d/K83devfsadm	link to init.d script
/dev/.devfsadm_dev.lock	update lock file
/dev/.devfsadm_daemon.lock	daemon lock file

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES **Trusted Solaris 8 4/01 Reference Manual**

/etc/security/tsol/minor_perm.adjunct, may not exist or may have different contents or interpretations in a future release.

add_drv(1M), drvconfig(1M), modload(1M), modunload(1M), rem_drv(1M)

devlinks(1M), disks(1M), modinfo(1M), ports(1M), tapes(1M), path_to_inst(4), attributes(5)

NOTES This document does not constitute an API. /etc/minor_perm, /etc/security/tsol/minor_perm.adjunct, /etc/name_to_major, /etc/driver_classes, and /devices may not exist or may have different contents or interpretations in a future release. The existence of this notice does not imply that any other documentation that lacks this notice constitutes an API.

NAME	devfsadm, devfsadmd – administration command for /dev and /devices												
SYNOPSIS	<pre>/usr/sbin/devfsadm [-C] [-c <i>device_class</i>] [-i <i>driver_name</i>] [-n] [-r <i>root_dir</i>] [-s] [-t <i>table_file</i>] [-v]</pre> <pre>/usr/lib/devfsadm/devfsadmd</pre>												
DESCRIPTION	<p>The devfsadm command maintains the /dev and /devices namespaces. It replaces the previous suite of devfs administration tools including drvconfig(1M), disks(1M), tapes(1M), ports(1M), audlinks(1M), and devlinks(1M).</p> <p>The default operation is to attempt to load every driver in the system and attach to all possible device instances. devfsadm then creates device special files in /devices and logical links in /dev.</p> <p>devfsadmd(1M) is the daemon version of devfsadm(1M). The daemon is started by the /etc/rc* scripts during system startup and is responsible for handling both reconfiguration boot processing and updating /dev and /devices in response to dynamic reconfiguration event notifications from the kernel.</p> <p>For compatibility purposes, drvconfig(1M), disks(1M), tapes(1M), ports(1M), audlinks(1M), and devlinks(1M) are implemented as links to devfsadm.</p> <p>In addition to managing /dev and /devices, devfsadm also maintains the path_to_inst(4) database.</p>												
OPTIONS	<p>The following options are supported:</p> <table> <tr> <td>-C</td><td>Cleanup mode. Prompts devfsadm to invoke cleanup routines that are not normally invoked to remove dangling logical links. If -c is also used, devfsadm only cleans up for the listed devices' classes.</td></tr> <tr> <td>-c <i>device_class</i></td><td>Restrict operations to devices of class <i>device_class</i>. The following values are defined for <i>device_class</i>: disk, tape, port, audio, and pseudo. This option may be specified more than once to specify multiple device classes.</td></tr> <tr> <td>-i <i>driver_name</i></td><td>Configure only the devices for the named driver, <i>driver_name</i>.</td></tr> <tr> <td>-n</td><td>Do not attempt to load drivers or add new nodes to the kernel device tree.</td></tr> <tr> <td>-s</td><td>Suppress any changes to /dev or /devices. This is useful with the -v option for debugging.</td></tr> <tr> <td>-t <i>table_file</i></td><td>Read an alternate devlink.tab file. devfsadm normally reads /etc/devlink.tab.</td></tr> </table>	-C	Cleanup mode. Prompts devfsadm to invoke cleanup routines that are not normally invoked to remove dangling logical links. If -c is also used, devfsadm only cleans up for the listed devices' classes.	-c <i>device_class</i>	Restrict operations to devices of class <i>device_class</i> . The following values are defined for <i>device_class</i> : disk, tape, port, audio, and pseudo. This option may be specified more than once to specify multiple device classes.	-i <i>driver_name</i>	Configure only the devices for the named driver, <i>driver_name</i> .	-n	Do not attempt to load drivers or add new nodes to the kernel device tree.	-s	Suppress any changes to /dev or /devices. This is useful with the -v option for debugging.	-t <i>table_file</i>	Read an alternate devlink.tab file. devfsadm normally reads /etc/devlink.tab.
-C	Cleanup mode. Prompts devfsadm to invoke cleanup routines that are not normally invoked to remove dangling logical links. If -c is also used, devfsadm only cleans up for the listed devices' classes.												
-c <i>device_class</i>	Restrict operations to devices of class <i>device_class</i> . The following values are defined for <i>device_class</i> : disk, tape, port, audio, and pseudo. This option may be specified more than once to specify multiple device classes.												
-i <i>driver_name</i>	Configure only the devices for the named driver, <i>driver_name</i> .												
-n	Do not attempt to load drivers or add new nodes to the kernel device tree.												
-s	Suppress any changes to /dev or /devices. This is useful with the -v option for debugging.												
-t <i>table_file</i>	Read an alternate devlink.tab file. devfsadm normally reads /etc/devlink.tab.												

devfsadmd(1M)

-r *root_dir* Presume that the /dev and /devices directory trees are found under *root_dir*, not directly under root (/). No other use or assumptions are made about *root_dir*.

-v Print changes to /dev and /devices in verbose mode.

EXIT STATUS The following exit values are returned:

0 Successful completion.

1 An error occurred.

FILES

/devices	device nodes directory
/dev	logical symbolic links to /devices
/usr/lib/devfsadm/devfsadmd	devfsadm daemon
/etc/init.d/devfsadm	daemon start/stop script
/etc/rcS.d/S50devfsadm	link to init.d script
/etc/rc0.d/K83devfsadm	link to init.d script
/dev/.devfsadm_dev.lock	update lock file
/dev/.devfsadm_daemon.lock	daemon lock file

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES **Trusted Solaris 8 4/01 Reference Manual**

/etc/security/tsol/minor_perm.adjunct, may not exist or may have different contents or interpretations in a future release.

add_drv(1M), drvconfig(1M), modload(1M), modunload(1M), rem_drv(1M)

devlinks(1M), disks(1M), modinfo(1M), ports(1M), tapes(1M), path_to_inst(4), attributes(5)

NOTES This document does not constitute an API. /etc/minor_perm, /etc/security/tsol/minor_perm.adjunct, /etc/name_to_major, /etc/driver_classes, and /devices may not exist or may have different contents or interpretations in a future release. The existence of this notice does not imply that any other documentation that lacks this notice constitutes an API.

NAME	device_clean – Device clean programs				
SYNOPSIS	<code>/etc/security/lib/device-clean-program character-media-label-string [-A -D]</code>				
DESCRIPTION	<p>An allocatable device may optionally have a device clean program. Device clean programs are specified in the <i>device-clean</i> field in the <code>device_allocate(4)</code> file. Device clean programs are invoked by <code>allocate(1)</code> and <code>deallocate(1)</code> to clean device states, registers, and any residual information in the device before it is allocated to a user as required by the <i>object reuse</i> policy, and also to ensure proper media labeling by asking the user to confirm the correct labeled media is inserted in the device on allocation and by asking the user to confirm removal of the media and affix correct label on the media.</p>				
ATTRIBUTES	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWcsu</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
OPTIONS	<p><i>character-media-label-string</i> Provide CMW Label of the device. This information is used by most device clean programs in a prompt to remind the user to affix a correct label to the removable media.</p> <p>-A The device clean program is invoked from <code>allocate(1)</code> command before the device is allocated to a user.</p> <p>-D The device clean program is invoked from <code>deallocate(1)</code> command after the device is deallocated from a user.</p>				
FILES	<code>/etc/security/device_allocate</code> Mandatory access control file for devices				
Trusted Solaris 8 4/01 Reference Manual	<code>allocate(1)</code> , <code>deallocate(1)</code> , <code>device_allocate(4)</code> <code>attributes(5)</code>				

devpolicy(1M)

NAME	devpolicy – Configure device policy					
SYNOPSIS	devpolicy [-s -v] [-f <i>policyfile</i>] [-r <i>rootdir</i>]					
DESCRIPTION	<p>devpolicy reads the /etc/security/tsol/device_policy file and, for each device node in the /devices tree, constructs device policy information and downloads the information to the kernel.</p> <p>To be successful, devpolicy requires the trusted path attribute and the sys_devices privilege. If device policy has been downloaded by an earlier invocation of the command, devpolicy will fail. If a device has two or more device nodes that are assigned different policies in the device_policy file, devpolicy displays a warning.</p>					
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:					
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWtsr</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsr
ATTRIBUTE TYPE	ATTRIBUTE VALUE					
Availability	SUNWtsr					
OPTIONS	<p>-s Silent mode; suppresses non-critical warning messages.</p> <p>-v Verbose mode; displays all warning messages, including messages for unknown devices.</p> <p>-f <i>policyfile</i> Read <i>policyfile</i> instead of /etc/security/tsol/device_policy.</p> <p>-r <i>rootdir</i> Find devices under <i>rootdir</i> instead of /devices.</p>					
EXIT STATUS	<p>0 Successful.</p> <p>>0 An error occurred.</p>					
FILES	<p>/etc/security/tsol/device_policy Security policy configuration file for devices</p>					
Trusted Solaris 8 4/01 Reference Manual	<p>drvconfig(1M), device_policy(4) attributes(5)</p>					

NAME	dfmounts – Display mounted resource information									
SYNOPSIS	dfmounts [-F <i>FSType</i>] [-h] [-o <i>specific_options</i>] [<i>restriction...</i>]									
DESCRIPTION	<p>dfmounts shows the local resources shared through a distributed file system <i>FSType</i> along with a list of clients that have the resource mounted. If <i>restriction</i> is not specified, dfmounts shows file systems that are currently shared on any NFS server. <i>specific_options</i> as well as the availability and semantics of <i>restriction</i> are specific to particular distributed file system types.</p> <p>If dfmounts is entered without arguments, all remote resources currently mounted on the local system are displayed, regardless of file system type.</p>									
dfmounts Output	<p>The output of dfmounts consists of an optional header line (suppressed with the -h flag) followed by a list of lines containing whitespace-separated fields. For each resource, the fields are:</p> <p><i>resource server pathname clients ...where:</i></p> <table><tr><td><i>resource</i></td><td>Specifies the resource name that must be given to the mount(1M) command.</td></tr><tr><td><i>server</i></td><td>Specifies the system from which the resource was mounted.</td></tr><tr><td><i>pathname</i></td><td>Specifies the pathname that must be given to the share(1M) command.</td></tr><tr><td><i>clients</i></td><td>Is a comma-separated list of systems that have mounted the resource. Clients are listed in the form <i>domain.</i>, <i>domain.system</i>, or <i>system</i>, depending on the file system type.</td></tr></table> <p>A field may be null. Each null field is indicated by a hyphen (–) unless the remainder of the fields on the line are also null; in which case, the hyphen may be omitted.</p> <p>Fields with whitespace are enclosed in quotation marks (" ").</p>		<i>resource</i>	Specifies the resource name that must be given to the mount(1M) command.	<i>server</i>	Specifies the system from which the resource was mounted.	<i>pathname</i>	Specifies the pathname that must be given to the share(1M) command.	<i>clients</i>	Is a comma-separated list of systems that have mounted the resource. Clients are listed in the form <i>domain.</i> , <i>domain.system</i> , or <i>system</i> , depending on the file system type.
<i>resource</i>	Specifies the resource name that must be given to the mount(1M) command.									
<i>server</i>	Specifies the system from which the resource was mounted.									
<i>pathname</i>	Specifies the pathname that must be given to the share(1M) command.									
<i>clients</i>	Is a comma-separated list of systems that have mounted the resource. Clients are listed in the form <i>domain.</i> , <i>domain.system</i> , or <i>system</i> , depending on the file system type.									
OPTIONS	<table><tr><td>-F <i>FSType</i></td><td>Specify filesystem type. Defaults to the first entry in /etc/dfs/fstypes. <i>Note:</i> currently the only valid <i>FSType</i> is nfs.</td></tr><tr><td>-h</td><td>Suppress header line in output.</td></tr><tr><td>-o <i>specific_options</i></td><td>Specify options specific to the filesystem provided by the -F option. <i>Note:</i> currently no options are supported.</td></tr></table>		-F <i>FSType</i>	Specify filesystem type. Defaults to the first entry in /etc/dfs/fstypes. <i>Note:</i> currently the only valid <i>FSType</i> is nfs.	-h	Suppress header line in output.	-o <i>specific_options</i>	Specify options specific to the filesystem provided by the -F option. <i>Note:</i> currently no options are supported.		
-F <i>FSType</i>	Specify filesystem type. Defaults to the first entry in /etc/dfs/fstypes. <i>Note:</i> currently the only valid <i>FSType</i> is nfs.									
-h	Suppress header line in output.									
-o <i>specific_options</i>	Specify options specific to the filesystem provided by the -F option. <i>Note:</i> currently no options are supported.									
FILES	/etc/dfs/fstypes	file system types								
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:									

dfmounts(1M)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**
Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.6
Reference Manual

The output fields show the resource and pathname that must be given to the Trusted Solaris versions of the mount and share commands.

dfshares(1M), mount(1M), share(1M), unshare(1M)

attributes(5)

NAME	dfshares – List available resources from remote or local systems								
SYNOPSIS	dfshares [-F <i>FSType</i>] [-h] [-o <i>specific_options</i>] [<i>server...</i>]								
DESCRIPTION	<p>dfshares provides information about resources available to the host through a distributed file system of type <i>FSType</i>. <i>specific_options</i> as well as the semantics of <i>server</i> are specific to particular distributed file systems.</p> <p>If dfshares is entered without arguments, all resources currently shared on the local system are displayed, regardless of file system type.</p> <p>The output of dfshares consists of an optional header line (suppressed with the -h flag) followed by a list of lines containing whitespace-separated fields. For each resource, the fields are:</p> <pre>resource server access transport</pre> <p>where</p> <table> <tr> <td><i>resource</i></td><td>Specifies the resource name that must be given to the mount(1M) command.</td></tr> <tr> <td><i>server</i></td><td>Specifies the name of the system that is making the resource available.</td></tr> <tr> <td><i>access</i></td><td>Specifies the access permissions granted to the client systems, either ro (for read-only) or rw (for read/write). If dfshares cannot determine access permissions, a hyphen (-) is displayed.</td></tr> <tr> <td><i>transport</i></td><td>Specifies the transport provider over which the resource is shared.</td></tr> </table> <p>A field may be null. Each null field is indicated by a hyphen (-) unless the remainder of the fields on the line are also null; in which case, the hyphen may be omitted.</p>	<i>resource</i>	Specifies the resource name that must be given to the mount(1M) command.	<i>server</i>	Specifies the name of the system that is making the resource available.	<i>access</i>	Specifies the access permissions granted to the client systems, either ro (for read-only) or rw (for read/write). If dfshares cannot determine access permissions, a hyphen (-) is displayed.	<i>transport</i>	Specifies the transport provider over which the resource is shared.
<i>resource</i>	Specifies the resource name that must be given to the mount(1M) command.								
<i>server</i>	Specifies the name of the system that is making the resource available.								
<i>access</i>	Specifies the access permissions granted to the client systems, either ro (for read-only) or rw (for read/write). If dfshares cannot determine access permissions, a hyphen (-) is displayed.								
<i>transport</i>	Specifies the transport provider over which the resource is shared.								
OPTIONS	<p>-F <i>FSType</i> Specify filesystem type. Defaults to the first entry in /etc/dfs/fstypes.</p> <p>-h Suppress header line in output.</p> <p>-o <i>specific_options</i> Specify options specific to the filesystem provided by the -F option.</p>								
FILES	/etc/dfs/fstypes								
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:								

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

dfshares(1M)

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**
Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

The output list shows the resource name that must be given to the Trusted Solaris version of the mount command.

dfmounts(1M), mount(1M), share(1M), unshare(1M)

attributes(5)

NAME	dispadmin – Process scheduler administration
SYNOPSIS	<p>dispadmin -l</p> <p>dispadmin -c <i>class</i> -g [-r <i>res</i>]</p> <p>dispadmin -c <i>class</i> -s <i>file</i></p>
DESCRIPTION	<p>The dispadmin command displays or changes process scheduler parameters while the system is running.</p> <p>dispadmin does limited checking on the values supplied in <i>file</i> to verify that they are within their required bounds. The checking, however, does not attempt to analyze the effect that the new values have on the performance of the system. Inappropriate values can have a negative effect on system performance. (See <i>System Administration Guide, Volume 1</i>.)</p>
OPTIONS	<p>-l Lists the scheduler classes currently configured in the system.</p> <p>-c <i>class</i> Specifies the class whose parameters are to be displayed or changed. Valid <i>class</i> values are: RT for the real-time class, TS for the time-sharing class, and IA for the inter-active class. The time-sharing and inter-active classes share the same scheduler, so changes to the scheduling parameters of one will change those of the other.</p> <p>-g Gets the parameters for the specified class and writes them to the standard output. Parameters for the real-time class are described in <code>rt_dptbl(4)</code>. Parameters for the time-sharing and inter-active classes are described in <code>ts_dptbl(4)</code>.</p> <p>-r <i>res</i> When using the -g option you may also use the -r option to specify a resolution to be used for outputting the time quantum values. If no resolution is specified, time quantum values are in milliseconds. If <i>res</i> is specified it must be a positive integer between 1 and 100000000 inclusive, and the resolution used is the reciprocal of <i>res</i> in seconds. For example, a <i>res</i> value of 10 yields time quantum values expressed in tenths of a second; a <i>res</i> value of 1000000 yields time quantum values expressed in microseconds. If the time quantum cannot be expressed as an integer in the specified resolution, it is rounded up to the next integral multiple of the specified resolution.</p> <p>-s <i>file</i> Sets scheduler parameters for the specified class using the values in <i>file</i>. These values overwrite the current values in memory—they become the parameters that control scheduling of processes in the specified class. The values in <i>file</i> must be in the format output by the -g option. Moreover, the values must describe a table that is the same size (has same number of priority levels) as the table being overwritten. The <code>sys_config</code> privilege is required for the -s option to succeed.</p> <p>Note: The -g and -s options are mutually exclusive: you may not retrieve the table at the same time you are overwriting it.</p>

dispadmin(1M)

EXAMPLES

EXAMPLE 1 Retrieving the current scheduler parameters for the real-time class.

The following command retrieves the current scheduler parameters for the real-time class from kernel memory and writes them to the standard output. Time quantum values are in microseconds.

```
dispadmin -c RT -g -r 1000000
```

EXAMPLE 2 Overwriting the current scheduler parameters for the real-time class.

The following command overwrites the current scheduler parameters for the real-time class with the values specified in `rt.config`.

```
dispadmin -c RT -s rt.config
```

EXAMPLE 3 Retrieving the current scheduler parameters for the time-sharing class.

The following command retrieves the current scheduler parameters for the time-sharing class from kernel memory and writes them to the standard output. Time quantum values are in nanoseconds.

```
dispadmin -c TS -g -r 1000000000
```

EXAMPLE 4 Overwriting the current scheduler parameters for the time-sharing class.

The following command overwrites the current scheduler parameters for the time-sharing class with the values specified in `ts.config`.

```
dispadmin -c TS -s ts.config
```

SUMMARY OF TRUSTED ATTRIBUTES CHANGES

To succeed with the `-s` option, this command needs the `sys_config` privilege.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8 4/01 Reference Manual Reference Manual DIAGNOSTICS

`priocntl(2)`

`priocntl(1)`, `rt_dptbl(4)`, `ts_dptbl(4)`, `attributes(5)`

`dispadmin` prints an appropriate diagnostic message if it fails to overwrite the current scheduler parameters due to lack of required permissions or a problem with the specified input file.

NAME	dl_booting, dl_restore – Inform the kernel that a machine is in the state of disklessly booting or in the normal state				
SYNOPSIS	<pre>/usr/sbin/dl_booting [hostname ip_address]</pre> <pre>/usr/sbin/dl_restore [hostname ip_address]</pre>				
DESCRIPTION	<p>dl_booting informs the kernel that the machine specified by <i>hostname</i> or <i>ip_address</i> is in the state of booting disklessly. Hence, until the kernel is notified that the machine has reverted to the normal state, it must be viewed as an unlabeled host, and only processes with the PAF_DISKLESS_BOOT process attribute can communicate with the machine while it is in the booting state. In the normal state, packets exchanged are properly labeled.</p> <p>dl_restore informs the kernel that the machine specified by the <i>hostname</i> or IP address is now in the normal state.</p> <p>To succeed, both dl_booting and dl_restore must inherit the sys_net_config privilege.</p>				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWtsu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWtsu				
SUMMARY OF TRUSTED SOLARIS CHANGES	To succeed, both dl_booting and dl_restore must inherit the sys_net_config privilege.				
Trusted Solaris 8 4/01 Reference Manual	chstate(2)				
Solaris 9 4/01 Reference Manual	attributes(5)				

dl_restore(1M)

NAME dl_booting, dl_restore – Inform the kernel that a machine is in the state of disklessly booting or in the normal state

SYNOPSIS **/usr/sbin/dl_booting** [*hostname* | *ip_address*]
/usr/sbin/dl_restore [*hostname* | *ip_address*]

DESCRIPTION dl_booting informs the kernel that the machine specified by *hostname* or *ip_address* is in the state of booting disklessly. Hence, until the kernel is notified that the machine has reverted to the normal state, it must be viewed as an unlabeled host, and only processes with the PAF_DISKLESS_BOOT process attribute can communicate with the machine while it is in the booting state. In the normal state, packets exchanged are properly labeled.

dl_restore informs the kernel that the machine specified by the *hostname* or IP address is now in the normal state.

To succeed, both dl_booting and dl_restore must inherit the sys_net_config privilege.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWtsu

SUMMARY OF TRUSTED SOLARIS CHANGES To succeed, both dl_booting and dl_restore must inherit the sys_net_config privilege.

Trusted Solaris 8
4/01 Reference Manual
SunOS 5.8
Reference Manual

chstate(2)
attributes(5)

NAME	dminfo – Report information about a device entry in a device maps file	
SYNOPSIS	dminfo [-v] [-a] [-f <i>pathname</i>] dminfo [-v] [-a] [-f <i>pathname</i>] -n <i>dev</i> -name... dminfo [-v] [-a] [-f <i>pathname</i>] -d <i>dev</i> -path... dminfo [-v] [-a] [-f <i>pathname</i>] -t <i>dev</i> -type... dminfo [-v] [-f <i>pathname</i>] -u <i>dm</i> -entry	
DESCRIPTION	dminfo reports and updates information about the device_maps(4) file.	
OPTIONS	<p>-v Verbose. Print the requested entry or entries, one line per entry, on the standard output. If no entries are specified, all are printed.</p> <p>-a Succeed if any of the requested entries are found. If used with -v, all entries that match the requested case(s) are printed.</p> <p>-f <i>pathname</i> Use a device_maps file with <i>pathname</i> instead of /etc/security/device_maps.</p> <p>-n <i>dev-name</i> Search by <i>dev-name</i>. Search device_maps(4) for a <i>device_name</i> field matching <i>dev-name</i>. This option cannot be used with -d, -t, or -u.</p> <p>-d <i>dev-path</i> Search by <i>dev-path</i>. Search device_maps(4) for a device special pathname in the <i>device_list</i> field matching the <i>dev-path</i> argument. This option cannot be used with -n, -t, or -u.</p> <p>-t <i>dev-type</i> Search by <i>dev-type</i>. Search device_maps(4) for a <i>device_type</i> field matching the given <i>dev-type</i>. This option cannot be used with -d, -n, or -u.</p> <p>-u <i>dm-entry</i> Update the device_maps(4) file. This option is provided to add entries to the device_maps(4) file. The <i>dm-entry</i> must be a complete device_maps file entry. The <i>dm-entry</i> has fields, as in the device_maps file. It uses the colon (:) as a field separator, and white space as the <i>device_list</i> subfield separators. The <i>dm-entry</i> is not made if any fields are missing, or if the <i>dm-entry</i> would be a duplicate. This option requires the trusted path and write access to the device_maps file.</p>	
DIAGNOSTICS	dminfo returns an exit code of 0 if successful, 1 if the request failed, and 2 if the invocation syntax was incorrect.	
SUMMARY OF TRUSTED SOLARIS CHANGES	The -u option requires the trusted path and write access to the /etc/security/device_maps file. The calling process may use the privileges file_mac_write and file_dac_write to override access restrictions.	
FILES	/etc/security/device_maps	List of physical devices associated with a device name and type.

dminfo(1M)

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

device_maps(4)

attributes(5)

NAME	drvconfig – Configure the /devices directory
SYNOPSIS	drvconfig [-bn] [-a <i>alias_name</i>] [-c <i>class_name</i>] [-i <i>drivername</i>] [-m <i>major_num</i>] [-r <i>rootdir</i>]
DESCRIPTION	<p>devfsadm(1M) is now the preferred command for /dev and /devices and should be used instead of drvconfig.</p> <p>The default operation of drvconfig is to create the /devices directory tree that describes, in the filesystem namespace, the hardware layout of a particular machine. Hardware devices present on the machine and powered on as well as pseudo-drivers are represented under /devices. Normally this command is run automatically after a new driver has been installed (with add_drv(1M)) and the system has been rebooted.</p>
/etc/minor_perm file	<p>drvconfig reads the /etc/minor_perm file to obtain permission information and applies the permissions only to nodes that it has just created. It does not change permissions on already existing nodes. The format of the /etc/minor_perm file is as follows:</p> <pre>name : minor_name permissions owner group</pre> <p><i>minor_name</i> may be the actual name of the minor node, or contain shell metacharacters to represent several minor nodes (see sh(1)).</p> <p>For example:</p> <pre>sd:* 0640 root sys zs:[a-z],cu 0600 uucp uucp mm:kmem 0640 root bin</pre> <p>The first line sets all devices exported by the sd node to 0640 permissions, owned by root, with group sys. In the second line, devices such as a, cu and z, cu exported by the zs driver are set to 0600 permission, owned by uucp, with group uucp. In the third line the kmem device exported by the mm driver is set to 0640 permission, owned by root, with group bin.</p>
/etc/security/tsol/minor_perm.adjunct file	<p>drvconfig reads the /etc/security/tsol/minor_perm.adjunct file to obtain label information and applies the labels to nodes that it has just created. drvconfig does not change labels on already existing nodes. The format of the file is:</p> <pre>name : minor_name [SL]</pre> <p><i>minor_name</i> is the name of the minor node; shell metacharacters may be used to represent several minor nodes (see sh(1)). Labels can be represented in hex or string format.</p> <pre>SD:* [admin_high] mm:zero [admin_low]</pre>

drvconfig(1M)

	The above example sets all devices exported by the <code>sd</code> node to have a sensitivity label of <code>ADMIN_HIGH</code> . The <code>zero</code> device exported by the <code>mm</code> driver is set to have a sensitivity label of <code>ADMIN_LOW</code> .	
OPTIONS	The following options may be of use to system administrators and driver developers:	
	<code>-i drivename</code>	Only configure the devices for the named driver. The following options are used by the implementation of <code>add_drv(1M)</code> and <code>rem_drv(1M)</code> , and may not be supported in future versions of the Solaris and Trusted Solaris environments.
	<code>-b</code>	Add a new major number to name binding into the kernel's internal <code>name_to_major</code> tables. This option is not normally used directly, but is used by other utilities such as <code>add_drv(1M)</code> . Use of the <code>-b</code> option requires that <code>-i</code> and <code>-m</code> be used also. No <code>/devices</code> entries are created.
	<code>-n</code>	Do not try to load and attach any drivers, or if the <code>-i</code> option is given, do not try to attach the driver named <i>drivename</i> .
	<code>-a alias_name</code>	Add the name <i>alias_name</i> to the list of aliases that this driver is known by. This option, if used, must be used with the <code>-m major_num</code> , the <code>-b</code> and the <code>-i drivename</code> options.
	<code>-c class_name</code>	The driver being added to the system exports the class <i>class_name</i> . This option is not normally used directly, but is used by other utilities. It is only effective when used with the <code>-b</code> option.
	<code>-m major_num</code>	Specify the major number <i>major_num</i> for this driver to add to the kernel's <code>name_to_major</code> binding tables.
	<code>-r rootdir</code>	Build the device tree under the directory specified by <i>rootdir</i> instead of the default <code>/devices</code> directory.
EXIT STATUS	<code>0</code>	Successful completion.
	non-zero	An error occurred.
FILES	<code>/devices</code>	Device nodes directory
	<code>/etc/minor_perm</code>	Minor mode permissions
	<code>/etc/security/tsol/minor_perm.adjunct</code>	Default label
	<code>/etc/name_to_major</code>	Major number binding
	<code>/etc/driver_classes</code>	Driver class binding file
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:	

NOTES

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

The /etc/security/tsol/minor_perm.adjunct file is used to record the sensitivity label of devices.

add_drv(1M), modload(1M), modunload(1M), rem_drv(1M)

sh(1), devlinks(1M), devfsadm(1M), disks(1M), modinfo(1M), ports(1M), tapes(1M), path_to_inst(4)attributes(5)

This document does not constitute an API. /etc/minor_perm, /etc/security/tsol/minor_perm.adjunct, /etc/name_to_major, /etc/driver_classes, and /devices may not exist or may have different contents or interpretations in a future release. The existence of this notice does not imply that any other documentation that lacks this notice constitutes an API.

drvconfig(1M)

du(1M)

NAME	du – Summarize disk usage
SYNOPSIS	<pre>/usr/bin/du [-adkLr] [-s -o] [-M] [file...] /usr/xpg4/bin/du [-a -s] [-krx] [file...]</pre>
DESCRIPTION	<p>The <code>du</code> utility writes to standard output the size of the file space allocated to, and the size of the file space allocated to each subdirectory of, the file hierarchy rooted in each of the specified files. The size of the file space allocated to a file of type directory is defined as the sum total of space allocated to all files in the file hierarchy rooted in the directory plus the space allocated to the directory itself.</p> <p>Files with multiple links will be counted and written for only one entry. The directory entry that is selected in the report is unspecified. By default, file sizes are written in 512-byte units, rounded up to the next 512-byte unit.</p>
/usr/xpg4/bin/du	When <code>du</code> cannot obtain file attributes or read directories (see <code>stat(2)</code>), it will report an error condition and the final exit status will be affected.
OPTIONS	<p>The following options are supported for <code>/usr/bin/du</code> and <code>/usr/xpg4/bin/du</code>:</p> <ul style="list-style-type: none">-a In addition to the default output, report the size of each file not of type directory in the file hierarchy rooted in the specified file. Regardless of the presence of the <code>-a</code> option, non-directories given as <i>file</i> operands will always be listed.-k Write the files sizes in units of 1024 bytes, rather than the default 512-byte units.-s Instead of the default output, report only the total sum for each of the specified files.
/usr/bin/du	<p>The following options are supported for <code>/usr/bin/du</code> only:</p> <ul style="list-style-type: none">-d Do not cross filesystem boundaries. For example, <code>du -d /</code> reports usage only on the root partition.-L Process symbolic links by using the file or directory which the symbolic link references, rather than the link itself.-o Do not add child directories' usage to a parent's total. Without this option, the usage listed for a particular directory is the space taken by the files in that directory, as well as the files in all directories beneath it. This option does nothing if <code>-s</code> is used.-r Generate messages about directories that cannot be read, files that cannot be opened, and so forth, rather than being silent (the default).-M Process all accessible single-level directories while descending multilevel directories.
/usr/xpg4/bin/du	The following options are supported for <code>/usr/xpg4/bin/du</code> only:

	<p>-r By default, generate messages about directories that cannot be read, files that cannot be opened, and so forth.</p> <p>-x When evaluating file sizes, evaluate only those files that have the same device as the file specified by the <i>file</i> operand.</p>						
OPERANDS	<p>The following operand is supported:</p> <p><i>file</i> The path name of a file whose size is to be written. If <i>file</i> is not specified, the current directory is used.</p>						
OUTPUT	The output from <code>du</code> consists of the amount of the space allocated to a file and the name of the file.						
USAGE	See <code>largefile(5)</code> for the description of the behavior of <code>du</code> when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).						
ENVIRONMENT VARIABLES	See <code>environ(5)</code> for descriptions of the following environment variables that affect the execution of <code>du</code> : <code>LC_CTYPE</code> , <code>LC_MESSAGES</code> , and <code>NLSPATH</code> .						
EXIT STATUS	<p>The following exit values are returned:</p> <p>0 Successful completion.</p> <p>>0 An error occurred.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
/usr/bin/du	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWcsu</td></tr> <tr> <td>CSI</td><td>enabled</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu	CSI	enabled
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWcsu						
CSI	enabled						
/usr/xpg4/bin/du	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWxcu4</td></tr> <tr> <td>CSI</td><td>enabled</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWxcu4	CSI	enabled
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWxcu4						
CSI	enabled						
SUMMARY OF TRUSTED SQUARES CHANGES	<p>The <code>-M</code> option processes SLDs when descending multilevel directories.</p> <p><code>stat(2)</code></p> <p><code>ls(1)</code>, <code>attributes(5)</code>, <code>environ(5)</code>, <code>largefile(5)</code>, <code>XPG4(5)</code></p> <p><i>System Administration Guide, Volume 1</i></p>						

du(1M)

NOTES | A file with two or more links is counted only once. If, however, there are links between files in different directories where the directories are on separate branches of the file system hierarchy, `du` will count the excess files more than once.

Files containing holes will result in an incorrect block count.

NAME	eeprom – EEPROM display and load utility	
SPARC	/usr/platform/ <i>platform-name</i> /sbin/eeprom [-] [-f <i>device</i>] [<i>parameter</i> [=value] ...]	
IA	/usr/platform/ <i>platform-name</i> /sbin/eeprom [-] [-f <i>device</i>] [-I] [<i>mmu-modlist</i>] [<i>parameter</i> [=value]]	
DESCRIPTION	<p>eeprom displays or changes the values of parameters in the EEPROM. It processes parameters in the order given. When processing a <i>parameter</i> accompanied by a <i>value</i>, eeprom makes the indicated alteration to the EEPROM; otherwise it displays the <i>parameter</i>'s value. When given no parameter specifiers, eeprom displays the values of all EEPROM parameters. A ' - ' (hyphen) flag specifies that parameters and values are to be read from the standard input (one <i>parameter</i> or <i>parameter=value</i> per line).</p> <p>eeprom verifies the EEPROM checksums and complains if they are incorrect.</p> <p><i>platform-name</i> is the name of the platform implementation and can be found using the -i option of uname(1).</p>	
SPARC	SPARC based systems implement firmware password protection with eeprom using the security-mode, security-password, and security-#badlogins properties.	
IA	<p>EEPROM storage is simulated using a file residing in the platform specific boot area. The /platform/<i>platform-name</i>/boot/solaris/bootenv.rc file simulates EEPROM storage.</p> <p>Because IA based systems typically implement password protection in the system BIOS, there is no support for password protection in the eeprom program. While it is possible to set the security-mode, security-password, and security-#badlogins properties on IA based systems, these properties have no special meaning or behavior on IA based systems.</p>	
OPTIONS	-f <i>device</i>	Use <i>device</i> as the EEPROM device.
IA Only	-I	Initialize boot properties on an IA based system. Only init(1M) run-level initialization scripts should use this option.
IA Only	<i>acpi-user-options</i>	A configuration variable that controls the use of ACPI. A value of 0x0 attempts to use ACPI if it is available on the system. A value of 0x2 disables the use of ACPI. The default value is 0x0.
	<i>mmu-modlist</i>	A colon-separated list of candidate modules that implement memory management. If <i>mmu-modlist</i> is defined, it overrides the default list derived from the memory configuration on IA based systems. Instead, the first module in the list that is found in /platform/ <i>platform-name</i> /kernel/mmu is used.

eeeprom(1M)

**NVRAM
CONFIGURATION
PARAMETERS**

Not all OpenBoot systems support all parameters. Defaults may vary depending on the system and the PROM revision.

auto-boot?	If true, boot automatically after power-on or reset. Defaults to true.
ansi-terminal?	Configuration variable used to control the behavior of the terminal emulator. The value <code>false</code> makes the terminal emulator stop interpreting ANSI escape sequences, instead just echoing them to the output device. Defaults to true.
boot-command	Command executed if <code>auto-boot?</code> is true. The default value is <code>boot</code> .
boot-device	Device from which to boot. <i>boot-device</i> may contain 0 or more device specifiers separated by spaces. Each device specifier may be either a prom device alias or a prom device path. The boot prom will attempt to open each successive device specifier in the list beginning with the first device specifier. The first device specifier which opens successfully will be used as the device to boot from. Defaults to <code>disk net</code> .
boot-file	File to boot (an empty string lets the secondary booter choose default). Defaults to an empty string.
boot-from	Boot device and file (OpenBoot PROM version 1.x only). Defaults to <code>vmunix</code> .
boot-from-diag	Diagnostic boot device and file (OpenBoot PROM version 1.x only). Defaults to <code>le() unix</code> .
comX-noprobe	Where X is the number of the serial port, prevents device probe on serial port X.
diag-device	Diagnostic boot source device. Defaults to <code>net</code> .
diag-file	File from which to boot in diagnostic mode. Defaults to an empty string.
diag-level	Diagnostics level. Values include <code>off</code> , <code>min</code> , <code>max</code> , and <code>menus</code> . There may be additional platform-specific values. When set to <code>off</code> , POST is not called. If POST is called, the value is made available as an argument to, and is interpreted by POST. Defaults to platform-dependent.
diag-switch?	If true, run in diagnostic mode. Defaults to true.
fcode-debug?	If true, include name parameter for plug-in device FCodes. Defaults to false.
hardware-revision	System version information.

eeeprom(1M)

input-device	Input device used at power-on (usually keyboard, ttya, or ttyb). . Defaults to keyboard.
keyboard-click?	If true enable keyboard click. Defaults to false.
keymap	Keymap for custom keyboard.
last-hardware-update	System update information.
load-base	Default load address for client programs. Defaults to 16384.
local-mac-address?	If true, network drivers use their own MAC address, not system's. Defaults to false.
mfg-mode	Manufacturing mode argument for POST. Possible values include off or chamber. The value is passed as an argument to POST. Defaults to off.
mfg-switch?	If true, repeat system self-tests until interrupted with STOP-A . Defaults to false.
nvrामrc	Contents of NVRAMRC. Defaults to empty.
oem-banner	Custom OEM banner (enabled by setting oem-banner? to true). Defaults to an empty string.
oem-banner?	If true, use custom OEM banner. Defaults to false.
oem-logo	Byte array custom OEM logo (enabled by setting oem-logo? to true). Displayed in hexadecimal.
oem-logo?	If true, use custom OEM logo (else, use Sun logo). Defaults to false.
output-device	Output device used at power-on (usually screen, ttya, or ttyb). Defaults to screen.
sbus-probe-list	Which SBus slots are probed and in what order. Defaults to 0123.
screen-#columns	Number of on-screen columns (characters/line). Defaults to 80.
screen-#rows	Number of on-screen rows (lines). Defaults to 34.
scsi-initiator-id	SCSI bus address of host adapter, range 0-7. Defaults to 7.
sd-targets	Map SCSI disk units (OpenBoot PROM version 1.x only). The default is 31204567, which means that unit 0 maps to target 3, unit 1 maps to target 1, and so on.
security-#badlogins	Number of incorrect security password attempts.

eeeprom(1M)

	This property has no special meaning or behavior on IA based systems.
security-mode	Firmware security level (options: none, command, or full). If set to command or full, system will prompt for PROM security password. Defaults to none.
security-password	<p>This property has no special meaning or behavior on IA based systems.</p> <p>Firmware security password (never displayed). Can be set only when security-mode is set to command or full.</p> <p>This property has no special meaning or behavior on IA based systems.</p> <pre>example# eeeprom security-password= Changing PROM password: New password: Retype new password:</pre>
selftest-#megs	Megabytes of RAM to test. Ignored if diag-switch? is true. Defaults to 1.
skip-vme-loopback?	If true, POST does not do VMEbus loopback tests. Defaults to false.
st-targets	Map SCSI tape units (OpenBoot PROM version 1.x only). Defaults to 45670123, which means that unit 0 maps to target 4, unit 1 maps to target 5, and so on.
sunmon-compatible?	If true, display Restricted Monitor prompt (>). Defaults to false.
testarea	One-byte scratch field, available for read/write test. Defaults to 0.
tpe-link-test?	Enable 10baseT link test for built-in twisted pair Ethernet. Defaults to true.
ttya-mode	<p>TTYA (baud rate, #bits, parity, #stop, handshake). Defaults to 9600, 8, n, 1, -.</p> <p>Fields, in left-to-right order, are:</p> <pre>baud rate: 110, 300, 1200, 4800, 9600 . . . data bits: 5, 6, 7, 8 parity: n(none), e(even), o(odd), m(mark), s(space) stop bits: 1, 1.5, 2</pre>

	handshake:	–(none), h(hardware:rts/cts), s(software:xon/xoff)
tttyb-mode	TTYB (baud rate, #bits, parity, #stop, handshake). Defaults to 9600, 8, n, 1, –.	
	Fields, in left-to-right order, are:	
	baud rate:	110, 300, 1200, 4800, 9600 . . .
	data bits:	5, 6, 7, 8
	stop bits:	1, 1.5, 2
	parity:	n(none), e(even), o(odd), m(mark), s(space)
	handshake:	–(none), h(hardware:rts/cts), s(software:xon/xoff)
tttya-ignore-cd	If true, operating system ignores carrier-detect on TTYA. Defaults to true.	
tttyb-ignore-cd	If true, operating system ignores carrier-detect on TTYB. Defaults to true.	
tttya-rts-dtr-off	If true, operating system does not assert DTR and RTS on TTYA. Defaults to false.	
tttyb-rts-dtr-off	If true, operating system does not assert DTR and RTS on TTYB. Defaults to false.	
use-nvramrc?	If true, execute commands in NVRAMRC during system start-up. Defaults to false.	
version2?	If true, hybrid (1.x/2.x) PROM comes up in version 2.x. Defaults to true.	
watchdog-reboot?	If true, reboot after watchdog reset. Defaults to false.	

EXAMPLES

EXAMPLE 1 Changing the number of megabytes of RAM.

The following example demonstrates the method for changing from one to two the number of megabytes of RAM that the system will test.

```
example# eeeprom selftest-#megs
selftest-#megs=1
```

```
example# eeeprom selftest-#megs=2
```

```
example# eeeprom selftest-#megs
selftest-#megs=2
```

eeeprom(1M)

SUMMARY OF
TRUSTED
SOLARIS
CHANGES

ATTRIBUTES

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

EXAMPLE 1 Changing the number of megabytes of RAM. (Continued)

EXAMPLE 2 Setting the auto-boot? parameter to true.

The following example demonstrates the method for setting the auto-boot? parameter to true.

```
example# eeeprom auto-boot?=true
```

When the eeeprom command is executed in user mode, the parameters with a trailing question mark (?) need to be enclosed in double quotation marks (" ") to prevent the shell from interpreting the question mark. Preceding the question mark with an escape character (\) will also prevent the shell from interpreting the question mark.

```
example% eeeprom "auto-boot?"=true
```

For administrative users who alter the EEPROM contents, this command must be invoked with effective user ID of 0.

/dev/openprom
device file

/usr/platform/*platform-name*/sbin/eeeprom
Platform-specific version of eeeprom. Use **uname -i.** to obtain *platform-name*.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

uname(1)

passwd(1), sh(1), attributes(5)

NAME	format – disk partitioning and maintenance utility
SYNOPSIS	format [-f <i>command-file</i>] [-l <i>log-file</i>] [-x <i>data-file</i>] [-d <i>disk-name</i>] [-t <i>disk-type</i>] [-p <i>partition-name</i>] [-s] [-m] [-M] [-e] [<i>disk-list</i>]
DESCRIPTION	<p>format enables you to format, label, repair and analyze disks on your system. Unlike previous disk maintenance programs, <i>format</i> runs under SunOS. Because there are limitations to what can be done to the system disk while the system is running, <i>format</i> is also supported within the memory-resident system environment. For most applications, however, running <i>format</i> under SunOS is the more convenient approach.</p> <p><i>format</i> first uses the disk list defined in <i>data-file</i> if the -x option is used. <i>format</i> then checks for the FORMAT_PATH environment variable, a colon-separated list of filenames and/or directories. In the case of a directory, <i>format</i> searches for a file named <i>format.dat</i> in that directory; a filename should be an absolute pathname, and is used without change. <i>format</i> adds all disk and partition definitions in each specified file to the working set. Multiple identical definitions are silently ignored. If FORMAT_PATH is not set, the path defaults to /etc/format.dat.</p> <p><i>disk-list</i> is a list of disks in the form c?t?d? or /dev/rdisk/c?t?d?s?. With the latter form shell wildcard specifications are supported. For example, specifying /dev/rdisk/c2* will cause <i>format</i> to work on all drives connected to controller c2 only. If no <i>disk-list</i> is specified, <i>format</i> lists all the disks present in the system.</p> <p>Removable media devices are listed only when users execute <i>format</i> in expert mode (option -e). This feature is provided for backward compatibility. Use <i>rmformat</i>(1) for rewritable removable media devices.</p>
OPTIONS	<p>The following options are supported:</p> <ul style="list-style-type: none"> -d <i>disk-name</i> Specify which disk should be made current upon entry into the program. The disk is specified by its logical name (for instance, -d c0t1d0). This can also be accomplished by specifying a single disk in the disk list. -e Enable SCSI expert menu. Note this option is not recommended for casual use. -f <i>command-file</i> Take command input from <i>command-file</i> rather than the standard input. The file must contain commands that appear just as they would if they had been entered from the keyboard. With this option, <i>format</i> does not issue continue? prompts; there is no need to specify y(es) or n(o) answers in the <i>command-file</i>. In non-interactive mode, <i>format</i> does not initially expect the input of a disk selection number. The user must specify the current working disk with the -d <i>disk-name</i> option when <i>format</i> is invoked, or specify <i>disk</i> and the disk selection number in the <i>command-file</i>.

format(1M)

	-l <i>log-file</i>	Log a transcript of the <code>format</code> session to the indicated <i>log-file</i> , including the standard input, the standard output and the standard error.
	-m	Enable extended messages. Provides more detailed information in the event of an error.
	-M	Enable extended and diagnostic messages. Provides extensive information on the state of a SCSI device's mode pages, during formatting.
	-p <i>partition-name</i>	Specify the partition table for the disk which is current upon entry into the program. The table is specified by its name as defined in the data file. This option can only be used if a disk is being made current, and its type is either specified or available from the disk label.
	-t <i>disk-type</i>	Specify the type of disk which is current upon entry into the program. A disk's type is specified by name in the data file. This option can only be used if a disk is being made current as described above.
	-s	Silent. Suppress all of the standard output. Error messages are still displayed. This is generally used in conjunction with the -f option.
	-x <i>data-file</i>	Use the list of disks contained in <i>data-file</i> .
USAGE	The <code>format</code> utility's main menu items allow you to do the following tasks:	
	analyze	Run read, write, and compare tests.
	backup	Search for backup labels.
	cache	Enable , disable and query the state of the write cache and read cache. This menu item only appears when <code>format</code> is invoked with the -e option, and is only supported on SCSI devices..
	current	Display the device name, the disk geometry, and the pathname to the disk device.
	defect	Retrieve and print defect lists.
	disk	Choose the disk that will be used in subsequent operations (known as the current disk).
	fdisk	Run the <code>fdisk(1M)</code> program to create an <code>fdisk</code> partition (IA based systems only).
	format	Format and verify the current disk.
	inquiry	Display the vendor, product name, and revision level of the current drive.
	label	Write a new label to the current disk.

ENVIRONMENT VARIABLES	partition	Create and modify slices.				
	quit	Exit the format menu.				
	repair	Repair a specific block on the disk.				
	save	Save new disk and slice information.				
	type	Select (define) a disk type.				
	verify	Read and display labels. Print information such as the number of cylinders, alternate cylinders, heads, sectors, and the partition table.				
	volname	Label the disk with a new eight character volume name.				
FILES	FORMAT_PATH	A colon-separated list of filenames and/or directories of disk and partition definitions. If a directory is specified, format searches for the file format.dat in that directory.				
ATTRIBUTES	/etc/format.dat	Default data file				
	See attributes(5) for descriptions of the following attributes:					
SUMMARY OF TRUSTED SOLARIS CHANGES SunOS 5.8 Reference Manual	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWcsu</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
	ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu					
	To succeed, the format command requires the sys_devices privilege and a UID of 0.					
	fmthard(1M), prtvtoc(1M), format.dat(4), attributes(5), sd(7D)					
	See Disk Management in System Administration Guide, Volume 1					
IA Only	fdisk(1M)					
WARNINGS	When the format utility is selected to format the Maxtor 207MB disk, the following message displays: Mode sense page(4) reports rpm value as 0, adjusting it to 3600 This is a drive bug that may also occur with older third party drives. The above message is not an error; the drive will still function correctly.					
NOTES	format provides a help facility you can use whenever format is expecting input. You can request help about what information is expected by simply entering a question mark (?) and format prints a brief description of what type of input is needed. If you enter a ? at the menu prompt, a list of available commands is displayed.					

format(1M)

For SCSI disks, formatting is done with both Primary and Grown defects list by default. However, if only Primary list is extracted in defect menu before formatting, formatting will be done with Primary list only.

Changing the state of the caches is only supported on SCSI devices, and not all SCSI devices support changing or saving the state of the caches.

NAME	fsdb_ufs – ufs File System Debugger
SYNOPSIS	fsdb -F ufs [<i>generic_options</i>] [<i>specific_options</i>] <i>special</i>
DESCRIPTION	<p>The <code>fsdb_ufs</code> command is an interactive tool that can be used to patch up a damaged UFS file system. It has conversions to translate block and i-numbers into their corresponding disk addresses. Also included are mnemonic offsets to access different parts of an inode. These greatly simplify the process of correcting control block entries or descending the file system tree.</p> <p><code>fsdb</code> contains several error-checking routines to verify inode and block addresses. These can be disabled if necessary by invoking <code>fsdb</code> with the <code>-o</code> option or by the use of the <code>o</code> command.</p> <p><code>fsdb</code> reads a block at a time and will therefore work with raw as well as block I/O devices. A buffer management routine is used to retain commonly used blocks of data in order to reduce the number of read system calls. All assignment operations result in an immediate write-through of the corresponding block. Note that in order to modify any portion of the disk, <code>fsdb</code> must be invoked with the <code>w</code> option.</p> <p>Wherever possible, <code>adb</code>-like syntax was adopted to promote the use of <code>fsdb</code> through familiarity.</p>
OPTIONS	<p>The following option is supported:</p> <ul style="list-style-type: none"> -o Specify UFS file system specific options. These options can be any combination of the following separated by commas (with no intervening spaces). The options available are: <ul style="list-style-type: none"> ? Display usage o Override some error conditions p='string' Set prompt to string w Open for write
USAGE	<p>Numbers are considered hexadecimal by default. However, the user has control over how data is to be displayed or accepted. The <code>base</code> command will display or set the input/output base. Once set, all input will default to this base and all output will be shown in this base. The base can be overridden temporarily for input by preceding hexadecimal numbers with <code>'0x'</code>, preceding decimal numbers with <code>'0t'</code>, or octal numbers with <code>'0'</code>. Hexadecimal numbers beginning with <code>a-f</code> or <code>A-F</code> must be preceded with <code>'0x'</code> to distinguish them from commands.</p> <p>Disk addressing by <code>fsdb</code> is at the byte level. However, <code>fsdb</code> offers many commands to convert a desired inode, directory entry, block, superblock and so forth to a byte address. Once the address has been calculated, <code>fsdb</code> will record the result in dot (<code>.</code>).</p> <p>Several global values are maintained by <code>fsdb</code>:</p> <ul style="list-style-type: none"> ■ the current base (referred to as <code>base</code>),

fsdb_ufs(1M)

- the current address (referred to as `dot`),
- the current inode (referred to as `inode`),
- the current count (referred to as `count`),
- and the current type (referred to as `type`).

Most commands use the preset value of `dot` in their execution. For example,

```
> 2:inode
```

will first set the value of `dot` to 2, `'.'`, will alert the start of a command, and the `inode` command will set `inode` to 2. A count is specified after a `'.'`. Once set, `count` will remain at this value until a new command is encountered which will then reset the value back to 1 (the default). So, if

```
> 2000,400/X
```

is typed, 400 hex longs are listed from 2000, and when completed, the value of `dot` will be `2000 + 400 * sizeof (long)`. If a `RETURN` is then typed, the output routine will use the current values of `dot`, `count`, and `type` and display 400 more hex longs. A `'*'` will cause the entire block to be displayed.

End of fragment, block and file are maintained by `fsdb`. When displaying data as fragments or blocks, an error message will be displayed when the end of fragment or block is reached. When displaying data using the `db`, `ib`, `directory`, or `file` commands an error message is displayed if the end of file is reached. This is mainly needed to avoid passing the end of a directory or file and getting unknown and unwanted results.

An example showing several commands and the use of `RETURN` would be:

```
> 2:ino; 0:dir?d
      or
> 2:ino; 0:db:block?d
```

The two examples are synonymous for getting to the first directory entry of the root of the file system. Once there, any subsequent `RETURN` (or `+`, `-`) will advance to subsequent entries. Note that

```
> 2:inode; :ls
      or
> :ls /
```

is again synonymous.

Expressions

The symbols recognized by `fsdb` are:

<code>RETURN</code>	update the value of <code>dot</code> by the current value of <code>type</code> and display using the current value of <code>count</code> .
---------------------	--

#	numeric expressions may be composed of +, -, *, and % operators (evaluated left to right) and may use parentheses. Once evaluated, the value of <code>dot</code> is updated.
, <i>count</i>	count indicator. The global value of <code>count</code> will be updated to <code>count</code> . The value of <code>count</code> will remain until a new command is run. A count specifier of '*' will attempt to show a <i>blocks's</i> worth of information. The default for <code>count</code> is 1.
? <i>f</i>	display in structured style with format specifier <i>f</i> . See <code>FormattedOutput</code> .
/ <i>f</i>	display in unstructured style with format specifier <i>f</i> . See <code>FormattedOutput</code> .
.	the value of <code>dot</code> .
+ <i>e</i>	increment the value of <code>dot</code> by the expression <i>e</i> . The amount actually incremented is dependent on the size of <code>type</code> : <code>dot = dot + e * sizeof (type)</code> The default for <i>e</i> is 1.
- <i>e</i>	decrement the value of <code>dot</code> by the expression <i>e</i> . See +.
* <i>e</i>	multiply the value of <code>dot</code> by the expression <i>e</i> . Multiplication and division don't use <code>type</code> . In the above calculation of <code>dot</code> , consider the <code>sizeof (type)</code> to be 1.
% <i>e</i>	divide the value of <code>dot</code> by the expression <i>e</i> . See *.
< <i>name</i>	restore an address saved in register <i>name</i> . <i>name</i> must be a single letter or digit.
> <i>name</i>	save an address in register <i>name</i> . <i>name</i> must be a single letter or digit.
= <i>f</i>	display indicator. If <i>f</i> is a legitimate format specifier, then the value of <code>dot</code> is displayed using the format specifier <i>f</i> . See <code>FormattedOutput</code> . Otherwise, assignment is assumed. See =.
= [<i>s</i>] [<i>e</i>]	assignment indicator. The address pointed to by <code>dot</code> has its contents changed to the value of the expression <i>e</i> or to the ASCII representation of the quoted (") string <i>s</i> . This may be useful for changing directory names or ASCII file information.
=+ <i>e</i>	incremental assignment. The address pointed to by <code>dot</code> has its contents incremented by expression <i>e</i> .
=- <i>e</i>	decremental assignment. The address pointed to by <code>dot</code> has its contents decremented by expression <i>e</i> .

Commands	<p>A command must be prefixed by a ':' character. Only enough letters of the command to uniquely distinguish it are needed. Multiple commands may be entered on one line by separating them by a SPACE, TAB or ';'.</p> <p>In order to view a potentially unmounted disk in a reasonable manner, fsdb offers the <code>cd</code>, <code>pwd</code>, <code>ls</code> and <code>find</code> commands. The functionality of these commands substantially matches those of its UNIX counterparts. See individual commands for details. The '*', '?', and '['-'] wild card characters are available.</p> <p><code>base=b</code> display or set base. As stated above, all input and output is governed by the current base. If the <code>=b</code> is omitted, the current base is displayed. Otherwise, the current base is set to <i>b</i>. Note that this is interpreted using the old value of base, so to ensure correctness use the '0', '0t', or '0x' prefix when changing the base. The default for base is hexadecimal.</p> <p><code>block</code> convert the value of <code>dot</code> to a block address.</p> <p><code>cd dir</code> change the current directory to directory <i>dir</i>. The current values of <code>inode</code> and <code>dot</code> are also updated. If no <i>dir</i> is specified, then change directories to inode 2 ("/").</p> <p><code>cg</code> convert the value of <code>dot</code> to a cylinder group.</p> <p><code>directory</code> If the current <code>inode</code> is a directory, then the value of <code>dot</code> is converted to a directory slot offset in that directory and <code>dot</code> now points to this entry.</p> <p><code>file</code> the value of <code>dot</code> is taken as a relative block count from the beginning of the file. The value of <code>dot</code> is updated to the first byte of this block.</p> <p><code>find dir [-name n] [-inum i]</code> find files by name or i-number. <code>find</code> recursively searches directory <i>dir</i> and below for filenames whose i-number matches <i>i</i> or whose name matches pattern <i>n</i>. Note that only one of the two options (-name or -inum) may be used at one time. Also, the -print is not needed or accepted.</p> <p><code>fill=p</code> fill an area of disk with pattern <i>p</i>. The area of disk is delimited by <code>dot</code> and <code>count</code>.</p> <p><code>fragment</code> convert the value of <code>dot</code> to a fragment address. The only difference between the <code>fragment</code> command and the <code>block</code> command is the amount that is able to be displayed.</p> <p><code>inode</code> convert the value of <code>dot</code> to an inode address. If successful, the current value of <code>inode</code> will be updated as well as the value of <code>dot</code>. As a convenient shorthand, if 'inode' appears at the beginning of the line, the value of <code>dot</code> is set to the current <code>inode</code> and that <code>inode</code> is displayed in inode format.</p>
-----------------	---

log_chk	run through the valid log entries without printing any information and verify the layout.
log_delta	count the number of deltas into the log, using the value of dot as an offset into the log. No checking is done to make sure that offset is within the head/tail offsets.
log_head	display the header information about the file system logging. This shows the block allocation for the log and the data structures on the disk.
log_otodb	return the physical disk block number, using the value of dot as an offset into the log.
log_show	display all deltas between the beginning of the log (BOL) and the end of the log (EOL).
ls	[-R] [-l] <i>pat1 pat2</i> . . . list directories or files. If no file is specified, the current directory is assumed. Either or both of the options may be used (but, if used, <i>must</i> be specified before the filename specifiers). Also, as stated above, wild card characters are available and multiple arguments may be given. The long listing shows only the i-number and the name; use the <code>inode</code> command with '?i' to get more information.
override	toggle the value of override. Some error conditions may be overridden if override is toggled on.
prompt <i>p</i>	change the fsdb prompt to <i>p</i> . <i>p</i> must be surrounded by ("")s.
pwd	display the current working directory.
quit	quit fsdb.
sb	the value of <i>dot</i> is taken as a cylinder group number and then converted to the address of the superblock in that cylinder group. As a shorthand, ':sb' at the beginning of a line will set the value of <i>dot</i> to the superblock and display it in superblock format.
shadow	if the current inode is a shadow inode, then the value of <i>dot</i> is set to the beginning of the shadow inode data.
!	escape to shell

Inode Commands

In addition to the above commands, there are several commands that deal with inode fields and operate directly on the current `inode` (they still require the ':'). They may be used to more easily display or change the particular fields. The value of *dot* is only used by the ':db' and ':ib' commands. Upon completion of the command, the value of *dot* is changed to point to that particular field. For example,

```
> :ln+=1
```

fsdb_ufs(1M)

would increment the link count of the current *inode* and set the value of *dot* to the address of the link count field.

at access time.

bs block size.

ct creation time.

db use the current value of *dot* as a direct block index, where direct blocks number from 0 - 11. In order to display the block itself, you need to 'pipe' this result into the *block* or *fragment* command. For example,

```
> 1:db:block,20/X
```

would get the contents of data block field 1 from the *inode* and convert it to a block address. 20 longs are then displayed in hexadecimal. See *FormattedOutput*.

gid group id.

ib use the current value of *dot* as an indirect block index where indirect blocks number from 0 - 2. This will only get the indirect block itself (the block containing the pointers to the actual blocks). Use the *file* command and start at block 12 to get to the actual blocks.

ln link count.

mt modification time.

md mode.

maj major device number.

min minor device number.

nm although listed here, this command actually operates on the directory name field. Once poised at the desired directory entry (using the *directory* command), this command will allow you to change or display the directory name. For example,

```
> 7:dir:nm="foo"
```

will get the 7th directory entry of the current *inode* and change its name to *foo*. Note that names cannot be made larger than the field is set up for. If an attempt is made, the string is truncated to fit and a warning message to this effect is displayed.

si shadow *inode*.

sz file size.

uid user id.

Formatted Output | There are two styles and many format types. The two styles are structured and unstructured. Structured output is used to display inodes, directories, superblocks and the like. Unstructured displays raw data. The following shows the different ways of displaying:

?

c	display as cylinder groups
i	display as inodes
d	display as directories
s	display as superblocks
S	display as shadow inode data

/

b	display as bytes
c	display as characters
o O	display as octal shorts or longs
d D	display as decimal shorts or longs
x X	display as hexadecimal shorts or longs

The format specifier immediately follows the '/' or '?' character. The values displayed by '/b' and all '?' formats are displayed in the current base. Also, type is appropriately updated upon completion.

EXAMPLES

```
> 2000+400%(20+20)=D
will display 2010 in decimal (use of fsdb as a calculator for complex arithmetic).

> 386:ino?i
display i-number 386 in an inode format. This now becomes the current inode.

> :ln=4
changes the link count for the current inode to 4.

> :ln+=1
increments the link count by 1.

> :ct=X
display the creation time as a hexadecimal long.

> :mt=t
display the modification time in time format.

> 0:file/c
displays, in ASCII, block zero of the file associated with the current inode.

> 2:ino,*?d
displays the first blocks worth of directory entries for the root inode of this file system. It will stop prematurely if the EOF is reached.
```

fsdb_ufs(1M)

```
> 5:dir:inode; 0:file,*/c
    changes the current inode to that associated with the 5th directory entry (numbered
    from zero) of the current inode. The first logical block of the file is then displayed
    in ASCII.

> :sb
    displays the superblock of this file system.

> 1:cg?c
    displays cylinder group information and summary for cylinder group 1.

> 2:inode; 7:dir=3
    changes the i-number for the seventh directory slot in the root directory to 3.

> 2:db:block,*?d
    displays the third block of the current inode as directory entries.

> 7:dir:nm="name"
    changes the name field in the directory slot to name.

> 3c3:fragment,20:fill=0x20
    get fragment 3c3 and fill 20 type elements with 0x20.

> 2050=0xfffff
    set the contents of address 2050 to 0xffffffff. 0xffffffff may be truncated
    depending on the current type.

> 1c92434="this is some text"
    will place the ASCII for the string at 1c92434.

> 2:ino:si:ino;0:shadow,*?S
    displays all of the shadow inode data in the shadow inode associated with the root
    inode of this file system.
```

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES SunOS 5.8 Reference Manual WARNINGS

The following are added to handle shadow inodes: the *shadow* command, the *si* inode command, and the *S* format type.

clri(1M), *fsck_ufs*(1M), *dir_ufs*(4), *fs_ufs*(4), *attributes*(5)

Since *fsdb* reads the disk raw, extreme caution is advised in determining its availability of *fsdb* on the system. Suggested permissions are 600 and owned by bin.

NOTES

The old command line syntax for clearing i-nodes using the ufs-specific '*-z i-number*' option is still supported by the new debugger, though it is obsolete and will be removed in a future release. Use of this flag will result in correct operation, but

`fsdb_ufs(1M)`

an error message will be printed warning of the impending obsolescence of this option to the command. The equivalent functionality is available using the more flexible `clri(1M)` command.

ftpd(1M)

NAME	in.ftpd, ftpd – File transfer protocol server
SYNOPSIS	in.ftpd [-dl] [-t <i>timeout</i>]
DESCRIPTION	in.ftpd is the Internet File Transfer Protocol (FTP) server process. The server is invoked by the Internet daemon inetd(1M) each time a connection to the FTP service (see services(4)) is made.
OPTIONS	<p>-d Debugging information is logged to the system log daemon syslogd(1).</p> <p>-l Each FTP session is logged to the system log daemon syslogd(1).</p> <p>-t <i>timeout</i> Set the inactivity timeout period to <i>timeout</i> seconds. The FTP server will timeout an inactive session after 15 minutes.</p>
Requests	<p>The FTP server currently supports the following FTP requests; case is not distinguished.</p> <p>ABOR abort previous command</p> <p>ACCT specify account (ignored)</p> <p>ALLO allocate storage (vacuously)</p> <p>APPE append to a file</p> <p>CDUP change to parent of current working directory</p> <p>CWD change working directory</p> <p>DELE delete a file</p> <p>HELP give help information</p> <p>LIST give list files in a directory (ls -lg)</p> <p>MKD make a directory</p> <p>MODE specify data transfer <i>mode</i></p> <p>NLST give name list of files in directory (ls)</p> <p>NOOP do nothing</p> <p>PASS specify password</p> <p>PASV prepare for server-to-server transfer</p> <p>EPSV extended passive command request</p> <p>LPSV long passive command request</p> <p>PORT specify data connection port</p> <p>EPRT specify extended address for the transport connection</p> <p>LPRT specify “long” address for the transport connection</p>

PWD	print the current working directory
QUIT	terminate session
RETR	retrieve a file
RMD	remove a directory
RNFR	specify rename-from file name
RNTO	specify rename-to file name
STOR	store a file
STOU	store a file with a unique name
STRU	specify data transfer <i>structure</i>
TYPE	specify data transfer <i>type</i>
USER	specify user name
XCUP	change to parent of current working directory
XCWD	change working directory
XMKD	make a directory
XPWD	print the current working directory
XRMD	remove a directory

The remaining FTP requests specified in *RFC 959* are recognized, but not implemented.

The FTP server will abort an active file transfer only when the ABOR command is preceded by a Telnet “Interrupt Process” (IP) signal and a Telnet “Synch” signal in the command Telnet stream, as described in *RFC 959*. `in.ftpd` interprets file names according to the “globbing” conventions used by `sh(1)`. This allows users to utilize the metacharacters: `* ? [] { } ~`. `in.ftpd`’s `umask` (which it uses to create files during PUT operations) may be adjusted by adding the line

```
UMASK=nnn
```

```
to /etc/default/ftpd.
```

The banner returned by `in.ftpd` in the parenthetical portion of its greeting is configurable. The default is equivalent to “`uname -sr`” and will be used if no banner is set in `/etc/default/ftpd`. To set the banner, add a line of the form

```
BANNER="..."
```

```
to /etc/default/ftpd. Nonempty banner strings are fed to shells for evaluation.
```

The default banner may also be obtained by

ftpd(1M)

```
BANNER="'uname -s' 'uname -r'"
```

and no banner will be printed if `/etc/default/ftpd` contains

```
BANNER="
```

`in.ftpd` authenticates users according to five rules.

First, the user name must be in the password data base, `/etc/passwd`, and have a password that is not NULL. A password must always be provided by the client before any file operations may be performed. The PAM framework (see `SECURITY` below) is used to verify that the correct password was entered.

Second, if the user name appears in the file `/etc/ftpusers`, ftp access is denied. The default list of users in `/etc/ftpusers` includes all of the accounts in `passwd(4)`. See `ftpusers(4)`.

Third, ftp access is denied if the user's shell is not listed in `getusershell(3C)`

Fourth, if the user name is "anonymous" or "ftp", an entry for the user name ftp must be present in the password and shadow files. The user is then allowed to log in by specifying any password — by convention this is given as the user's e-mail address (such as `user@host.Sun.COM`). Do not specify a valid shell in the password entry of the ftp user, and do not give it a valid password (use NP in the encrypted password field of the shadow file).

Fifth, access is denied unless a user has the remote login authorization. If the `/etc/nologin` file exists, access is denied.

For anonymous ftp users, `in.ftpd` takes special measures to restrict the client's access privileges. The server performs a `chroot(2)` command to the home directory of the "ftp" user. In order that system security is not breached, it is recommended that the "ftp" subtree be constructed with care; the following rules are suggested.

`~ftp`

Make the home directory owned by root and unwritable by anyone.

`~ftp/bin`

Make this directory owned by root and unwritable by anyone. Make this a symbolic link to `~ftp/usr/bin`. The program `ls(1)` must be present to support the list commands. This program should have mode 111.

`~ftp/usr/lib`

Make this directory owned by root and unwritable by anyone. Copy the following shared libraries from `/usr/lib` into this directory:

```
ld.so.1*  
libc.so.1*  
libdl.so.1*  
libmp.so.2*  
libnsl.so.1*  
libsocket.so.1*
```

```
nss_compat.so.1*
nss_dns.so.1*
nss_files.so.1*
nss_nis.so.1*
nss_nisplus.so.1*
nss_xfn.so.1*
straddr.so*
straddr.so.2*

~ftp/etc
    Make this directory owned by root and unwritable by anyone. Copies of the files
    passwd(4), group(4), and netconfig(4) must be present for the ls(1) command
    to work properly. These files should be mode 444.

~ftp/pub
    Make this directory mode 755 and owned by root. Users should then place files
    which are to be accessible via the anonymous account in this directory.

~ftp/dev
    Make this directory owned by root and unwritable by anyone. First perform ls
    -lL on the device files listed below to determine their major and minor numbers,
    then use mknod to create them in this directory.

    /dev/zero
    /dev/tcp
    /dev/udp
    /dev/ticotsord

    Set the read and write mode on these nodes to 666 so that passive ftp will not fail
    with "permission denied" errors.

~ftp/usr/share/lib/zoneinfo
    Make this directory mode 555 and owned by root. Copy its contents from
    /usr/share/lib/zoneinfo. This enables ls -l to display time and date stamps
    correctly.
```

SECURITY in.ftpd uses pam(3PAM) for authentication, account management, and session management. The PAM configuration policy, listed through /etc/pam.conf, specifies the module to be used for in.ftpd. Here is a partial pam.conf file with entries for the in.ftpd command using the UNIX authentication, account management, and session management module.

ftp	auth	required	/usr/lib/security/pam_unix.so.1
ftp	account	required	/usr/lib/security/pam_unix.so.1
ftp	session	required	/usr/lib/security/pam_unix.so.1

ftpd(1M)

If there are no entries for the ftp service, then the entries for the "other" service will be used. Unlike login, passwd, and other commands, the ftp protocol will only support a single password. Using multiple modules will prevent `in.ftpd` from working properly.

USAGE The `in.ftpd` command is IPv6-enabled. See `ip6(7P)`.

EXAMPLES **EXAMPLE 1** Setting Up An Anonymous Ftp

To set up anonymous ftp, add the following entry to the `/etc/passwd` file. In this example, `/export/ftp` was chosen to be the anonymous ftp area, and the shell is the non-existent file `/nosuchshell`. This prevents users from logging in as the ftp user.

```
ftp:x:30000:30000:Anonymous FTP:/export/ftp:/nosuchshell
```

Add the following entry to the `/etc/shadow` file:

```
ftp:NP:6445:.....:
```

The following shell script sets up the anonymous ftp area. It presumes that names are resolved using NIS.

```
#!/bin/sh
# script to setup anonymous ftp area
#

# verify you are root
/usr/bin/id | grep -w 'uid=0' >/dev/null 2>&1
if [ "$?" != "0" ]; then
    echo
    exit 1
fi

# handle the optional command line argument
case $# in

    # the default location for the anon ftp comes from the passwd
    # file
    0) ftphome=`getent passwd ftp | cut -d: -f6`
       ;;

    1) if [ "$1" = "start" ]; then
        ftphome=`getent passwd ftp | cut -d: -f6`
        else
            ftphome=$1
        fi
        ;;

    *) echo "Usage: $0 [anon-ftp-root]"
       exit 1
       ;;

esac

if [ -z "${ftphome}" ]; then
    echo "$0: ftphome must be non-null"
    exit 2
```

EXAMPLE 1 Setting Up An Anonymous Ftp (Continued)

```

fi

case ${ftphome} in
    /*) # ok
        ;;

    *) echo "$0: ftphome must be an absolute pathname"
       exit 1
       ;;
esac

# This script assumes that ftphome is neither / nor /usr so ...
if [ -z "${ftphome}" -o "${ftphome}" = "/" -o "${ftphome}" = "/usr" ]; then
    echo "$0: ftphome must be non-null and neither / or /usr"
    exit 2
fi

# If ftphome does not exist but parent does, create ftphome
if [ ! -d ${ftphome} ]; then
    # lack of -p below is intentional
    mkdir ${ftphome}
fi
chown root ${ftphome}
chmod 555 ${ftphome}

echo Setting up anonymous ftp area ${ftphome}

# Ensure that the /usr directory exists
if [ ! -d ${ftphome}/usr ]; then
    mkdir -p ${ftphome}/usr
fi
# Now set the ownership and modes to match the man page
chown root ${ftphome}/usr
chmod 555 ${ftphome}/usr

# Ensure that the /usr/bin directory exists
if [ ! -d ${ftphome}/usr/bin ]; then
    mkdir -p ${ftphome}/usr/bin
fi
# Now set the ownership and modes to match the man page
chown root ${ftphome}/usr/bin
chmod 555 ${ftphome}/usr/bin

# this may not be the right thing to do
# but we need the bin -> usr/bin link
rm -f ${ftphome}/bin
ln -s usr/bin ${ftphome}/bin

# Ensure that the /usr/lib and /etc directories exist
if [ ! -d ${ftphome}/usr/lib ]; then
    mkdir -p ${ftphome}/usr/lib
fi
chown root ${ftphome}/usr/lib
chmod 555 ${ftphome}/usr/lib

```

EXAMPLE 1 Setting Up An Anonymous Ftp (Continued)

```

if [ ! -d ${ftphome}/usr/lib/security ]; then
    mkdir -p ${ftphome}/usr/lib/security
fi
chown root ${ftphome}/usr/lib/security
chmod 555 ${ftphome}/usr/lib/security

if [ ! -d ${ftphome}/etc ]; then
    mkdir -p ${ftphome}/etc
fi
chown root ${ftphome}/etc
chmod 555 ${ftphome}/etc

# a list of all the commands that should be copied to
# ${ftphome}/usr/bin
# /usr/bin/ls is needed at a minimum.
ftpcmd="
    /usr/bin/ls
"

# ${ftphome}/usr/lib needs to have all the libraries needed by the above
# commands, plus the runtime linker, and some name service libraries
# to resolve names. We just take all of them here.

ftplib="`ldd $ftpcmd | nawk ' $3 ~ /lib/ { print $3 }' | sort | uniq`"
ftplib="$ftplib /usr/lib/nss_* /usr/lib/straddr* /usr/lib/libmp.so*"
ftplib="$ftplib /usr/lib/libnsl.so.1 /usr/lib/libsocket.so.1 \\\
/usr/lib/ld.so.1"
ftplib="`echo $ftplib | tr ' ' '\`
' | sort | uniq`"

cp ${ftplib} ${ftphome}/usr/lib
chmod 555 ${ftphome}/usr/lib/*

cp /usr/lib/security/* ${ftphome}/usr/lib/security
chmod 555 ${ftphome}/usr/lib/security/*

cp ${ftpcmd} ${ftphome}/usr/bin
chmod 111 ${ftphome}/usr/bin/*

# you also might want to have separate minimal versions of passwd
# and group
cp /etc/passwd /etc/group /etc/netconfig /etc/pam.conf ${ftphome}/etc
chmod 444 ${ftphome}/etc/*

# need /etc/default/init for timezone to be correct
if [ ! -d ${ftphome}/etc/default ]; then
    mkdir ${ftphome}/etc/default
fi
chown root ${ftphome}/etc/default
chmod 555 ${ftphome}/etc/default
cp /etc/default/init ${ftphome}/etc/default
chmod 444 ${ftphome}/etc/default/init

```

EXAMPLE 1 Setting Up An Anonymous Ftp (Continued)

```
# Copy timezone database
mkdir -p ${ftphome}/usr/share/lib/zoneinfo
(cd ${ftphome}/usr/share/lib/zoneinfo
  (cd /usr/share/lib/zoneinfo; find . -print |
    cpio -o) 2>/dev/null | cpio -imdu 2>/dev/null
  find . -print | xargs chmod 555
  find . -print | xargs chown root
)

# Ensure that the /dev directory exists
if [ ! -d ${ftphome}/dev ]; then
  mkdir -p ${ftphome}/dev
fi

# make device nodes. ticotsord and udp are necessary for
# 'ls' to resolve NIS names.

for device in zero tcp udp ticotsord ticlts
do
  line=`ls -lL /dev/${device} | sed -e 's/,/\/'`
  major=`echo $line | awk '{print $5}'`
  minor=`echo $line | awk '{print $6}'`
  rm -f ${ftphome}/dev/${device}
  mknod ${ftphome}/dev/${device} c ${major} ${minor}
done

chmod 666 ${ftphome}/dev/*

## Now set the ownership and modes
chown root ${ftphome}/dev
chmod 555 ${ftphome}/dev

# uncomment the below if you want a place for people to store
# things, but beware the security implications
#if [ ! -d ${ftphome}/pub ]; then
#  mkdir -p ${ftphome}/pub
#fi
#chown root ${ftphome}/pub
#chmod 1755 ${ftphome}/pub
```

After running this script, edit the files in ~ftp/etc to make sure all non-public information is removed.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

ftpd(1M)

FILES	/etc/default/ftpd
	/etc/ftpusers file listing users for whom ftp login privileges are disallowed.
SUMMARY OF TRUSTED SOLARIS CHANGES Trusted Solaris 8 4/01 Reference Manual SunOS 5.10 Reference Manual	<p>Login is not allowed unless the user has the remote login authorization. If the /etc/nologin file exists, the user is not allowed to log in.</p> <p>inetd(1M), chroot(2), getsockopt(3SOCKET), inetd.conf(4)</p> <p>ftp(1), ld.so.1(1), ls(1), sh(1), aset(1M), mknod(1M), syslogd(1M), pam(3PAM), ftpusers(4), group(4), netconfig(4), netrc(4), pam.conf(4), passwd(4), services(4), attributes(5), pam_unix(5)</p> <p>Allman, M., Ostermann, S., and Metz, C., <i>RFC 2428, FTP Extensions for IPv6 and NATs</i>, The Internet Society, 1998.</p> <p>Postel, Jon, and Joyce Reynolds, <i>RFC 959, File Transfer Protocol (FTP)</i>, Network Information Center, SRI International, Menlo Park, Calif., October 1985.</p> <p>Piscitello, D., <i>RFC 1639, FTP Operation Over Big Address Records (FOOBAR)</i>, Network Working Group, June 1994.</p>
DIAGNOSTICS	in.ftpd logs various errors to syslogd, with a facility code of daemon.
Info Severity	<p>These messages are logged only if the -l flag is specified.</p> <p>FTPD: connection from <i>host</i> at <i>time</i> A connection was made to ftpd from the host <i>host</i> at the date and time <i>time</i>.</p> <p>FTPD: User <i>user</i> timed out after <i>timeout</i> seconds at <i>time</i> The user <i>user</i> was logged out because they had not entered any commands after <i>timeout</i> seconds; the logout occurred at the date and time <i>time</i>.</p>
Debug Severity	<p>These messages are logged only if the -d flag is specified.</p> <p>FTPD: command: <i>command</i> A command line containing <i>command</i> was read from the FTP client.</p> <p>lost connection The FTP client dropped the connection.</p> <p><— <i>replycode</i> <— <i>replycode</i>— A reply was sent to the FTP client with the reply code <i>replycode</i>. The next message logged will include the message associated with the reply. If a – follows the reply code, the reply is continued on later lines.</p>
NOTES	The anonymous ftp account is inherently dangerous and should be avoided when possible.

The name service caching daemon `/usr/sbin/nscd` may interfere with some of the functionality of anonymous ftp. The *sublogin* feature does not work unless caching for `passwd` is disabled in `/etc/nscd.conf`.

The server must run as the superuser to create sockets with privileged port numbers. It maintains an effective user ID of the logged in user, reverting to the superuser only when binding addresses to sockets. The possible security holes have been extensively scrutinized, but may be incomplete.

The file `/etc/ftpusers`, which is now included as part of Solaris, contains a list of users who cannot access the system; the default list of users in `/etc/ftpusers` includes all of the accounts in `passwd(4)`. See `ftpusers(4)`.

fuser(1M)

NAME	fuser – Identify processes using a file or file structure
SYNOPSIS	/usr/sbin/fuser [- [c f]ku] <i>files</i> [[- [c f]ku] <i>files</i>] ...
DESCRIPTION	<p>fuser displays the process IDs of the processes that are using the <i>files</i> specified as arguments.</p> <p>Each process ID is followed by a letter code. These letter codes are interpreted as follows: if the process is using the file as</p> <ul style="list-style-type: none"> c Indicates that the process is using the file as its current directory. m Indicates that the process is using a file mapped with <code>mmap()</code>. See <code>mmap(2)</code> for details. o Indicates that the process is using the file as an open file. r Indicates that the process is using the file as its root directory. t Indicates that the process is using the file as its text file. y Indicates that the process is using the file as its controlling terminal. <p>For block special devices with mounted file systems, all processes using any file on that device are listed. For all types of files (text files, executables, directories, devices, and so forth), only the processes using that file are reported.</p> <p>If more than one group of files are specified, the options may be respecified for each additional group of files. A lone dash cancels the options currently in force.</p> <p>The process IDs are printed as a single line on the standard output, separated by spaces and terminated with a single new line. All other output is written on standard error.</p> <p>To succeed, this command requires the <code>sys_mount</code> privilege. Any user with permission to read <code>/dev/kmem</code> and <code>/dev/mem</code> can use fuser.</p>
OPTIONS	<p>The following options are supported:</p> <ul style="list-style-type: none"> -c Reports on files that are mount points for file systems, and any files within that mounted file system. -f Print a report for the named file, not for files within a mounted file system. -k Sends the <code>SIGKILL</code> signal to each process. Since this option spawns kills for each process, the kill messages may not show up immediately. (See <code>kill(2)</code>). The <code>-k</code> option requires the <code>sys_mount</code> and <code>proc_owner</code> privileges if it is to terminate another user's process. -u Displays the user login name in parentheses following the process ID.
ENVIRONMENT VARIABLES	See <code>environ(5)</code> for descriptions of the following environment variables that affect the execution of fuser: <code>LANG</code> , <code>LC_ALL</code> , <code>LC_CTYPE</code> , <code>LC_MESSAGES</code> , and <code>NLSPATH</code> .

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**
Trusted Solaris 8
4/01 Reference
Manual
NOTES

To succeed, this command requires the `sys_mount` privilege. With the `-k` option, the command also needs the `proc_owner` privilege to terminate another user's process.

mount(1M), kill(2)

ps(1), signal(3C), attributes(5)

Because `fuser` works with a snapshot of the system image, it may miss processes that begin using a file while `fuser` is running. Also, processes reported as using a file may have stopped using it while `fuser` was running. These factors should discourage the use of the `-k` option.

getfsattr(1M)

NAME	getfsattr – display file system security attributes
SYNOPSIS	getfsattr [-l -L -m -p -P -s -S] [-F <i>fstype</i>] [-o <i>option</i>] [<i>pathname</i> ...]
DESCRIPTION	getfsattr displays the specified security attributes of the file system on which <i>pathname</i> resides. If no option is specified, all the file system security attributes are displayed. If no <i>pathname</i> is given, the attributes of all mounted file systems are displayed.
OPTIONS	<ul style="list-style-type: none"> -l Display the file system sensitivity level range in short form. -L Display the file system sensitivity level range in long form. -m Display the file system MLD prefix. -p Display the file system allowed privilege set. -P Display the file system forced privilege set. -s Display the file system sensitivity label in short form. -S Display the file system sensitivity label in long form. -F <i>fstype</i> Restrict the output to file systems of type <i>fstype</i>. -o <i>option</i> Specifies filesystem specific options. See the individual file system variants of getfsattr, such as getfsattr_ufs, for details.
EXIT STATUS	<p>getfsattr exits with one of the following values:</p> <ul style="list-style-type: none"> 0 Success 1 Usage error 2 Failure, error message is the system error number from getfsattr(2).
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWtsu

Trusted Solaris 8
4/01 Reference
Manual
getfsattr_ufs(1M), getfsattr(2), vfstab_adjunct(4)
attributes(5)

NAME	getfsattr_ufs – Display file system security attributes					
SYNOPSIS	getfsattr -F ufs [generic_options] -o {b c } { [mount_point] { device_name] ... }					
DESCRIPTION	getfsattr displays the security attributes specified by generic_options for the file system mounted on mount_point or on the device device_name.					
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:					
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWtsu</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE					
Availability	SUNWtsu					
OPTIONS	generic_options	Options supported by the generic getfsattr command. See getfsattr(1M) for a description of these options.				
	-o	Select the device type to read. A b requests use of the block device, while a c requests use of the character device.				
DIAGNOSTICS	getfsattr exits with one of the following values:					
	0	Success				
	1	Usage error				
	2	Unable to access the specified file system				
Trusted Solaris 8 4/01 Reference Manual	getfsattr(1M)					
SunOS 5.6 Reference Manual	attributes(5)					

halt(1M)

NAME	halt, poweroff – Stop the processor				
SYNOPSIS	<pre>/usr/sbin/halt [-dlnqy] /usr/sbin/poweroff [-dlnqy]</pre>				
DESCRIPTION	<p>halt and poweroff write out any pending information to the disks and then stop the processor. poweroff will have the machine remove power, if possible.</p> <p>halt and poweroff normally log the system shutdown to the system log daemon, syslogd(1M), and place a shutdown record in the login accounting file /var/adm/wtmp. These actions are inhibited if the -n or -q options are present.</p>				
OPTIONS	<p>The following options are supported:</p> <ul style="list-style-type: none">-d Force a system crash dump before rebooting. See dumpadm(1M) for information on configuring system crash dumps.-l Suppress sending a message about who executed halt to the system log daemon, syslogd(1M), about who executed halt-n Prevent the sync(4) before stopping.-q Quick halt. No graceful shutdown is attempted.-y Halt the system, even from a dialup terminal.				
FILES	/var/adm/wtmp History of user access and administration information				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWcsu</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>This command requires the sys_boot privilege and an effective uid of 0 in order to run.</p>				
Trusted Solaris 8 4/01 Reference Manual	init(1M), reboot(1M)				
Solaris 9 4/01 Reference Manual	dumpadm(1M), shutdown(1M), sync(1M), syslogd(1M), attributes(5)				
NOTES	<p>Unlike shutdown(1M) and init(1M), halt does not execute the rc0 scripts.</p> <p>The poweroff utility is equivalent to init 5.</p>				

NAME	hextoalabel – convert a hexadecimal label to its character-coded equivalent				
SYNOPSIS	<pre> /usr/sbin/hextoalabel [hexadecimal_CMW_label] /usr/sbin/hextoalabel -c [hexadecimal_clearance] /usr/sbin/hextoalabel -s [hexadecimal_sensitivity_label] </pre>				
DESCRIPTION	hextoalabel converts a hexadecimal label of the type specified into its standard formatted character-coded equivalent and writes the result to the standard output file. If no hexadecimal label is specified, one is read from standard input.				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWtsu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWtsu				
OPTIONS	<p>-c Identifies the hexadecimal label as a clearance.</p> <p>-s Identifies the hexadecimal label as a sensitivity label.</p>				
EXIT STATUS	<p>The following exit values are returned:</p> <p>0 On success.</p> <p>1 On failure, and writes diagnostics to the standard error file.</p>				
FILES	<p>/etc/security/tsol/label_encodings</p> <p>The label encodings file contains the classification names, words, constraints, and values for the defined labels of this system.</p>				
Trusted Solaris 8 4/01 Reference Manual	<p>label_encodings(4)</p> <p><i>Trusted Solaris administrator's document set</i></p>				
SunOS 5.8 Reference Manual DIAGNOSTICS	<p>attributes(5)</p> <p>label translation unavailable or <i>hexadecimal_label</i> not translatable by this process</p> <p>The label services are unavailable at this time for one of these reasons: either the label daemon is not running, or the label_encodings file is incorrect or unavailable, or this process is not allowed to translate <i>hexadecimal_label</i>. The sys_trans_label privilege may be used to override this last restriction.</p> <p>unable to translate <i>hexadecimal_label</i> as type specified <i>hexadecimal_label</i> does not match the hexadecimal format for the specified type.</p>				

ifconfig(1M)

NAME	ifconfig – configure network interface parameters
SYNOPSIS	<pre> /sbin/ifconfig interface [address_family] [address [/prefix_length] [dest_address]] [addif address [/prefix_length]] [removeif address [/prefix_length]] [arp -arp] [auth_algs authentication algorithm] [encr_algs encryption algorithm] [encr_auth_algs authentication algorithm] [auto-revarp] [broadcast address] [deprecated -deprecated] [destination dest_address] [[failover] [-failover]] [group [[name] ""]] [index {if_index}] [metric n] [modlist] [modinsert mod_name@pos] [modremove mod_name@pos] [mtu n] [netmask mask] [plumb] [unplumb] [private -private] [nud -nud] [set [address] [/netmask]] [[standby] [-standby]] [subnet subnet_address] [tdst tunnel_dest_address] [tsrc tunnel_src_address] [trailers -trailers] [up] [down] [xmit -xmit] /usr/sbin/ifconfig interface [address_family] [address [/prefix_length] [dest_address]] [addif address [/prefix_length]] [removeif address [/prefix_length]] [arp -arp] [auth_algs authentication algorithm] [encr_algs encryption algorithm] [encr_auth_algs authentication algorithm] [auto-revarp] [broadcast address] [deprecated -deprecated] [destination dest_address] [[failover] [-failover]] [group [[name] ""]] [index {if_index}] [metric n] [modlist] [modinsert mod_name@pos] [modremove mod_name@pos] [mtu n] [netmask mask] [plumb] [unplumb] [private -private] [nud -nud] [set [address] [/netmask]] [[standby] [-standby]] [subnet subnet_address] [tdst tunnel_dest_address] [tsrc tunnel_src_address] [trailers -trailers] [up] [down] [xmit -xmit] /sbin/ifconfig interface {auto-dhcp dhcp} [primary] [wait seconds] drop extend inform ping release start status /usr/sbin/ifconfig interface {auto-dhcp dhcp} [primary] [wait seconds] drop extend inform ping release start status </pre>
DESCRIPTION	<p>The command <code>ifconfig</code> is used to assign an address to a network interface and to configure network interface parameters. The <code>ifconfig</code> command must be used at boot time to define the network address of each interface present on a machine; it may also be used at a later time to redefine an interface's address or other operating parameters. If no option is specified, <code>ifconfig</code> displays the current configuration for a network interface. If an address family is specified, <code>ifconfig</code> reports only the details specific to that address family. <code>ifconfig</code> needs the <code>sys_net_config</code> privilege in order to modify the configuration of a network interface. Options appearing within braces (<code>{ }</code>) indicate that one of the options must be specified.</p> <p>The two versions of <code>ifconfig</code>, <code>/sbin/ifconfig</code> and <code>/usr/sbin/ifconfig</code>, behave differently with respect to name services. The order in which names are looked</p>

DHCP Configuration

up by `/sbin/ifconfig` when the system is booting is fixed and cannot be changed. In contrast, changing `/etc/nsswitch.conf` may affect the behavior of `/usr/sbin/ifconfig`. The system administrator may configure the source and lookup order in the tables by means of the name service switch. See `nsswitch.conf(4)` for more information.

The third and fourth forms of this command are used to control the Dynamic Host Configuration Protocol (“DHCP”) configuring of the interface. DHCP is only available on interfaces for which the address family is `inet`. In this mode, `ifconfig` is used to control operation of `dhcagent(1M)`, the DHCP client daemon. Once an interface is placed under DHCP control by using the `start` operand, `ifconfig` should not, in normal operation, be used to modify the address or characteristics of the interface. If the address of an interface under DHCP is changed, `dhcagent` will remove the interface from its control.

OPTIONS

Options that need to open network devices readable only by root and protected at `ADMIN_HIGH` (`ether`, `auto-revarp`, and `plumb`) are intended to be invoked at `ADMIN_HIGH` with effective user ID 0. These restrictions may be overridden by the `file_dac_read`, `file_dac_write`, and `file_mac_read` privileges.

The following options are supported:

`addif address`

Create the next unused logical interface on the specified physical interface.

`arp`

Enable the use of the Address Resolution Protocol (“ARP”) in mapping between network level addresses and link level addresses (default). This is currently implemented for mapping between IPv4 addresses and 10Mb/s Ethernet addresses.

`-arp`

Disable the use of the ARP.

`auth_algs authentication algorithm`

For a tunnel, enable IPsec AH with the authentication algorithm specified. The algorithm can be either a number or an algorithm name, including *any* to express no preference in algorithm. All IPsec tunnel properties must be specified on the same command line. To disable tunnel security, specify an `auth_alg` of *none*.

`auto-dhcp`

Use DHCP to automatically acquire an address for this interface. This option has a completely equivalent alias called `dhcp`.

`primary`

Defines the interface as the `primary`. The interface is defined as the preferred one for the delivery of client-wide configuration data. Only one interface can be the `primary` at any given time. If another interface is subsequently selected as the `primary`, it replaces the previous one. Nominating an interface as the `primary` one will not have much significance

ifconfig(1M)

	once the client work station has booted, as many applications will already have started and been configured with data read from the previous primary interface.
<code>wait <i>seconds</i></code>	The <code>ifconfig</code> command will wait until the operation either completes or for the interval specified, whichever is the sooner. If no wait interval is given, and the operation is one that cannot complete immediately, <code>ifconfig</code> will wait 30 seconds for the requested operation to complete. The symbolic value <code>forever</code> may be used as well, with obvious meaning.
<code>drop</code>	Remove the specified interface from DHCP control. Additionally, set the IP address to zero and mark the interface as "down".
<code>extend</code>	Attempt to extend the lease on the interface's IPv4 address. This is not required, as the agent will automatically extend the lease well before it expires.
<code>inform</code>	Obtain network configuration parameters from DHCP without obtaining a lease on an IP address. This is useful in situations where an IP address is obtained through mechanisms other than DHCP.
<code>ping</code>	Check whether the interface given is under DHCP control, which means that the interface is managed by the DHCP agent and is working properly. An exit status of 0 means success. This subcommand has no meaning when the named interface represents more than one interface.
<code>release</code>	Relinquish the IPv4 address on the interface, and mark the interface as "down."
<code>start</code>	Start DHCP on the interface.
<code>status</code>	Display the DHCP configuration status of the interface.
<code>auto-revarp</code>	Use the Reverse Address Resolution Protocol ("RARP") to automatically acquire an address for this interface.
<code>broadcast <i>address</i></code>	For IPv4 only. Specify the address to use to represent broadcasts to the network. The default broadcast address is the address with a host part of all 1's. A "+" (plus sign) given for the broadcast value causes the broadcast address to be reset to a default appropriate for the (possibly new) address and netmask. The arguments of <code>ifconfig</code> are interpreted left to right. Therefore

```
example% ifconfig -a netmask + broadcast +
and
```

```
example% ifconfig -a broadcast + netmask +
may result in different values being assigned for the broadcast addresses of the
interfaces.
```

deprecated

Marks the address as a deprecated address. Addresses marked as deprecated will not be used as source address for outbound packets unless either there are no other addresses available on this interface or the application has bound to this address explicitly. The status display shows DEPRECATED as part of flags.

-deprecated

Marks the address as not deprecated.

destination *dest_address*

Set the destination address for a point-to point interface.

dhcp

This option is an alias for option `auto-dhcp`.

down

Mark an interface "down". When an interface is marked "down", the system does not attempt to transmit messages through that interface. If possible, the interface is reset to disable reception as well. This action does not automatically disable routes using the interface.

encr_auth_algs *authentication algorithm*

For a tunnel, enable IPsec ESP with the authentication algorithm specified. It can be either a number or an algorithm name, including any or none, to indicate no algorithm preference. If an ESP encryption algorithm is specified but the authentication algorithm is not, the default value for the ESP authentication algorithm will be any.

encr_algs *encryption algorithm*

For a tunnel, enable IPsec ESP with the encryption algorithm specified. It can be either a number or an algorithm name. Note that all IPsec tunnel properties must be specified on the same command line. To disable tunnel security, specify the value of `encr_alg` as none. If an ESP authentication algorithm is specified, but the encryption algorithm is not, the default value for the ESP encryption will be null.

-failover

Mark the address as a non-failover address. Addresses marked this way will not failover when the interface fails. Status display shows "NOFAILOVER" as part of flags.

failover

Mark the address as a failover address. This address will failover when the interface fails. Status display does not show "NOFAILOVER" as part of flags.

ifconfig(1M)

group [*name* | ""]

Insert the interface in the multipathing group specified by *name*. To delete an interface from a group, use a null string "". When invoked on the logical interface with id zero, the status display shows the group name.

index *n*

Change the interface index for the interface. The value of *n* must be an interface index (*if_index*) that is not used on another interface. *if_index* will be a non-zero positive number that uniquely identifies the network interface on the system.

metric *n*

Set the routing metric of the interface to *n*; if no value is specified, the default is 0. The routing metric is used by the routing protocol. Higher metrics have the effect of making a route less favorable; metrics are counted as addition hops to the destination network or host.

modinsert *mod_name@pos*

Insert a module with name *mod_name* to the stream of the device at position *pos*. The position is relative to the stream head. Position 0 means directly under stream head.

Based upon the example in the **modlist** option, use the following command to insert a module with name *ipqos* under the *ip* module and above the *firewall* module:

```
example% ifconfig hme0 modinsert ipqos@2
```

A subsequent listing of all the modules in the stream of the device follows:

```
example% ifconfig hme0 modlist
0 arp
1 ip
2 ipqos
3 firewall
4 hme
```

modlist

List all the modules in the stream of the device.

The following example lists all the modules in the stream of the device:

```
example% ifconfig hme0 modlist
0 arp
1 ip
2 firewall
4 hme
```

modremove *mod_name@pos*

Remove a module with name *mod_name* from the stream of the device at position *pos*. The position is relative to the stream head.

Based upon the example in the **modinsert** option, use the following command to remove the *firewall* module from the stream after inserting the *ipqos* module:

```
example% ifconfig hme0 modremove firewall@3
```

A subsequent listing of all the modules in the stream of the device follows:

```
example% ifconfig hme0 modlist
0 arp
1 ip
2 ipqos
3 hme
```

Note that the core IP stack modules, for example, `ip` and `tun` modules, cannot be removed.

`mtu n`

Set the maximum transmission unit of the interface to *n*. For many types of networks, the `mtu` has an upper limit, for example, 1500 for Ethernet.

`netmask mask`

For IPv4 only. Specify how much of the address to reserve for subdividing networks into subnetworks. The mask includes the network part of the local address and the subnet part, which is taken from the host field of the address. The mask contains 1's for the bit positions in the 32-bit address which are to be used for the network and subnet parts, and 0's for the host part. The mask should contain at least the standard network portion, and the subnet field should be contiguous with the network portion. The mask can be specified in one of four ways:

1. with a single hexadecimal number with a leading 0x,
2. with a dot-notation address,
3. with a "+" (plus sign) address, or
4. with a pseudo host name/pseudo network name found in the network database `networks(4)`.

If a "+" (plus sign) is given for the netmask value, the mask is looked up in the `netmasks(4)` database. This lookup finds the longest matching netmask in the database by starting with the interface's IPv4 address as the key and iteratively masking off more and more low order bits of the address. This iterative lookup ensures that the `netmasks(4)` database can be used to specify the netmasks when variable length subnetmasks are used within a network number.

If a pseudo host name/pseudo network name is supplied as the netmask value, netmask data may be located in the `hosts` or `networks` database. Names are looked up by first using `gethostbyname(3NSL)`. If not found there, the names are looked up in `getnetbyname(3SOCKET)`. These interfaces may in turn use `nsswitch.conf(4)` to determine what data store(s) to use to fetch the actual value.

For both `inet` and `inet6`, the same information conveyed by *mask* can be specified as a *prefix_length* attached to the *address* parameter.

`nud`

Enables the neighbor unreachability detection mechanism on a point-to-go interface.

ifconfig(1M)

`-nud`

Disables the neighbor unreachability detection mechanism on a point-to-go interface.

`plumb`

Open the device associated with the physical interface name and set up the streams needed for IP to use the device. When used with a logical interface name, this command is used to create a specific named logical interface. An interface must be separately plumbed for use by IPv4 and IPv6. The *address_family* parameter controls whether the `ifconfig` command applies to IPv4 or IPv6.

Before an interface has been plumbed, , the interface will not show up in the output of the `ifconfig -a` command.

`private`

Tells the `in.routed` routing daemon that the interface should not be advertised.

`-private`

Specify unadvertised interfaces.

`removeif address`

Remove the logical interface on the physical interface specified that matches the *address* specified.

`set`

Set the *address*, *prefix_length* or both, for an interface.

`standby`

Marks the physical interface as a standby interface. If the interface is marked `STANDBY` and is part of the multipathing group, the interface will not be selected to send out packets unless some other interface in the group has failed and the network access has been failed over to this standby interface.

The status display shows "`STANDBY, INACTIVE`" indicating that that the interface is a standby and is also inactive. `IFF_INACTIVE` will be cleared when some other interface belonging to the same multipathing group fails over to this interface. Once a failback happens, the status display will return to `INACTIVE`.

`-standby`

Turns off standby on this interface.

`subnet`

Set the subnet *address* for an interface.

`tdst tunnel_dest_address`

Set the destination address of a tunnel. The address should not be the same as the *dest_address* of the tunnel, because no packets leave the system over such a tunnel.

trailers

This flag previously caused a nonstandard encapsulation of `inet` packets on certain link levels. Drivers supplied with this release no longer use this flag. It is provided for compatibility, but is ignored.

-trailers

Disable the use of a "trailer" link level encapsulation.

tsrc *tunnel_src_address*

Set the source address of a tunnel. This is the source address on an outer encapsulating IP header. It must be an address of another interface already configured using `ifconfig`.

unplumb

Destroy any streams associated with this physical interface and close the associated device. When used with a logical interface name, the logical interface is removed from the system. After this command is executed, the device name will no longer appear in the output of `ifconfig -a`. An interface must be "down" before it can be unplumbed.

up

Mark an interface "up". This happens automatically when setting the first address on an interface. The `up` option enables an interface after an `ifconfig down`, which reinitializes the hardware.

xmit

Enable an interface to transmit packets. This is the default behavior when the interface is up.

-xmit

Disable transmission of packets on an interface. The interface will continue to receive packets.

OPERANDS The *interface* operand, as well as address parameters that affect it, are described below.

interface

A string of the form, *name physical-unit*, for example, `le0` or `ie1`; or of the form *name physical-unit:logical-unit*, for example, `le0:1`; or of the form `ip.tunN`, for tunnels.

If the interface name starts with a dash (-), it is interpreted as a set of options which specify a set of interfaces. In such a case, `-a` must be part of the options and any of the additional options below can be added in any order. If one of these interface names is given, the commands following it are applied to all of the interfaces that match.

-a Apply the commands to all interfaces in the system.

ifconfig(1M)

	<ul style="list-style-type: none"> -d Apply the commands to all "down" interfaces in the system. -D Apply the commands to all interfaces not under DHCP (Dynamic Host Configuration Protocol) control. -u Apply the commands to all "up" interfaces in the system. -4 Apply the commands to all IPv4 interfaces. -6 Apply the commands to all IPv6 interfaces.
<i>address_family</i>	<p>The address family is specified by the <i>address_family</i> parameter. The <i>ifconfig</i> command currently supports the following families: <i>ether</i>, <i>inet</i>, and <i>inet6</i>. If no address family is specified, the default is <i>inet</i>.</p>
<i>address</i>	<p>For the IPv4 family (<i>inet</i>), the <i>address</i> is either a host name present in the host name data base (see <i>hosts(4)</i>) or in the Network Information Service (NIS) map <i>hosts</i>, or an IPv4 address expressed in the Internet standard "dot notation".</p> <p>For the IPv6 family (<i>inet6</i>), the <i>address</i> is either a host name present in the host name data base (see <i>ipnodes(4)</i>) or in the Network Information Service (NIS) map <i>ipnode</i>, or an IPv6 address expressed in the Internet standard colon-separated hexadecimal format represented as <i>x::x::x::x::x::x</i> where <i>x</i> is a hexadecimal number between 0 and FFFF.</p> <p>For the <i>ether</i> address family, the address is an Ethernet address represented as <i>x::x::x::x</i> where <i>x</i> is a hexadecimal number between 0 and FF.</p> <p>Some, though not all, of the Ethernet interface cards have their own addresses. To use cards that do not have their own addresses, refer to section 3.2.3(4) of the IEEE 802.3 specification for a definition of the locally administered address space. The use of interface groups should be restricted to those cards with their own addresses (see INTERFACE GROUPS).</p>
<i>prefix_length</i>	<p>For the IPv4 and IPv6 families (<i>inet</i> and <i>inet6</i>), the <i>prefix_length</i> is a number between 0 and the number of bits in the address. For <i>inet</i>, the number of bits in the address is 32; for <i>inet6</i>, the number of bits in the</p>

	address is 128. The <i>prefix_length</i> denotes the number of leading set bits in the netmask.
<i>dest_address</i>	If the <i>dest_address</i> parameter is supplied in addition to the <i>address</i> parameter, it specifies the address of the correspondent on the other end of a point-to-point link.
<i>tunnel_dest_address</i>	An address that is or will be reachable through an interface other than the tunnel being configured. This tells the tunnel where to send the tunneled packets. This address must not be the same as the <i>tunnel_dest_address</i> being configured.
<i>tunnel_src_address</i>	As address that is attached to an already configured interface that has been configured "up" with <i>ifconfig</i> .

LOGICAL INTERFACES

Solaris TCP/IP allows multiple logical interfaces to be associated with a physical network interface. This allows a single machine to be assigned multiple IP addresses, even though it may have only one network interface. Physical network interfaces have names of the form *driver-name physical-unit-number*, while logical interfaces have names of the form *driver-name physical-unit-number:logical-unit-number*. A physical interface is configured into the system using the *plumb* command. For example:

example% *ifconfig le0 plumb* Once a physical interface has been "plumbed", logical interfaces associated with the physical interface can be configured by separate *plumb* or *addif* options to the *ifconfig* command.

example% *ifconfig le0:1 plumb* allocates a specific logical interface associated with the physical interface *le0*. The command

example% *ifconfig le0 addif 192.9.200.1/24 up* allocates the next available logical unit number on the *le0* physical interface and assigns an *address* and *prefix_length*.

A logical interface can be configured with parameters (*address*, *prefix_length*, and so on) different from the physical interface with which it is associated. Logical interfaces that are associated with the same physical interface can be given different parameters as well. Each logical interface must be associated with an existing and "up" physical interface. So, for example, the logical interface *le0:1* can only be configured after the physical interface *le0* has been plumbed.

To delete a logical interface, use the *unplumb* or *removeif* options. For example,

example% *ifconfig le0:1 down unplumb* Will delete the logical interface *le0:1*.

INTERFACE GROUPS

If an interface (logical or physical) shares an IP prefix with another interface, these interfaces are collected into an *interface group*. IP uses an interface group to rotate source address selection when the source address is unspecified, and in the case of

ifconfig(1M)

multiple physical interfaces in the same group, to scatter traffic across different IP addresses on a per-IP-destination basis. See `netstat(1M)` for per-IP-destination information.

This feature may be enabled by using `ndd(1M)`.

One can also use the `group` keyword to form a multipathing group. When multipathing groups are used, the functionality of the `interface group` is subsumed into the functionality of the multipathing group. A multipathing group provides failure detection and repair detection for the interfaces in the group. See `in.mpathd(1M)` and *System Administration Guide, Volume 3*.

The interface groups formed using `ndd(1M)` will be made obsolete in the future. Accordingly, it is advisable to use form multipathing groups using the `group` keyword.

CONFIGURING IPv6 INTERFACES

When an IPv6 physical interface is plumbed and configured “up” with `ifconfig`, it is automatically assigned an IPv6 link-local address for which the last 64 bits are calculated from the MAC address of the interface.

`ifconfig le0 inet6 plumb up` The following example shows that the link-local address has a prefix of `fe80::/10`.

```
example% ifconfig le0 inet6
le0: flags=2000841<UP,RUNNING,MULTICAST,IPv6>
      mtu 1500 index 2
      inet6 fe80::a00:20ff:fe8e:f3ad/10
```

If an advertising IPv6 router exists on the link advertising prefixes, then the newly plumbed IPv6 interface will autoconfigure logical interface(s) depending on the prefix advertisements. For example, for prefix advertisements `fec0:0:0:55::/64` and `3ff0:0:0:55::/64`, the autoconfigured interfaces will look like:

```
le0:1: flags=2080841<UP,RUNNING,MULTICAST,ADDRCONF,IPv6>
      mtu 1500 index 2
      inet6 fec0::55:a00:20ff:fe8e:f3ad/64
le0:2: flags=2080841<UP,RUNNING,MULTICAST,ADDRCONF,IPv6>
      mtu 1500 index 2
      inet6 3ff0::55:a00:20ff:fe8e:f3ad/64
```

Even if there are no prefix advertisements on the link, you can still assign site-local and global addresses manually, for example:

```
example% ifconfig le0 inet6 addif fec0::55:a00:20ff:fe8e:f3ad/64 up
example% ifconfig le0 inet6 addif 3ff0::55:a00:20ff:fe8e:f3ad/64 up
```

To configure boot-time defaults for the interface `le0`, place the following entries in the `/etc/hostname6.le0` file:

```
addif fec0::55:a00:20ff:fe8e:f3ad/64 up
addif 3ff0::55:a00:20ff:fe8e:f3ad/64 up
```

Link-local addresses are only used for on-link communication and are not visible to other subnets.

**Configuring
IPv6/IPv4 tunnels**

An IPv6 over IPv4 tunnel interface can send and receive IPv6 packets encapsulated in an IPv4 packet. Create tunnels at both ends pointing to each other. IPv6 over IPv4 tunnels require the tunnel source and tunnel destination IPv4 and IPv6 addresses. Solaris 8 supports both automatic and configured tunnels. For automatic tunnels, an IPv4-compatible IPv6 address is used. The following demonstrates auto-tunnel configuration:

```
example% ifconfig ip.atun0 inet6 plumb
example% ifconfig ip.atun0 inet6 tsrc <IPv4-address> \
    ::<IPv4 address>/96 up
```

where IPv4-address is the IPv4 address of the interface through which the tunnel traffic will flow, and IPv4-address, ::<IPv4-address>, is the corresponding IPv4-compatible IPv6 address.

The following is an example of a configured tunnel:

```
example% ifconfig ip.tun0 inet6 plumb tsrc <my-ipv4-address> \
    tdst <peer-ipv4-address> up
```

This creates a configured tunnel between my-ipv4-address and peer-ipv4-address with corresponding link-local addresses. For tunnels with global or site-local addresses, the logical tunnel interfaces need to be configured in the following form:

```
ifconfig ip.tun0 inet6 addif <my-v6-address> <peer-v6-address> up
```

For example,

```
example% ifconfig ip.tun0 inet6 plumb tsrc 109.146.85.57 \
    tdst 109.146.85.212 up
example% ifconfig ip.tun0 inet6 addif 2::45 2::46 up
```

To show all IPv6 interfaces that are up and configured:

```
example% ifconfig -au6
ip.tun0: flags=2200851<UP,POINTOPOINT,RUNNING,MULTICAST,NUD,IPv6>
    mtu 1480 index 3
    inet tunnel src 109.146.85.57    tunnel dst 109.146.85.212
    inet6 fe80::6d92:5539/10 --> fe80::6d92:55d4
ip.tun0:1: flags=2200851<UP,POINTOPOINT,RUNNING,MULTICAST,NUD,IPv6>
    mtu 1480 index 3
    inet6 2::45/128 --> 2::46
```

EXAMPLES**EXAMPLE 1** Using the ifconfig Command

If your workstation is not attached to an Ethernet, the le0 interface should be marked "down" as follows:

```
example% ifconfig le0 down
```

EXAMPLE 2 Printing Addressing Information

To print out the addressing information for each interface, use the following command:

```
example% ifconfig -a
```

EXAMPLE 2 Printing Addressing Information *(Continued)***EXAMPLE 3** Resetting the Broadcast Address

To reset each interface's broadcast address after the netmasks have been correctly set, use the next command:

```
example% ifconfig -a broadcast +
```

EXAMPLE 4 Changing the Ethernet Address

To change the Ethernet address for interface `le0`, use the following command:

```
example% ifconfig le0 ether aa:1:2:3:4:5
```

EXAMPLE 5 Configuring an IP-in-IP Tunnel

To configure an IP-in-IP tunnel, first plumb it with the following command:

```
example% ifconfig ip.tun0 plumb
```

Then configure it as a point-to-point interface, supplying the tunnel source and the tunnel destination:

```
example% ifconfig ip.tun0 myaddr mydestaddr tsrc another_myaddr \
        tdst a_dest_addr up
```

Tunnel security properties must be configured on one invocation of `ifconfig`:

```
example% ifconfig ip.tun0 encr_auth_algs md5 encr_algs 3des
```

EXAMPLE 6 Requesting a Service Without Algorithm Preference

To request a service without any algorithm preferences, specify any:

```
example% inconfig ip.tun0 encr_auth_algs any encr_algs any
```

EXAMPLE 7 Disabling All Security

To disable all security, specify any security service with none as the algorithm value:

```
example% ifconfig ip.tun0 auth_algs none
```

OR

```
example% ifconfig ip.tun0 encr_algs none
```

FILES `/etc/netmasks` netmask data

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

`/usr/sbin`

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

ifconfig(1M)

/sbin

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level for options modlist, modinsert, and modremove	Evolving

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsr
Stability Level for options modlist, modinsert, and modremove	Evolving

SUMMARY OF TRUSTED SOLARIS CHANGES

The ifconfig command needs the sys_net_config privilege to configure network interfaces; without privilege, ifconfig displays the status of network interfaces. The ether, auto-revarp, and plumb options need to open ADMIN_HIGH network devices readable only by root; these options are intended to be invoked at ADMIN_HIGH with an effective user ID of 0. Alternately, file_dac_read, file_dac_write, and file_mac_read privileges may be used to override these restrictions.

Trusted Solaris 8
4/01 Reference
Manual

in.routed(1M), ndd(1M), netstat(1M), nsswitch.conf(4)

dhcinfo(1), dhcpagent(1M), in.mpathd(1M), ethers(3SOCKET), gethostbyname(3NSL), getnetbyname(3SOCKET), hosts(4), netmasks(4), networks(4), attributes(5), arp(7P), ipsec(7P), ipsec(7P), tun(7M)

System Administration Guide, Volume 3

DIAGNOSTICS

ifconfig sends messages that indicate if:

- the specified interface does not exist
- the requested address is unknown
- the user is not privileged and tried to alter an interface's configuration

NOTES

It is recommended that the names broadcast, down, private, trailers, up, and the other possible option names not be selected when choosing host names. Choosing any one of these names as host names will cause bizarre problems that can be extremely difficult to diagnose.

inetd(1M)

NAME	inetd – Internet services daemon
SYNOPSIS	inetd [-d] [-s] [-t] [-r <i>count interval</i>] [<i>configuration-file</i>]
DESCRIPTION	<p>inetd is the server process for the Internet standard services. It is usually started up at system boot time. The <i>configuration-file</i> lists the services that inetd is to provide. If no <i>configuration-file</i> is given on the command line, inetd reads its configuration information from the file <code>/etc/inetd.conf</code>. See <code>inetd.conf(4)</code> for more information on the format of this file. inetd listens for service requests on the TCP or UDP ports associated with each of the service listed in the configuration file. When a request arrives, inetd executes the server program associated with the service.</p> <p>A service can be configured to be "wait" wait-status, in which case, inetd waits for the server process to exit before starting a second server process. RPC services can also be started by inetd.</p> <p>inetd provides a number of simple Internet services internally. These include <code>echo</code>, <code>discard</code>, <code>chargen</code> (character generator), <code>daytime</code> (human-readable time), and <code>time</code> (machine-readable time, in the form of the number of seconds since midnight, January 1, 1900).</p> <p>inetd rereads its <i>configuration-file</i> once when it is started and again whenever it receives a hangup signal, <code>SIGHUP</code>. New services can be activated, and existing services deleted or modified by editing the <i>configuration-file</i>, then sending inetd a <code>SIGHUP</code> signal.</p> <p>Then inetd reads the <i>configuration-file</i> and attempts to <code>bind()</code> to the service to start listening to it. That attempt may fail if another standalone server or "wait" wait-status server started by inetd is already listening for this service. inetd will defer implementing the newly read configuration for that service and will attempt periodically to start listening, after logging an error on console. The retry interval is currently 10 minutes.</p>
OPTIONS	<p>-d Runs inetd in the foreground and enables debugging output.</p> <p>-s Allows you to run inetd "standalone," outside the Service Access Facility (SAF). If the -s option is omitted, inetd will attempt to contact the service access controller (SAC) and will exit if SAC is not already running. See <code>sac(1M)</code>.</p> <p>-t Instructs inetd to trace the incoming connections for all of its TCP services. It does this by logging the client's IP address and TCP port number, along with the name of the service, using the <code>syslog(3C)</code> facility. UDP services can not be traced. When tracing is enabled, inetd uses the syslog facility code "daemon" and "notice" priority level.</p> <p>-r Allows inetd to detect and then suspend "broken" connectionless datagram services servers, for example, UDP, and RPC/CLTS. Without this</p>

detection, a buggy server that fails before consuming the service request will be continuously restarted and will tax system resources too much. The `-r` flag has the form:

`-r count interval` *count* and *interval* are decimal numbers that represent the maximum *count* of invocations per *interval* of seconds a service may be started before the service is considered “broken.”

Once considered “broken,” a server is suspended for ten minutes. After ten minutes, `inetd` again enables service, hoping the server operates correctly.

If the `-r` flag is not specified, `inetd` behaves as though `-r40 60` was specified.

OPERANDS *configuration-file* Lists the services `inetd` is to provide.

EXIT STATUS `inetd` does not return an exit status.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

`inetd` starts servers at the correct sensitivity label based upon the sensitivity label of the client request.

A number of new configuration options are defined in `inetd.conf(4)`. See that man page for more detail.

`inetd` registers RPC servers as multilevel servers with `rpcbind`.

If there is an entry for a server in the `inetd` profile and that entry specifies privileges, the server will inherit the specified privileges from `inetd`. To support this inheritance, `inetd` must have all privileges.

If there is an entry for a server in the `inetd` profile entry and that entry specifies minimum and maximum sensitivity labels, `inetd` will verify that the sensitivity label of the client is within the specified min/max range. If the label is not within the range, the server will not be executed.

Trusted Solaris 8 4/01 Reference Manual

`in.ftpd(1M)`, `in.rexecd(1M)`, `in.rshd(1M)`, `in.tftpd(1M)`, `inetd.conf(4)`

`sac(1M)`, `syslog(3C)`, `attributes(5)`

Postel, Jon, *RFC 862: Echo Protocol*, Network Information Center, SRI International, Menlo Park, CA, May 1983.

inetd(1M)

Postel, Jon, *RFC 863: Discard Protocol*, Network Information Center, SRI International, Menlo Park, CA, May 1983.

Postel, Jon, *RFC 864: Character Generator Protocol*, Network Information Center, SRI International, Menlo Park, CA, May 1983.

Postel, Jon, *RFC 867: Daytime Protocol*, Network Information Center, SRI International, Menlo Park, CA, May 1983.

Postel, Jon, and Ken Harrenstien, *RFC 868: Time Protocol*, Network Information Center, SRI International, Menlo Park, CA, May 1983.

WARNINGS Do not configure `udp` services as `nowait`. This will cause a race condition where the `inetd` program selects on the socket and the server program reads from the socket. Many server programs will be forked and performance will be severely compromised.

NOTES For RPC services, `inetd` listens on all the transports (not only `tcp` and `udp`) as specified for each service in the `inetd.conf(4)` file.

NAME	in.ftpd, ftpd – File transfer protocol server
SYNOPSIS	in.ftpd [-dl] [-t <i>timeout</i>]
DESCRIPTION	in.ftpd is the Internet File Transfer Protocol (FTP) server process. The server is invoked by the Internet daemon inetd(1M) each time a connection to the FTP service (see services(4)) is made.
OPTIONS	<p>-d Debugging information is logged to the system log daemon syslogd(1).</p> <p>-l Each FTP session is logged to the system log daemon syslogd(1).</p> <p>-t <i>timeout</i> Set the inactivity timeout period to <i>timeout</i> seconds. The FTP server will timeout an inactive session after 15 minutes.</p>
Requests	<p>The FTP server currently supports the following FTP requests; case is not distinguished.</p> <p>ABOR abort previous command</p> <p>ACCT specify account (ignored)</p> <p>ALLO allocate storage (vacuously)</p> <p>APPE append to a file</p> <p>CDUP change to parent of current working directory</p> <p>CWD change working directory</p> <p>DELE delete a file</p> <p>HELP give help information</p> <p>LIST give list files in a directory (ls -lg)</p> <p>MKD make a directory</p> <p>MODE specify data transfer <i>mode</i></p> <p>NLST give name list of files in directory (ls)</p> <p>NOOP do nothing</p> <p>PASS specify password</p> <p>PASV prepare for server-to-server transfer</p> <p>EPSV extended passive command request</p> <p>LPSV long passive command request</p> <p>PORT specify data connection port</p> <p>EPRT specify extended address for the transport connection</p> <p>LPRT specify “long” address for the transport connection</p>

in.ftpd(1M)

PWD	print the current working directory
QUIT	terminate session
RETR	retrieve a file
RMD	remove a directory
RNFR	specify rename-from file name
RNTO	specify rename-to file name
STOR	store a file
STOU	store a file with a unique name
STRU	specify data transfer <i>structure</i>
TYPE	specify data transfer <i>type</i>
USER	specify user name
XCUP	change to parent of current working directory
XCWD	change working directory
XMKD	make a directory
XPWD	print the current working directory
XRMD	remove a directory

The remaining FTP requests specified in *RFC 959* are recognized, but not implemented.

The FTP server will abort an active file transfer only when the ABOR command is preceded by a Telnet "Interrupt Process" (IP) signal and a Telnet "Synch" signal in the command Telnet stream, as described in *RFC 959*. `in.ftpd` interprets file names according to the "globbing" conventions used by `sh(1)`. This allows users to utilize the metacharacters: `* ? [] { } ~` `in.ftpd`'s `umask` (which it uses to create files during PUT operations) may be adjusted by adding the line

```
UMASK=nnn
```

to `/etc/default/ftpd`.

The banner returned by `in.ftpd` in the parenthetical portion of its greeting is configurable. The default is equivalent to `"uname -sr"` and will be used if no banner is set in `/etc/default/ftpd`. To set the banner, add a line of the form

```
BANNER="..."
```

to `/etc/default/ftpd`. Nonempty banner strings are fed to shells for evaluation.

The default banner may also be obtained by

```
BANNER="'uname -s' 'uname -r'"
```

and no banner will be printed if `/etc/default/ftpd` contains

```
BANNER="
```

`in.ftpd` authenticates users according to five rules.

First, the user name must be in the password data base, `/etc/passwd`, and have a password that is not NULL. A password must always be provided by the client before any file operations may be performed. The PAM framework (see `SECURITY` below) is used to verify that the correct password was entered.

Second, if the user name appears in the file `/etc/ftpusers`, ftp access is denied. The default list of users in `/etc/ftpusers` includes all of the accounts in `passwd(4)`. See `ftpusers(4)`.

Third, ftp access is denied if the user's shell is not listed in `getusershell(3C)`

Fourth, if the user name is "anonymous" or "ftp", an entry for the user name ftp must be present in the password and shadow files. The user is then allowed to log in by specifying any password — by convention this is given as the user's e-mail address (such as `user@host.Sun.COM`). Do not specify a valid shell in the password entry of the ftp user, and do not give it a valid password (use NP in the encrypted password field of the shadow file).

Fifth, access is denied unless a user has the remote login authorization. If the `/etc/nologin` file exists, access is denied.

For anonymous ftp users, `in.ftpd` takes special measures to restrict the client's access privileges. The server performs a `chroot(2)` command to the home directory of the "ftp" user. In order that system security is not breached, it is recommended that the "ftp" subtree be constructed with care; the following rules are suggested.

`~ftp`

Make the home directory owned by root and unwritable by anyone.

`~ftp/bin`

Make this directory owned by root and unwritable by anyone. Make this a symbolic link to `~ftp/usr/bin`. The program `ls(1)` must be present to support the list commands. This program should have mode 111.

`~ftp/usr/lib`

Make this directory owned by root and unwritable by anyone. Copy the following shared libraries from `/usr/lib` into this directory:

```
ld.so.1*
libc.so.1*
libdl.so.1*
libmp.so.2*
libnsl.so.1*
libsocket.so.1*
```

in.ftpd(1M)

```
nss_compat.so.1*
nss_dns.so.1*
nss_files.so.1*
nss_nis.so.1*
nss_nisplus.so.1*
nss_xfn.so.1*
straddr.so*
straddr.so.2*
```

~ftp/etc

Make this directory owned by root and unwritable by anyone. Copies of the files `passwd(4)`, `group(4)`, and `netconfig(4)` must be present for the `ls(1)` command to work properly. These files should be mode 444.

~ftp/pub

Make this directory mode 755 and owned by root. Users should then place files which are to be accessible via the anonymous account in this directory.

~ftp/dev

Make this directory owned by root and unwritable by anyone. First perform `ls -lL` on the device files listed below to determine their major and minor numbers, then use `mknod` to create them in this directory.

```
/dev/zero
/dev/tcp
/dev/udp
/dev/ticotsord
```

Set the read and write mode on these nodes to 666 so that passive ftp will not fail with “permission denied” errors.

~ftp/usr/share/lib/zoneinfo

Make this directory mode 555 and owned by root. Copy its contents from `/usr/share/lib/zoneinfo`. This enables `ls -l` to display time and date stamps correctly.

SECURITY

`in.ftpd` uses `pam(3PAM)` for authentication, account management, and session management. The PAM configuration policy, listed through `/etc/pam.conf`, specifies the module to be used for `in.ftpd`. Here is a partial `pam.conf` file with entries for the `in.ftpd` command using the UNIX authentication, account management, and session management module.

ftp	auth	required	/usr/lib/security/pam_unix.so.1
ftp	account	required	/usr/lib/security/pam_unix.so.1
ftp	session	required	/usr/lib/security/pam_unix.so.1

If there are no entries for the ftp service, then the entries for the "other" service will be used. Unlike login, passwd, and other commands, the ftp protocol will only support a single password. Using multiple modules will prevent in.ftpd from working properly.

USAGE The in.ftpd command is IPv6-enabled. See ip6(7P).

EXAMPLES **EXAMPLE 1** Setting Up An Anonymous Ftp

To set up anonymous ftp, add the following entry to the /etc/passwd file. In this example, /export/ftp was chosen to be the anonymous ftp area, and the shell is the non-existent file /nosuchshell. This prevents users from logging in as the ftp user.

```
ftp:x:30000:30000:Anonymous FTP:/export/ftp:/nosuchshell
```

Add the following entry to the /etc/shadow file:

```
ftp:NP:6445:.....:
```

The following shell script sets up the anonymous ftp area. It presumes that names are resolved using NIS.

```
#!/bin/sh
# script to setup anonymous ftp area
#

# verify you are root
/usr/bin/id | grep -w 'uid=0' >/dev/null 2>&1
if [ "$?" != "0" ]; then
    echo
    exit 1
fi

# handle the optional command line argument
case $# in

    # the default location for the anon ftp comes from the passwd
    # file
    0) ftphome=`getent passwd ftp | cut -d: -f6`
        ;;

    1) if [ "$1" = "start" ]; then
        ftphome=`getent passwd ftp | cut -d: -f6`
        else
            ftphome=$1
        fi
        ;;

    *) echo "Usage: $0 [anon-ftp-root]"
        exit 1
        ;;

esac

if [ -z "${ftphome}" ]; then
    echo "$0: ftphome must be non-null"
    exit 2
```

EXAMPLE 1 Setting Up An Anonymous Ftp (Continued)

```

fi

case ${ftphome} in
    /*) # ok
        ;;

    *) echo "$0: ftphome must be an absolute pathname"
       exit 1
       ;;
esac

# This script assumes that ftphome is neither / nor /usr so ...
if [ -z "${ftphome}" -o "${ftphome}" = "/" -o "${ftphome}" = "/usr" ]; then
    echo "$0: ftphome must be non-null and neither / or /usr"
    exit 2
fi

# If ftphome does not exist but parent does, create ftphome
if [ ! -d ${ftphome} ]; then
    # lack of -p below is intentional
    mkdir ${ftphome}
fi
chown root ${ftphome}
chmod 555 ${ftphome}

echo Setting up anonymous ftp area ${ftphome}

# Ensure that the /usr directory exists
if [ ! -d ${ftphome}/usr ]; then
    mkdir -p ${ftphome}/usr
fi
# Now set the ownership and modes to match the man page
chown root ${ftphome}/usr
chmod 555 ${ftphome}/usr

# Ensure that the /usr/bin directory exists
if [ ! -d ${ftphome}/usr/bin ]; then
    mkdir -p ${ftphome}/usr/bin
fi
# Now set the ownership and modes to match the man page
chown root ${ftphome}/usr/bin
chmod 555 ${ftphome}/usr/bin

# this may not be the right thing to do
# but we need the bin -> usr/bin link
rm -f ${ftphome}/bin
ln -s usr/bin ${ftphome}/bin

# Ensure that the /usr/lib and /etc directories exist
if [ ! -d ${ftphome}/usr/lib ]; then
    mkdir -p ${ftphome}/usr/lib
fi
chown root ${ftphome}/usr/lib
chmod 555 ${ftphome}/usr/lib

```

EXAMPLE 1 Setting Up An Anonymous Ftp (Continued)

```

if [ ! -d ${ftphome}/usr/lib/security ]; then
    mkdir -p ${ftphome}/usr/lib/security
fi
chown root ${ftphome}/usr/lib/security
chmod 555 ${ftphome}/usr/lib/security

if [ ! -d ${ftphome}/etc ]; then
    mkdir -p ${ftphome}/etc
fi
chown root ${ftphome}/etc
chmod 555 ${ftphome}/etc

# a list of all the commands that should be copied to
# ${ftphome}/usr/bin
# /usr/bin/ls is needed at a minimum.
ftpcmd="
    /usr/bin/ls
"

# ${ftphome}/usr/lib needs to have all the libraries needed by the above
# commands, plus the runtime linker, and some name service libraries
# to resolve names. We just take all of them here.

ftplib="`ldd $ftpcmd | awk ' $3 ~ /lib/ { print $3 }' | sort | uniq`"
ftplib="$ftplib /usr/lib/nss_* /usr/lib/straddr* /usr/lib/libmp.so*"
ftplib="$ftplib /usr/lib/libnsl.so.1 /usr/lib/libsocket.so.1 \\\
/usr/lib/ld.so.1"
ftplib="`echo $ftplib | tr ' ' '\
' | sort | uniq`"

cp ${ftplib} ${ftphome}/usr/lib
chmod 555 ${ftphome}/usr/lib/*

cp /usr/lib/security/* ${ftphome}/usr/lib/security
chmod 555 ${ftphome}/usr/lib/security/*

cp ${ftpcmd} ${ftphome}/usr/bin
chmod 111 ${ftphome}/usr/bin/*

# you also might want to have separate minimal versions of passwd
# and group
cp /etc/passwd /etc/group /etc/netconfig /etc/pam.conf ${ftphome}/etc
chmod 444 ${ftphome}/etc/*

# need /etc/default/init for timezone to be correct
if [ ! -d ${ftphome}/etc/default ]; then
    mkdir ${ftphome}/etc/default
fi
chown root ${ftphome}/etc/default
chmod 555 ${ftphome}/etc/default
cp /etc/default/init ${ftphome}/etc/default
chmod 444 ${ftphome}/etc/default/init

```

EXAMPLE 1 Setting Up An Anonymous Ftp (Continued)

```

# Copy timezone database
mkdir -p ${ftphome}/usr/share/lib/zoneinfo
(cd ${ftphome}/usr/share/lib/zoneinfo
  (cd /usr/share/lib/zoneinfo; find . -print |
    cpio -o) 2>/dev/null | cpio -imdu 2>/dev/null
  find . -print | xargs chmod 555
  find . -print | xargs chown root
)

# Ensure that the /dev directory exists
if [ ! -d ${ftphome}/dev ]; then
  mkdir -p ${ftphome}/dev
fi

# make device nodes. ticotsord and udp are necessary for
# 'ls' to resolve NIS names.

for device in zero tcp udp ticotsord ticlts
do
  line=`ls -lL /dev/${device} | sed -e 's/,//'`
  major=`echo $line | awk '{print $5}'`
  minor=`echo $line | awk '{print $6}'`
  rm -f ${ftphome}/dev/${device}
  mknod ${ftphome}/dev/${device} c ${major} ${minor}
done

chmod 666 ${ftphome}/dev/*

## Now set the ownership and modes
chown root ${ftphome}/dev
chmod 555 ${ftphome}/dev

# uncomment the below if you want a place for people to store
# things, but beware the security implications
# if [ ! -d ${ftphome}/pub ]; then
#   mkdir -p ${ftphome}/pub
# fi
# chown root ${ftphome}/pub
# chmod 1755 ${ftphome}/pub

```

After running this script, edit the files in `~ftp/etc` to make sure all non-public information is removed.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

		in.ftpd(1M)
FILES	/etc/default/ftpd /etc/ftpusers	file listing users for whom ftp login privileges are disallowed.
SUMMARY OF TRUSTED SOLARIS CHANGES Trusted Solaris 8 4/01 Reference Manual Sun Microsystems Reference Manual	Login is not allowed unless the user has the remote login authorization. If the /etc/nologin file exists, the user is not allowed to log in. inetd(1M), chroot(2), getsockopt(3SOCKET), inetd.conf(4) ftp(1), ld.so.1(1), ls(1), sh(1), aset(1M), mknod(1M), syslogd(1M), pam(3PAM), ftpusers(4), group(4), netconfig(4), netrc(4), pam.conf(4), passwd(4), services(4), attributes(5), pam_unix(5) Allman, M., Ostermann, S., and Metz, C., <i>RFC 2428, FTP Extensions for IPv6 and NATs</i> , The Internet Society, 1998. Postel, Jon, and Joyce Reynolds, <i>RFC 959, File Transfer Protocol (FTP)</i> , Network Information Center, SRI International, Menlo Park, Calif., October 1985. Piscitello, D., <i>RFC 1639, FTP Operation Over Big Address Records (FOOBAR)</i> , Network Working Group, June 1994.	
DIAGNOSTICS	in.ftpd logs various errors to syslogd, with a facility code of daemon.	
Info Severity	These messages are logged only if the -l flag is specified. FTPD: connection from <i>host</i> at <i>time</i> A connection was made to ftpd from the host <i>host</i> at the date and time <i>time</i> . FTPD: User <i>user</i> timed out after <i>timeout</i> seconds at <i>time</i> The user <i>user</i> was logged out because they had not entered any commands after <i>timeout</i> seconds; the logout occurred at the date and time <i>time</i> .	
Debug Severity	These messages are logged only if the -d flag is specified. FTPD: command: <i>command</i> A command line containing <i>command</i> was read from the FTP client. lost connection The FTP client dropped the connection. <— <i>replycode</i> <— <i>replycode</i> — A reply was sent to the FTP client with the reply code <i>replycode</i> . The next message logged will include the message associated with the reply. If a – follows the reply code, the reply is continued on later lines.	
NOTES	The anonymous ftp account is inherently dangerous and should be avoided when possible.	

in.ftpd(1M)

The name service caching daemon `/usr/sbin/nscd` may interfere with some of the functionality of anonymous ftp. The *sublogin* feature does not work unless caching for `passwd` is disabled in `/etc/nscd.conf`.

The server must run as the superuser to create sockets with privileged port numbers. It maintains an effective user ID of the logged in user, reverting to the superuser only when binding addresses to sockets. The possible security holes have been extensively scrutinized, but may be incomplete.

The file `/etc/ftpusers`, which is now included as part of Solaris, contains a list of users who cannot access the system; the default list of users in `/etc/ftpusers` includes all of the accounts in `passwd(4)`. See `ftpusers(4)`.

NAME	init, telinit – Process control initialization								
SYNOPSIS	<p>/sbin/init [0123456abcQqSs]</p> <p>/etc/telinit [0123456abcQqSs]</p>								
DESCRIPTION	init is a general process spawner. Its primary role is to create processes from information stored in the file <code>/etc/inittab</code> .								
Run Level Defined	At any given time, the system is in one of eight possible run levels. A run level is a software configuration under which only a selected group of processes exists. Processes spawned by <code>init</code> for each of these run levels are defined in <code>/etc/inittab</code> . <code>init</code> can be in one of eight run levels, 0–6 and S or s (S and s are identical). The run level changes when a privileged user runs <code>/sbin/init</code> . This sends appropriate signals to the original <code>init</code> spawned by the operating system at boot time, saying which run level to invoke.								
init and System Booting	<p>When the system is booted, <code>init</code> is invoked and the following occurs. First, it reads <code>/etc/default/init</code> to set environment variables. This is typically where TZ (time zone) and locale-related environments such as LANG or LC_CTYPE get set.</p> <p><code>init</code> then looks in <code>/etc/inittab</code> for the <code>initdefault</code> entry [see <code>inittab(4)</code>]. If the <code>initdefault</code> entry:</p> <table> <tr> <td>exists</td><td><code>init</code> usually uses the run level specified in that entry as the initial run level to enter.</td></tr> <tr> <td>does not exist</td><td><code>/etc/inittab</code>, <code>init</code> asks the user to enter a run level from the system console.</td></tr> <tr> <td>S or s</td><td><code>init</code> goes to the single-user state. In this state, the system console device (<code>/dev/console</code>) is opened for reading and writing and the command <code>/sbin/su</code>, (see <code>su(1M)</code>), is invoked. Use either <code>init</code> or <code>telinit</code> to change the run level of the system. Note that if the shell is terminated (using an end-of-file), <code>init</code> only re-initializes to the single-user state if <code>/etc/inittab</code> does not exist.</td></tr> <tr> <td>0–6</td><td><code>init</code> enters the corresponding run level. Run levels 0, 5, and 6 are reserved states for shutting the system down. Run levels 2, 3, and 4 are available as multi-user operating states.</td></tr> </table> <p>If this is the first time since power up that <code>init</code> has entered a run level other than single-user state, <code>init</code> first scans <code>/etc/inittab</code> for <code>boot</code> and <code>bootwait</code> entries (see <code>inittab(4)</code>). These entries are performed before any other processing of <code>/etc/inittab</code> takes place, providing that the run level entered matches that of the entry. In this way any special initialization of the operating system, such as mounting</p>	exists	<code>init</code> usually uses the run level specified in that entry as the initial run level to enter.	does not exist	<code>/etc/inittab</code> , <code>init</code> asks the user to enter a run level from the system console.	S or s	<code>init</code> goes to the single-user state. In this state, the system console device (<code>/dev/console</code>) is opened for reading and writing and the command <code>/sbin/su</code> , (see <code>su(1M)</code>), is invoked. Use either <code>init</code> or <code>telinit</code> to change the run level of the system. Note that if the shell is terminated (using an end-of-file), <code>init</code> only re-initializes to the single-user state if <code>/etc/inittab</code> does not exist.	0–6	<code>init</code> enters the corresponding run level. Run levels 0, 5, and 6 are reserved states for shutting the system down. Run levels 2, 3, and 4 are available as multi-user operating states.
exists	<code>init</code> usually uses the run level specified in that entry as the initial run level to enter.								
does not exist	<code>/etc/inittab</code> , <code>init</code> asks the user to enter a run level from the system console.								
S or s	<code>init</code> goes to the single-user state. In this state, the system console device (<code>/dev/console</code>) is opened for reading and writing and the command <code>/sbin/su</code> , (see <code>su(1M)</code>), is invoked. Use either <code>init</code> or <code>telinit</code> to change the run level of the system. Note that if the shell is terminated (using an end-of-file), <code>init</code> only re-initializes to the single-user state if <code>/etc/inittab</code> does not exist.								
0–6	<code>init</code> enters the corresponding run level. Run levels 0, 5, and 6 are reserved states for shutting the system down. Run levels 2, 3, and 4 are available as multi-user operating states.								

init(1M)

	<p>file systems, can take place before users are allowed onto the system. <code>init</code> then scans <code>/etc/inittab</code> and executes all other entries that are to be processed for that run level.</p> <p>To spawn each process in <code>/etc/inittab</code>, <code>init</code> reads each entry and for each entry that should be respawned, it forks a child process. After it has spawned all of the processes specified by <code>/etc/inittab</code>, <code>init</code> waits for one of its descendant processes to die, a powerfail signal, or a signal from another <code>init</code> or <code>telinit</code> process to change the system's run level. When one of these conditions occurs, <code>init</code> re-examines <code>/etc/inittab</code>.</p>												
inittab Additions	<p>New entries can be added to <code>/etc/inittab</code> at any time; however, <code>init</code> still waits for one of the above three conditions to occur before re-examining <code>/etc/inittab</code>. To get around this, <code>init Q</code> or <code>init q</code> command wakes <code>init</code> to re-examine <code>/etc/inittab</code> immediately.</p> <p>When <code>init</code> comes up at boot time and whenever the system changes from the single-user state to another run state, <code>init</code> sets the <code>ioctl(2)</code> states of the console to those modes saved in the file <code>/etc/ioctl.syscon</code>. <code>init</code> writes this file whenever the single-user state is entered.</p>												
Run Level Changes	<p>When a run level change request is made, <code>init</code> sends the warning signal (<code>SIGTERM</code>) to all processes that are undefined in the target run level. <code>init</code> waits five seconds before forcibly terminating these processes by sending a kill signal (<code>SIGKILL</code>).</p> <p>When <code>init</code> receives a signal telling it that a process it spawned has died, it records the fact and the reason it died in <code>/var/adm/utmpx</code> and <code>/var/adm/wtmpx</code> if it exists (see <code>who(1)</code>). A history of the processes spawned is kept in <code>/var/adm/wtmpx</code>.</p> <p>If <code>init</code> receives a powerfail signal (<code>SIGPWR</code>) it scans <code>/etc/inittab</code> for special entries of the type <code>powerfail</code> and <code>powerwait</code>. These entries are invoked (if the run levels permit) before any further processing takes place. In this way <code>init</code> can perform various cleanup and recording functions during the powerdown of the operating system.</p>												
/etc/defaults/init File	<p>Default values can be set for the following flags in <code>/etc/default/init</code>. For example: <code>TZ=US/Pacific</code></p> <table><tr><td><code>TZ</code></td><td>Either specifies the timezone information (see <code>ctime(3C)</code>) or the name of a timezone information file <code>/usr/share/lib/zoneinfo</code>.</td></tr><tr><td><code>LC_CTYPE</code></td><td>Character characterization information.</td></tr><tr><td><code>LC_MESSAGES</code></td><td>Message translation.</td></tr><tr><td><code>LC_MONETARY</code></td><td>Monetary formatting information.</td></tr><tr><td><code>LC_NUMERIC</code></td><td>Numeric formatting information.</td></tr><tr><td><code>LC_TIME</code></td><td>Time formatting information.</td></tr></table>	<code>TZ</code>	Either specifies the timezone information (see <code>ctime(3C)</code>) or the name of a timezone information file <code>/usr/share/lib/zoneinfo</code> .	<code>LC_CTYPE</code>	Character characterization information.	<code>LC_MESSAGES</code>	Message translation.	<code>LC_MONETARY</code>	Monetary formatting information.	<code>LC_NUMERIC</code>	Numeric formatting information.	<code>LC_TIME</code>	Time formatting information.
<code>TZ</code>	Either specifies the timezone information (see <code>ctime(3C)</code>) or the name of a timezone information file <code>/usr/share/lib/zoneinfo</code> .												
<code>LC_CTYPE</code>	Character characterization information.												
<code>LC_MESSAGES</code>	Message translation.												
<code>LC_MONETARY</code>	Monetary formatting information.												
<code>LC_NUMERIC</code>	Numeric formatting information.												
<code>LC_TIME</code>	Time formatting information.												

	LC_ALL	If set, all other LC_* environmental variables take on this value.
	LANG	If LC_ALL is not set, and any particular LC_* is also not set, the value of LANG is used for that particular environmental variable.
telinit	telinit, which is linked to /sbin/init, is used to direct the actions of init. It takes a one-character argument and signals init to take the appropriate action.	
SECURITY	init uses pam(3PAM) for session management. The PAM configuration policy, listed through /etc/pam.conf, specifies the session management module to be used for init. Here is a partial pam.conf file with entries for init using the UNIX session management module.	
	<pre>init session required /usr/lib/security/pam_unix.so.1</pre>	
	If there are no entries for the init service, then the entries for the "other" service will be used.	
OPTIONS	0	Go into firmware.
	1	Put the system in system administrator mode. All local file systems are mounted. Only a small set of essential kernel processes are left running. This mode is for administrative tasks such as installing optional utility packages. All files are accessible and no users are logged in on the system.
	2	Put the system in multi-user mode. All multi-user environment terminal processes and daemons are spawned. This state is commonly referred to as the multi-user state.
	3	Extend multi-user mode by making local resources available over the network.
	4	Is available to be defined as an alternative multi-user environment configuration. It is not necessary for system operation and is usually not used.
	5	Shut the machine down so that it is safe to remove the power. Have the machine remove power, if possible.
	6	Stop the operating system and reboot to the state defined by the initdefault entry in /etc/inittab.
	a, b, c	Process only those /etc/inittab entries having the a, b, or c run level set. These are pseudo-states, which may be defined to run certain commands, but which do not cause the current run level to change.
	Q, q	Re-examine /etc/inittab.
	S, s	Enter single-user mode. This is the only run level that doesn't require the existence of a properly formatted /etc/inittab file. If this file does not exist, then by default, the only legal run level

init(1M)

that `init` can enter is the single-user mode. When in single-user mode, the filesystems required for basic system operation will be mounted. When the system comes down to single-user mode, these file systems will remain mounted (even if provided by a remote file server), and any other local filesystems will also be left mounted. During the transition down to single-user mode, all processes started by `init` or `init.d` scripts that should only be running in multi-user mode are killed. In addition, any process that has a `utmpx` entry will be killed. This last condition insures that all port monitors started by the SAC are killed and all services started by these port monitors, including `ttymon` login services, are killed.

FILES	<code>/etc/inittab</code>	Controls process dispatching by <code>init</code> .
	<code>/var/adm/utmpx</code>	User access and administration information
	<code>/var/adm/wtmpx</code>	History of user access and administration information
	<code>/etc/ioctl.syscon</code>	System console states.
	<code>/dev/console</code>	System console device.
	<code>/etc/default/init</code>	Environment variables.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES Since SunOS 4.0 Reference Manual

`init` requires privilege to run in the Trusted Solaris environment.

`login(1)`, `kill(2)`, `inittab(4)`

`sh(1)`, `stty(1)`, `who(1)`, `shutdown(1M)`, `su(1M)`, `ttymon(1M)`, `ioctl(2)`, `ctime(3C)`, `pam(3PAM)`, `pam.conf(4)`, `utmpx(4)`, `attributes(5)`, `pam_unix(5)`, `termio(7I)`

DIAGNOSTICS

If `init` finds that it is respawning an entry from `/etc/inittab` more than ten times in two minutes, assumes that there is an error in the command string in the entry, and generates an error message on the system console. It will then refuse to respawn this entry until either five minutes has elapsed or it receives a signal from a user-spawned `init` or `telinit`. This prevents `init` from eating up system resources when someone makes a typographical error in the `inittab` file, or a program is removed that is referenced in `/etc/inittab`.

NOTES

`init` and `telinit` can be run only by a privileged user.

The `S` or `s` state must not be used indiscriminately in `/etc/inittab`. When modifying this file, it is best to avoid adding this state to any line other than `initdefault`.

If a default state is not specified in the `initdefault` entry in `/etc/inittab`, state `6` is entered. Consequently, the system will loop by going to firmware and rebooting continuously.

If the `utmpx` file cannot be created when booting the system, the system will boot to state `"s"` regardless of the state specified in the `initdefault` entry in `/etc/inittab`. This can occur if the `/var` file system is not accessible.

init.wbem(1M)

NAME	init.wbem – start and stop the CIM Boot Manager
SYNOPSIS	<code>/etc/init.d/init.wbem start stop status</code>
DESCRIPTION	<p>The <code>init.wbem</code> utility is run automatically during installation and each time the system is rebooted. This utility manipulates the CIM Object Manager (CIMOM) and the Solaris Management Console (SMC) server, both of which run combined in a single process. <code>init.wbem</code> can be used to start, stop, or retrieve status from the server.</p>
CIM Object Manager	<p>The CIM Object Manager manages CIM objects on a WBEM-enabled system. A CIM object is a computer representation, or model, of a managed resource, such as a printer, disk drive, or CPU. CIM objects are stored internally as Java classes.</p> <p>When a WBEM client application accesses information about a CIM object, the CIM Object Manager contacts either the appropriate provider for that object or the CIM Object Manager Repository. Providers are classes that communicate with managed objects to access data.</p> <p>When a WBEM client application requests data from a managed resource that is not available from the CIM Object Manager Repository, the CIM Object Manager forwards the request to the provider for that managed resource. The provider dynamically retrieves the information.</p> <p>At startup, the CIM Object Manager performs the following functions:</p> <ul style="list-style-type: none">■ Listens for RMI connections on RMI port 5987 and for XML/HTTP connections on HTTP port 80.■ Sets up a connection to the CIM Object Manager Repository.■ Waits for incoming requests. <p>During normal operations, the CIM Object Manager performs the following functions:</p> <ul style="list-style-type: none">■ Performs security checks to authenticate user login and authorization to access namespaces.■ Performs syntactical and semantic checking of CIM data operations to ensure that they comply with the latest CIM Specification.■ Routes requests to the appropriate provider or to the CIM Object Manager Repository.■ Delivers data from providers and from the CIM Object Manager Repository to WBEM client applications. <p>A WBEM client application contacts the CIM Object Manager to establish a connection when it needs to perform WBEM operations, such as creating a CIM class or updating a CIM instance. When a WBEM client application connects to a CIM Object Manager, it gets a reference to the CIM Object Manager, which it then uses to request services and operations.</p>

	init.wbem(1M)						
Solaris Management Console Server	The SMC server is the back end to the front end console, <code>smc(1M)</code> . It provides tools for the console to download and performs common services for the console and its tools to use, such as authentication, authorization, logging, messaging, and persistence.						
System Booting	The <code>init.wbem</code> script is installed in the <code>/etc/init.d</code> directory. A link to it exists in <code>/etc/rc2.d/S90wbem</code> , which is run with the <code>start</code> option when init state 2 is entered (normally at boot time). Other links to it exist in <code>/etc/rc0.d/K36wbem</code> , <code>/etc/rc1.d/K36wbem</code> , and <code>/etc/rcS.d/K36wbem</code> , which are run with the <code>stop</code> option when init states 0, 1, and S are entered (normally at system halt, or when entering “system administrator mode” or single user mode).						
OPTIONS	<p>The following options are supported:</p> <table> <tr> <td><code>start</code></td><td>Starts the CIMOM and SMC server on the local host.</td></tr> <tr> <td><code>stop</code></td><td>Stops the CIMOM and SMC server on the local host.</td></tr> <tr> <td><code>status</code></td><td>Gets the status of the CIMOM and SMC server on the local host.</td></tr> </table>	<code>start</code>	Starts the CIMOM and SMC server on the local host.	<code>stop</code>	Stops the CIMOM and SMC server on the local host.	<code>status</code>	Gets the status of the CIMOM and SMC server on the local host.
<code>start</code>	Starts the CIMOM and SMC server on the local host.						
<code>stop</code>	Stops the CIMOM and SMC server on the local host.						
<code>status</code>	Gets the status of the CIMOM and SMC server on the local host.						
NOTES	<p>When the <code>init.wbem</code> script is run, it does not run the CIMOM and SMC server directly. The server process is in Java and is too heavyweight to be run immediately at system boot time. Instead, a lightweight process is run which listens on the ports the CIMOM and the SMC server normally use, running the two servers the first time it gets a connection on either port, thus acting similarly to <code>inetd(1M)</code>.</p> <p>Because Java programs cannot inherit file descriptors as other programs can, there is a small time period from when the first connection is made until the server is fully operational where client connections may be dropped. WBEM clients are immune to this, as they will retry until the server comes online. SMC clients are not immune, and it may be necessary to manually reconnect, though this should not happen in the common case.</p>						
EXAMPLES	<p>EXAMPLE 1 Restarting the SMC server after enabling users on untrusted clients to assume a role</p> <p>The <code>start</code> and <code>stop</code> options can be used to restart the SMC server, as is necessary when enabling users on untrusted clients to assume a role via the SMC login dialog.</p> <ol style="list-style-type: none"> 1. Assume the <code>secadmin</code> role. 2. Edit the file <code>/usr/sadm/lib/smc/bin/smcwbemserver</code> and change <pre>com.sun.management.viperimpl.server.ViperWbemServer "\$@" & to com.sun.management.viperimpl.server.ViperWbemServer "\$@" -u &</pre> 3. Restart the SMC server. <pre>\$ /etc/init.d/init.wbem stop \$ /etc/init.d/init.wbem start</pre> <p>See the <i>Trusted Solaris Administrator's Procedures</i> for details of this procedure.</p>						

init.wbem(1M)

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWwbcor

**Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual**

inetd(1M), smc(1M)

mofcomp(1M), smcconf(1M), wbemadmin(1M), wbemlogviewer(1M),
attributes(5)

NAME	in.named, named – Internet domain name server	
SYNOPSIS	in.named [-d <i>debuglevel</i>] [-q] [-r] [-f] [-p <i>remote/local-port</i>] [-w <i>dirname</i>] [[-b -c] <i>configfile</i>]	
DESCRIPTION	<p>in.named is the Internet domain name server. in.named spawns the named-xfer process whenever it needs to perform a zone transfer. See named-xfer(1M).</p> <p>The in.named name service is used by hosts on the Internet to provide access to the Internet distributed naming database. See <i>RFC 1034</i> and <i>RFC 1035</i> for more information on the Internet domain name system.</p> <p>With no arguments, in.named reads the default configuration file <i>/etc/named.conf</i> for any initial data, and listens for queries. Any additional arguments beyond those shown in the SYNOPSIS section are interpreted as the names of configuration files. If multiple configuration files are specified, only the last is used.</p> <p>The name server reads the configuration file to obtain instructions on where to find its initial data.</p> <p>In the Trusted Solaris environment, in.named listens for input requests on a multilevel port (MLP) and sends responses to the DNS client at the sensitivity label of the client's request. Thus, though in.named runs at the sensitivity label ADMIN_LOW, it can accept requests at any sensitivity label. in.named can also serve DNS clients and communicate with other DNS name servers on either Trusted Solaris hosts or non-trusted hosts.</p> <p>The DNS name server running on a Trusted Solaris machine is viewed as a supplier of public information, and the name database that it maintains is considered trusted. in.named requires the trusted path attribute, and it requires that the <i>/etc/named.boot</i> file, zone files, and other configuration files that it uses be at the sensitivity label ADMIN_LOW. As part of the name database, these files and their contents are also considered trusted; thus in.named can query any DNS name server specified in the files. The DNS name servers specified in these files may reside on either Trusted Solaris hosts or non-trusted hosts.</p>	
OPTIONS	<p>-b <i>configfile</i> Use <i>configfile</i> rather than <i>/etc/named.conf</i>. This option allows filenames to begin with a leading dash.</p> <p>-c <i>configfile</i> Use <i>configfile</i> rather than <i>/etc/named.conf</i>. This option allows filenames to begin with a leading dash.</p> <p>-d <i>level</i> Print debugging information. <i>level</i> is a number indicating the level of messages printed.</p> <p>-f Run this process in the foreground. The process will not fork(2). By default, in.named runs as a daemon in the background.</p> <p>-p <i>remote/local-port</i> Use different port numbers. The default is the standard port number as returned</p>	

in.named(1M)

	<p>bygetservbyname(3SOCKET) for service domain.</p> <p>The <code>-p</code> argument can specify up to two port numbers. The specification of two port numbers requires a <code>''</code> (slash) separator. In this case, the first port is used when contacting remote servers, and the second one is the service port bound by the local instance of <code>in.named</code>. This option is used mostly for debugging purposes.</p>
<code>-q</code>	Trace all incoming queries. Note: this option is ignored in favor of the boot file directive, <code>options query-log</code> , when both options are used.
<code>-r</code>	Turns recursion off in the server. Answers can come only from local (primary or secondary) zones. This option can be used on root servers. Note: This option will probably be eventually abandoned in favor of the boot file directive, <code>options no-recursion</code> .
<code>-w <i>dirname</i></code>	Change the current working directory of <code>in.named</code> to <i>dirname</i> .
<p><code>/etc/named.conf</code> File Directives</p>	<p>The following is a simple configuration file <code>/etc/named.conf</code> containing directives to guide the <code>in.named</code> process at startup time.</p> <pre>options { directory "/usr/local/adm/named"; pid-file "/var/named/named.pid"; named-xfer "/usr/sbin/named-xfer"; forwarders { 10.0.0.78; 10.2.0.78; }; transfers-in 10; forward only; fake-iquery yes; pollfd-chunk-size 20; }; logging { category lame-servers { null; }; category cname { null; }; }; zone "." in { type hint; file "root.cache"; }; zone "cc.berkeley.edu" in { type slave; file "128.32.137.3"; masters { 128.32.137.8; }; };</pre>

```

zone "6.32.128.in-addr.arpa" in {
    type slave;
    file "128.32.137.3";
    masters { 128.32.137.8; };
};
zone "0.0.127.in-addr.arpa" in {
    type master;
    file "master/db.127";
};

zone "berkeley.edu" in {
    type master;
    file "berkeley.edu.zone";
};

zone "32.128.in-addr.arpa" in {
    type master;
    file "ucbhosts.rev";
};

```

}; The configuration file consists of sections and comments. Sections end with a ';' and contain statements which are enclosed in '{ }' and may span multiple lines. The following sections are supported: options, zone, server, logging, acl, include, and key.

Comments Syntax

The following are examples of comments syntax in BIND 8.1:

```

/* This is a BIND comment as in C */
// This is a BIND comment as in C++
# This is a BIND comment as in common Unix shells and perl

```

WARNING: you cannot use the semicolon character (;) to start a comment.

Options Section

The syntax of the options section is as follows:

```

options {
    [ directory path_name; ]
    [ named-xfer path_name; ]
    [ pid-file path_name; ]
    [ auth-nxdomain yes_or_no; ]
    [ fake-iquery yes_or_no; ]
    [ fetch-glue yes_or_no; ]
    [ multiple-cnames yes_or_no; ]
    [ notify yes_or_no; ]
    [ recursion yes_or_no; ]
    [ forward ( only | first ); ]
    [ forwarders { [ in_addr ; [ in_addr ; ... ] ] }; ]
    [ check-names ( master | slave | response ) ( warn | fail | ignore ); ]
    [ allow-query { address_match_list }; ]
    [ allow-transfer { address_match_list }; ]
    [ listen-on [ port ip_port ] { address_match_list }; ]
    [ query-source [ address ( ip_addr | * ) ] [ port ( ip_port | * ) ] ; ]
    [ max-transfer-time-in number; ]
    [ transfer-format ( one-answer | many-answers ); ]
    [ transfers-in number; ]
}

```

in.named(1M)

```
[ transfers-out number; ]
[ transfers-per-ns number; ]
[ coresize size_spec ; ]
[ datasize size_spec ; ]
[ files size_spec ; ]
[ stacksize size_spec ; ]
[ clean-interval number; ]
[ interface-interval number; ]
[ scan-interval number; ]
[ topology { address_match_list }; ]
};
```

Definitions and Use of Options

The options section sets up global options to be used by BIND. This section may appear at only once in a configuration file; if more than one occurrence is found, the first occurrence determines the actual options used, and a warning will be generated. If there is no options section, an options block with each option set to its default will be used.

Pathnames

directory	The working directory of the server. Any non-absolute pathnames in the configuration file will be taken as relative to this directory. The default location for most server output files (for example, "named.run") is this directory. If a directory is not specified, the working directory defaults to ".", the directory from which the server was started. The directory specified should be an absolute path.
named-xfer	The pathname to the named-xfer program that the server uses for inbound zone transfers. If not specified, the default is operating system dependent, for example, "/usr/sbin/named-xfer").
pid-file	The pathname of the file the server writes its process ID in. If not specified, the default is operating system dependent, but is usually "/var/run/named.pid" or "/etc/named.pid". The pid-file is used by programs like "ndc" that want to send signals to the running nameserver.

Boolean Options

auth-nxdomain	If yes, then the AA bit is always set on NXDOMAIN responses, even if the server is not actually authoritative. The default is yes. Do not turn off auth-nxdomain unless you are sure you know what you are doing, as some older software will not like it.
fake-iquery	If yes, the server will simulate the obsolete DNS query type IQUERY. The default is no.
fetch-glue	If yes (the default), the server will fetch "glue" resource records it does not have when constructing the additional data section of a response. fetch-glue no can be used in conjunction with recursion no to prevent the server's cache from growing or becoming corrupted (at the cost of requiring more work from the client).

in.named(1M)

	multiple-cnames	If yes, then multiple CNAME resource records will be allowed for a domain name. The default is no. Allowing multiple CNAME records is against standards and is not recommended. Multiple CNAME support is available because previous versions of BIND allowed multiple CNAME records, and these records have been used for load balancing by a number of sites.
	notify	If yes (the default), DNS NOTIFY messages are sent when a zone the server is authoritative for changes. The use of NOTIFY speeds convergence between the master and its slaves. Slave servers that receive a NOTIFY message and understand it will contact the master server for the zone and see if they need to do a zone transfer, and if they do, they will initiate it immediately. The notify option may also be specified in the zone section, in which case it overrides the options notify statement.
	recursion	If yes, and a DNS query requests recursion, then the server will attempt to do all the work required to answer the query. If recursion is not on, the server will return a referral to the client if it doesn't know the answer. The default is yes. See also fetch-glue above.
Forwarding		The forwarding facility can be used to create a large sitewide cache on a few servers, reducing traffic over links to external name servers. It can also be used to allow queries by servers that do not have direct access to the Internet, but wish to look up exterior names anyway. Forwarding occurs only on those queries for which the server is not authoritative, and it does not have the answer in its cache.
	forward	This option is only meaningful if the forwarders list is not empty. A value of first, the default, causes the server to query the forwarders first, and if that doesn't answer the question, the server will then look for the answer itself. If only is specified, the server will only query the forwarders.
	forwarders	Specifies the IP addresses to be used for forwarding. The default is the empty list (no forwarding).
		Future versions of BIND 8 will provide a more powerful forwarding system. The syntax described above will continue to be supported.
Name Checking		The server can check domain names based upon their expected client contexts. For example, a domain name used as a hostname can be checked for compliance with the valid hostnames defined in the RFCs. Three checking methods are available:
	ignore	No checking is done.
	warn	Names are checked against their expected client contexts. Invalid names are logged, but processing continues normally.
	fail	Names are checked against their expected client contexts. Invalid names are logged, and the offending data is rejected.

in.named(1M)

The server can check names in three areas: master zone files, slave zone files, and in responses to queries the server has initiated. If `check-names response fail` has been specified, and answering the client's question would require sending an invalid name to the client, the server will send a `REFUSED` response code to the client.

The defaults are:

```
check-names master fail;
check-names slave warn;
check-names response ignore;
```

`check-names` may also be specified in the zone section, in which case it overrides the options `check-names` statement. When used in a zone section, the area is not specified (because it can be deduced from the zone type).

Access Control

Access to the server can be restricted based on the IP address of the requesting system. See `address_match_list` for details on how to specify IP address lists.

`allow-query` Specifies which hosts are allowed to ask ordinary questions. `allow-query` may also be specified in the zone section, in which case it overrides the options `allow-query` statement. If not specified, the default is to allow queries from all hosts.

`allow-transfer` Specifies which hosts are allowed to receive zone transfers from the server. `allow-transfer` may also be specified in the zone section, in which case it overrides the options `allow-transfer` statement. If not specified, the default is to allow transfers from all hosts.

Interfaces

The interfaces and ports that the server will answer queries from may be specified using the `listen-on` option. `listen-on` takes an optional port, and an `address_match_list`. The server will listen on all interfaces allowed by the address match list. If a port is not specified, port 53 will be used.

Multiple `listen-on` statements are allowed. For example,

```
listen-on { 5.6.7.8; };
listen-on port 1234 { !1.2.3.4; 1.2/16; };
```

If no `listen-on` is specified, the server will listen on port 53 on all interfaces.

Query Address

If the server does not know the answer to a question, it will query other name servers. `query-source` specifies the address and port used for such queries. If address is `*` or is omitted, a wildcard IP address (`INADDR_ANY`) will be used. If port is `*` or is omitted, a random unprivileged port will be used. The default is:

```
query-source address * port *;
```

Note: `query-source` currently applies only to UDP queries; TCP queries always use a wildcard IP address and a random unprivileged port.

Zone Transfers	max-transfer-time-in	Inbound zone transfers (named-xfer processes) running longer than this many minutes will be terminated. The default is 120 minutes.
	transfer-format	The server supports two zone transfer methods. <code>one-answer</code> uses one DNS message per resource record transferred. <code>many-answers</code> packs as many resource records as possible into a message. <code>many-answers</code> is more efficient, but is only known to be understood by BIND 8.1 and patched versions of BIND 4.9.5. The default is <code>one-answer</code> . <code>transfer-format</code> may be overridden on a per-server basis by using the server section.
	transfers-in	The maximum number of inbound zone transfers that can be running concurrently. The default value is 10. Increasing <code>transfers-in</code> may speed up the convergence of slave zones, but it also may increase the load on the local system.
	transfers-out	This option will be used in the future to limit the number of concurrent outbound zone transfers. It is checked for syntax, but is otherwise ignored.
	transfers-per-ns	The maximum number of inbound zone transfers (named-xfer processes) that can be concurrently transferring from a given remote name server. The default value is 2. Increasing <code>transfers-per-ns</code> may speed up the convergence of slave zones, but it also may increase the load on the remote name server. <code>transfers-per-ns</code> may be overridden on a per-server basis by using the transfers statement in the server section.
Resource Limits	The server's usage of many system resources can be limited. Some operating systems do not support some of the limits, and a warning will be generated if an unsupported limit is set in the configuration file.	
	Scaled values are allowed when specifying resource limits. For example, 1G can be used instead of 1073741824 to specify a limit of one gigabyte, unlimited requests unlimited use, or the maximum available amount. Default uses the limit that was in force when the server was started. See <code>ulimit(1)</code> for a discussion of <code>ulimit -a</code> (ksh only) for defaults.	
	coresize	The maximum size of a core dump. The default is system dependent.
	datasize	The maximum amount of data memory the server may use. The default is system dependent.

in.named(1M)

	<p>files The maximum number of files that the server may have open concurrently. The default is system dependent.</p> <p>stacksize The maximum amount of stack memory the server may use. The default is system dependent.</p>
Topology	<p>All other things being equal, when the server chooses a name server to query from a list of name servers, it prefers the one that is topologically closest to itself. The topology statement takes an <code>address_match_list</code> and interprets it in a special way. Each top-level list element is assigned a distance. Non-negated elements get a distance based on their position in the list, where the closer the match is to the start of the list, the shorter the distance is between it and the server. A negated match will be assigned the maximum distance from the server. If there is no match, the address will get a distance which is further than any non-negated list element, and closer than any negated element. For example,</p> <pre>topology { 10/8; !1.2.3/24; { 1.2/16; 3/8; }; };</pre> <p>will prefer servers on network 10 the most, followed by hosts on network 1.2.0.0 (netmask 255.255.0.0) and network 3, with the exception of hosts on network 1.2.3 (netmask 255.255.255.0), which is preferred least of all. The default topology is</p> <pre>topology { localhost; localnets; };</pre>
The Server Section	<p>The syntax of the server section is as follows:</p> <pre>server ip_addr { [bogus yes_or_no;] [transfers number;] [transfer-format (one-answer many-answers);] [keys { key_id [key_id ...] };] };</pre> <p>The server statement defines the characteristics to be associated with a remote name server.</p> <p>If you discover that a server is giving out bad data, marking it as bogus will prevent further queries to it. The default value is no.</p> <p>The server supports two zone transfer methods. The first, <code>one-answer</code>, uses one DNS message per resource record transferred. <code>many-answers</code> packs as many resource records as possible into a message. <code>many-answers</code> is more efficient, but is only known to be understood by BIND 8.1 and patched versions of BIND 4.9.5. You can specify which method to use for a server with the <code>transfer-format</code> option. If <code>transfer-format</code> is not specified, the <code>transfer-format</code> specified by the options statement will be used.</p> <p>The <code>transfers</code> will be used in a future release of the server to limit the number of concurrent inbound zone transfers from the specified server. It is checked for syntax but is otherwise ignored.</p>

The Zone Section

The keys statement is intended for future use by the server. It is checked for syntax but is otherwise ignored.

The syntax of the zone section is as follows:

```
zone domain_name [ ( in | hs | hesiod | chaos ) ] {
    type master;
    file path_name;
    [ check-names ( warn | fail | ignore ); ]
    [ allow-update { address_match_list }; ]
    [ allow-query { address_match_list }; ]
    [ allow-transfer { address_match_list }; ]
    [ notify yes_or_no; ]
    [ also-notify { ip_addr; [ ip_addr; ... ] }; ]
};

zone domain_name [ ( in | hs | hesiod | chaos ) ] {
    type ( slave | stub );
    [ file path_name; ]
    masters { ip_addr; [ ip_addr; ... ] };
    [ check-names ( warn | fail | ignore ); ]
    [ allow-update { address_match_list }; ]
    [ allow-query { address_match_list }; ]
    [ allow-transfer { address_match_list }; ]
    [ max-transfer-time-in number; ]
    [ notify yes_or_no; ]
    [ also-notify { ip_addr; [ ip_addr; ... ] }; ]
};

zone . [ ( in | hs | hesiod | chaos ) ] {
    type hint;
    file path_name;
    [ check-names ( warn | fail | ignore ); ]
};
```

Zone types are defined as follows:

- | | |
|--------|--|
| master | The master copy of the data in a zone. |
| slave | A slave zone is a replica of a master zone. The masters list specifies one or more IP addresses that the slave contacts to update its copy of the zone. If file is specified, then the replica will be written to the file. Use of file is recommended, since it often speeds server startup and eliminates a needless waste of bandwidth. |
| stub | A stub zone is like a slave zone, except that it replicates only the NS records of a master zone instead of the entire zone. |
| hint | The initial set of root name servers is specified using a hint zone. When the server starts up, it uses the root hints to find a root name server and get the most recent list of root name servers. |

Note: previous releases of BIND used the term primary for a master zone, secondary for a slave zone, and cache for a hint zone.

in.named(1M)

The zone's name may optionally be followed by a class. If a class is not specified, class in is used.

Zone options are described as follows:

check-names	See Name Checking.
allow-query	See the description of allow-query in the Access Control section.
allow-update	Specifies which hosts are allowed to submit dynamic DNS updates to the server. The default is to deny updates from all hosts.
allow-transfer	See the description of allow-transfer in the Access Control section.
max-transfer-time-in	See the description of max-transfer-time-in in the Zone Transfers section.
notify	See the description of notify in the Boolean Options section.
also-notify	also-notify is only meaningful if notify is active for this zone.

The set of machines that will receive a DNS NOTIFY message for this zone is made up of all the listed name servers for the zone (other than the primary master) plus any IP addresses specified with also-notify. also-notify is not meaningful for stub zones. The default is the empty list.

The Logging Section

The syntax of the logging section is as follows:

```
logging {
    [ channel channel_name {
        ( file path_name
          [ versions ( number | unlimited ) ]
          [ size size_spec ]
        | syslog ( kern | user | mail | daemon | auth | syslog | lpr |
                  news | uucp | cron | authpriv | ftp |
                  local0 | local1 | local2 | local3 |
                  local4 | local5 | local6 | local7 )
        | null );

        [ severity ( critical | error | warning | notice |
                    info | debug [ level ] | dynamic ); ]
        [ print-category yes_or_no; ]
        [ print-severity yes_or_no; ]
        [ print-time yes_or_no; ]
    }; ]

    [ category category_name {
        channel_name; [ channel_name; ... ]
    }; ]
    ...
}
```

in.named(1M)

```
};
```

The logging statement configures a wide variety of logging options for the name server. Its channel phrase associates output methods, format options and severity levels with a name that can then be used with the category phrase to select how various classes of messages are logged.

Only one logging statement is used to define as many channels and categories as are wanted. If there are multiple logging statements in a configuration, the first defined determines the logging, and warnings are issued for the others. If there is no logging statement, the default logging configuration will be:

```
logging {
    category default { default_syslog; default_debug; };
    category panic { default_syslog; default_stderr; };
    category packet { default_debug; };
    category eventlib { default_debug; };
};
```

The Channel Phrase

All log output goes to one or more "channels"; you can make as many of them as you want.

Every channel definition must include a clause that says whether messages selected for the channel go to a file, to a particular syslog facility, or are discarded. It can optionally also limit the message severity level that will be accepted by the channel (default is "info"), and whether to include a named-generated time stamp, the category name and/or severity level (default is not to include any).

The word `null` as the destination option for the channel will cause all messages sent to it to be discarded; other options for the channel are meaningless.

The file clause can include limitations both on how large the file is allowed to become, and how many versions of the file will be saved each time the file is opened.

The size option for files is simply a hard ceiling on log growth. If the file ever exceeds the size, then named will just not write anything more to it until the file is reopened; exceeding the size does not automatically trigger a reopen. The default behavior is to not limit the size of the file.

If you use the version logfile option, then named will retain that many backup versions of the file by renaming them when opening. For example, if you choose to keep 3 old versions of the file "lamers.log" then just before it is opened lamers.log.1 is renamed to lamers.log.2, lamers.log.0 is renamed to lamers.log.1, and lamers.log is renamed to lamers.log.0. No rolled versions are kept by default. The unlimited keyword is synonymous with 99 in current BIND releases.

The argument for the `syslog()` clause is a `syslog()` facility as described in the `syslog(3C)` manual page. How `syslogd(1M)` will handle messages sent to this facility is described in the `syslog.conf(4)` manual page. If you have a system which

in.named(1M)

uses a very old version of `syslog()` that only uses two arguments to the `openlog()` function, then this clause is silently ignored.

The severity clause works like the "priorities" to `syslog()`, except that they can also be used if you are writing straight to a file rather than using `syslog()`. Messages which are not at least of the severity level given will not be selected for the channel; messages of higher severity levels will be accepted.

If you are using `syslog()`, then the `syslog.conf` priorities will also determine what eventually passes through. For example, defining a channel facility and severity as `daemon` and `debug` but only logging `daemon.warning` by way of `syslog.conf` will cause messages of severity `info` and `notice` to be dropped. If the situation were reversed, with `named` writing messages of only `warning` or higher, then `syslogd` would print all messages it received from the channel.

The server can supply extensive debugging information when it is in debugging mode. If the server's global debug level is greater than zero, then debugging mode will be active. The global debug level is set either by starting the server with the `-d` option followed by a positive integer, or by sending the server the `SIGUSR1` signal (for example, by using `"ndc trace"`). The global debug level can be set to zero, and debugging mode turned off, by sending the server the `SIGUSR2` signal (`"ndc notrace"`). All debugging messages in the server have a debug level, and higher debug levels give more detailed output. Channels that specify a specific debug severity, for example:

```
channel specific_debug_level {
    file "foo";
    severity debug 3;
};
```

will get debugging output of level 3 or less any time the server is in debugging mode, regardless of the global debugging level. Channels with dynamic severity use the server's global level to determine what messages to print.

If `print-time` has been turned on, then the date and time will be logged. `print-time` may be specified for a `syslog()` channel, but is usually pointless since `syslog()` also prints the date and time. If `print-category` is requested, then the category of the message will be logged as well. Finally, if `print-severity` is on, then the severity level of the message will be logged. The `print-options` may be used in any combination, and will always be printed in the following order: `time`, `category`, `severity`. Here is an example where all three `print-options` are on:

```
28-Apr-1997 15:05:32.863 default: notice: Ready to answer queries.
```

There are four predefined channels that are used for default logging for `in.named` as follows. How they are used is described in the next section.

```
channel default_syslog {
    syslog daemon;      # send to syslog's daemon facility
    severity info;      # only send priority info and higher
};
```

```

channel default_debug {
    file "named.run";      # write to named.run in the working directory
    severity dynamic;      # log at the server's current debug level
};

channel default_stderr { # writes to stderr
    file "<stderr>";      # this is illustrative only;
    # there's currently   # no way of specifying an internal file
                        # descriptor in the configuration language.
    severity info;        # only send priority info and higher
};

channel null {
    null;                  # toss anything sent to this channel
};

```

Once a channel is defined, it cannot be redefined. Thus you cannot alter the built-in channels directly, but you can modify the default logging by pointing categories at channels you have defined.

The Category Phase

There are many categories, so you can send the logs you want to see wherever you want, without seeing logs you do not want. If you do not specify a list of channels for a category, then log messages in that category will be sent to the default category instead. If do not specify a default category, the following "default default" is used:

```
category default { default_syslog; default_debug; };
```

For example, if you want to log security events to a file, but you also want keep the default logging behavior, specify the following:

```

channel my_security_channel {
    file "my_security_file";
    severity info;
};
category security { my_security_channel; default_syslog; default_debug; };

```

To discard all messages in a category, specify the null channel:

```

category lame-servers { null; };
category cname { null; };

```

The following categories are available:

default	The catch-all. Many things still are not classified into categories, and they all end up here. Also, if you do not specify any channels for a category, the default category is used instead. If you do not define the default category, the following definition is used:
---------	--

```
category default { default_syslog; default_debug; };
```

in.named(1M)

config	High-level configuration file processing.
parser	Low-level configuration file processing.
queries	A short log message is generated for every query the server receives.
lame-servers	Messages like "Lame server on ..."
statistics	Statistics.
panic	<p>If the server has to shut itself down due to an internal problem, it will log the problem in this category as well as in the problem's native category. If you do not define the panic category, the following definition is used:</p> <pre>category panic { default_syslog; default_stderr; };</pre>
update	Dynamic updates.
ncache	Negative caching.
xfer-in	Zone transfers the server is receiving.
xfer-out	Zone transfers the server is sending.
db	All database operations.
eventlib	<p>Debugging info from the event system. Only one channel may be specified for this category, and it must be a file channel. If you do not define the eventlib category, the following definition is used:</p> <pre>category eventlib { default_debug; };</pre>
packet	<p>Dumps of packets received and sent. Only one channel may be specified for this category, and it must be a file channel. If you do not define the packet category, the following definition is used:</p> <pre>category packet { default_debug; };</pre>
notify	The NOTIFY protocol.
cname	Messages like "... points to a CNAME".
security	Approved/unapproved requests.
os	Operating system problems.
insist	Internal consistency check failures.
maintenance	Periodic maintenance events.
load	Zone loading messages.

response-checks Messages arising from response checking, such as "Malformed response ...", "wrong ans. name ...", "unrelated additional info ...", "invalid RR type ...", and "bad referral ...".

The Key Section

The syntax of the key section is as follows:

```
key key_id {
    algorithm algorithm_id;
    secret secret_string;
};
```

The key section defines a key ID which can be used in a server section to associate an authentication method with a particular name server.

A key ID must be created with the key statement before it can be used in a server definition.

The `algorithm_id` is a string that specifies a security/authentication algorithm. `secret_string` is the secret to be used by the algorithm.

The key statement is intended for future use by the server. It is checked for syntax but is otherwise ignored.

The Include Section

The syntax of the include section is as follows:

```
include path_name;
```

The include statement inserts the specified file at the point that the include statement is encountered. It cannot be used within another statement, though, so a line such as `acl internal_hosts { "include internal_hosts.acl" }` is not allowed. Use include to break the configuration up into easily-managed chunks. For example:

```
include "/etc/security/keys.bind";
include "/etc/acls.bind";
```

could be used at the top of a BIND configuration file in order to include any ACL or key information.

Be careful not to type `#include`, like you would in a C program, because `#` is used to start a comment.

The ACL Format

The syntax of the ACL section is as follows:

```
acl name {
    address_match_list
};
```

The `acl` statement creates a named address match list. It gets its name from a primary use of address match lists: Access Control Lists (ACLs).

in.named(1M)

Note that an address match list's name must be defined with `acl` before it can be used elsewhere; no forward references are allowed.

The following ACLs are built-in:

<code>any</code>	Allows all hosts.
<code>none</code>	Denies all hosts.
<code>localhost</code>	Allows the IP addresses of all interfaces on the system.
<code>localnets</code>	Allows any host on a network for which the system has an interface.

Zone File Format

The zone files are also known as the authoritative master files (data files) for a zone. In the boot file, references were made to these files as part of the specification of any primary directives.

Two classes of entries populate the zone files, directives and resource records. The start of the zone file is likely to contain one or two directives that establish a context that modifies the way subsequent records are interpreted.

Resource records for a zone determine how a zone is managed by establishing zone characteristics. For example, one type of zone record establishes the zone's mailbox information.

The very first record of each zone file should be a Start-of-Authority record (SOA) for a zone. A multiple-line SOA record is presented below. The meaning of the values in this sample will become clearer with the help of a list that describes the purpose of each field in the zone record (see the SOA list subitem under the `rr-type` list item in, *Format of Resource Records in Zone Files*).

```
@ IN SOA ucbvax.Berkeley.EDU. rwh.ucbvax.Berkeley.EDU. (
1989020501 ;serial
10800      ;refresh
3600       ;retry
3600000    ;expire
86400 )    ;minimum
```

Resource records normally end at the end of a line, but may be continued across lines between opening and closing parentheses (as demonstrated by the preceding sample).

Comments are introduced by semicolons. They continue to the end of the line.

Directives in Zone Files

There are two control directives that help determine how the zone file is processed, `$INCLUDE` and `$ORIGIN`.

The `$INCLUDE` directive refers to still another file within which zone characteristics are described. Such files typically contain groups of resource records, but they may also contain further directives.

The \$ORIGIN directive establishes a current origin that is appended to any domain values that do not end with a '.' (dot). The placeholder domain represents the first resource record field as shown in Format of Resource Records in Zone Files. The format for these directives is:

```
$INCLUDE filename opt-current-domain
$ORIGIN current-domain
```

where:

current-domain	Specifies the value of the current origin that remains in effect for this configuration file unless a subsequent \$ORIGIN directive overrides it for the remaining portion of the file.
filename	Specifies a file, the contents of which are, in effect, incorporated into the configuration file at the location of the corresponding \$INCLUDE directive.
opt-current-domain	Optionally defines a current origin that is applicable only to the records residing in the specified file in the corresponding \$INCLUDE directive. This directive overrides the origin given in a preceding \$ORIGIN directive, but only for the scope of the included text. See also current-domain. Neither the opt-current-domain argument of \$INCLUDE nor the \$ORIGIN directive in the included file can affect the current origin in effect for the remaining records in the main configuration file (as defined by those \$ORIGIN directives that reside there).

Format of Resource Records in Zone Files

The format of the resource records is:

```
domain opt-ttl opt-class rr-type rr-data...where:
```

domain	<p>Specifies the domain being described by the current line and any following lines that lack a value for this field. Beware of any domain values that you enter without full qualification, because the value of the current origin will be appended to them. The value of the current origin is appended when domain does not end with a dot.</p> <p>A domain value specified as the symbol @ is replaced with the value of the current origin. The current-domain or any locally-overriding opt-current-domain value is used as its replacement. (For a discussion of these placeholders, see the earlier discussion of the \$ORIGIN and \$INCLUDE directives.)</p> <p>A domain value specified as a '.' (dot) represents the root.</p>
--------	--

in.named(1M)

opt-ttl	Specifies the number of seconds corresponding to the <code>time-to-live</code> value applicable to the zone characteristic that is defined in the remaining fields. This field is optional. It defaults to zero. Zero is interpreted as the minimum value specified in the SOA record for the zone.
opt-class	Specifies the object address type; currently only one type is supported, <code>IN</code> , for objects connected to the Internet.
rr-type rr-data ...	Specifies values that describe a zone characteristic. Permissible <code>rr-type</code> and other field values are listed below. The field values are listed in the order that they must appear.
A address	
Specifies the host address (in <code>dotted-quad</code> format). DCE or AFS server.	
CNAME canonical-name	
Specifies in a <code>domain-name</code> format the canonical name for the alias (domain).	
HINFO cpu-type OS-type	
Host information supplied in terms of a CPU type and an OS type.	
MX preference mail-exchanger	
Specifies in <code>domain-name</code> format a mail exchanger preceded by a preference value (between 0 and 32767), with lower numeric values representing higher logical preferences.	
NS authoritative-server	
Specifies in <code>domain-name</code> format an authoritative name server.	
NULL	
Specifies a null zone record.	
PTR domain-pointer	
Specifies in <code>domain-name</code> format a domain name pointer.	
RP mailbox txt-referral	
Offers details about how to reach a responsible person for the domain name.	
<code>retry expire ttl</code>	
SOA host-domain maintainer-addr serial- no refresh	
Establishes the start of a zone of authority in terms of the domain of the originating host (<code>host-domain</code>), the domain address of the maintainer (<code>maintainer-addr</code>), a serial number (<code>serial-no</code>), the refresh period in seconds (<code>refresh</code>), the retry	

in.named(1M)

period in seconds (retry), the expiration period in seconds (expire), and the minimum time-to-live period in seconds (ttl). See RFC 1035.

The serial number should be changed each time the master file is changed. Secondary servers check the serial number at intervals specified by the refresh time in seconds; if the serial number changes, a zone transfer will be done to load the new data.

If a master server cannot be contacted when a refresh is due, the retry time specifies the interval at which refreshes should be attempted. If a master server cannot be contacted within the interval given by the expire time, all data from the zone is discarded by secondary servers. The minimum value is the time-to-live used by records in the file with no explicit time-to-live value.

The serial number can be given as a dotted number. However, this is a very unwise thing to do, since the translation to normal integers is via concatenation rather than multiplication and addition. You could spell out the year, month, day of month, and 0..99 version number and still fit it inside the unsigned 32-bit size of this field. This strategy should work for the foreseeable future (but is questionable after the year 4293).

For more detailed information, see RFC 883.

rr-data ... See the description of rr-type.

Consult *Name Server Operations Guide for BIND* for further information about the supported types of resource records.

EXIT STATUS The in.named process returns the following exit values:

- | | |
|---|------------------------|
| 0 | Successful completion. |
| 1 | An error occurred. |

FILES In the Trusted Solaris environment, these files have a sensitivity label of ADMIN_LOW:

- | | |
|-----------------------|--------------------------------------|
| /etc/named.conf | Name server configuration boot file. |
| /etc/named.pid | The process ID (on older systems). |
| /var/tmp/named.run | Debug output. |
| /var/tmp/named.stats | Nameserver statistics data. |
| /var/tmp/nameddump.db | Dump of the name servers database. |
| /var/tmp/named.pid | The process ID (on newer systems). |

in.named(1M)

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**

in.named accepts requests at any sensitivity label and replies at the sensitivity label of the client's request. in.named can serve DNS clients and can communicate with other DNS servers that are on Trusted Solaris hosts or non-trusted hosts.

Files used by in.named should be protected from unauthorized access by having the sensitivity label ADMIN_LOW.

Invoking in.named requires the trusted path attribute, an effective UID of 0, a process sensitivity label of ADMIN_LOW, and the following privileges: net_mac_read, net_privaddr, net_upgrade_sl, proc_setclr, sys_trans_label, sys_net_config, and sys_config.

**Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual**

fork(2), resolver(3RESOLV), listen(3SOCKET), resolv.conf(4)

kill(1), named-xfer(1M), syslogd(1M), signal(3C), syslog(3C),
getservbyname(3SOCKET), syslog.conf(4), attributes(5)

Braden, R. (Editor), *Requirements for Internet Hosts - Applications and Support*, RFC 1123, Internet Engineering Task Force - Network Working Group, October 1989.

Mockapetris, Paul, *Domain Names - Concepts and Facilities*, RFC 1034, , Network Information Center, SRI International, Menlo Park, Calif., November 1987.

Mockapetris, Paul, *Domain Names - Implementation and Specification*, RFC 1035, Network Information Center, SRI International, Menlo Park, Calif., November 1987.

Mockapetris, Paul, *Domain System Changes and Observations*, RFC 973, Network Information Center, SRI International, Menlo Park, Calif., January 1986.

Partridge, Craig, *Mail Routing and the Domain System*, RFC 974, Network Information Center, SRI International, Menlo Park, Calif., January 1986.

Vixie, Paul, Dunlap, Keven J., Karels, Michael J., *Name Server Operations Guide for BIND* (public domain), Internet Software Consortium, 1995.

NOTES

The following signals have the specified effect when sent to the server process using the kill(1) command:

SIGHUP	Causes the server to read /etc/named.conf and reload the database.
SIGHUP	Also causes the server to check the serial number on all secondary zones. Normally the serial numbers are only checked at the

in.named(1M)

	intervals specified by the SOA record at the start of each zones-definition file.
SIGINT	Dumps the current database and cache to <code>/var/tmp/nameddump.db</code> .
SIGIOT	Dumps statistical data into <code>/var/tmp/named.stats</code> . Statistical data are appended to the file.
SIGUSR1	Turns on debugging at the lowest level when received the first time; receipt of each additional SIGUSR1 signal causes the server to increment the debug level.
SIGUSR2	Turns off debugging completely.
SIGWINCH	Toggles logging of all incoming queries through the syslog system daemon. See <code>syslogd(1M)</code> .

in.rarpd(1M)

NAME	in.rarpd, rarpd – DARPA Reverse Address Resolution Protocol server
SYNOPSIS	<pre>/usr/sbin/in.rarpd [-d] -a</pre> <pre>/usr/sbin/in.rarpd [-d] <i>device unit</i></pre>
DESCRIPTION	<p><code>in.rarpd</code> starts a daemon that responds to Reverse Address Resolution Protocol (RARP) requests. The daemon forks a copy of itself that runs in background. It must be started from the trusted path, with a UID of 0 and the label <code>ADMIN_LOW</code>. To succeed, it must inherit the <code>sys_net_conf</code> and <code>net_broadcast</code> privileges.</p> <p>RARP is used by machines at boot time to discover their Internet Protocol (IP) address. The booting machine provides its Ethernet address in a RARP request message. Using the <code>ethers</code> and <code>hosts</code> databases, <code>in.rarpd</code> maps this Ethernet address into the corresponding IP address which it returns to the booting machine in an RARP reply message. The booting machine must be listed in both databases for <code>in.rarpd</code> to locate its IP address. <code>in.rarpd</code> issues no reply when it fails to locate an IP address.</p> <p><code>in.rarpd</code> uses the STREAMS-based Data Link Provider Interface (DLPI) message set to communicate directly with the datalink device driver.</p>
OPTIONS	<p>The following options are supported:</p> <ul style="list-style-type: none"> -a Get the list of available network interfaces from IP using the <code>SIOCGIFADDR</code> ioctl and start a RARP daemon process on each interface returned. -d Print assorted debugging messages while executing.
EXAMPLES	<p>EXAMPLE 1 Starting an <code>in.rarpd</code> daemon for each network interface name returned from <code>/dev/ip</code>:</p> <p>The following command starts an <code>in.rarpd</code> for each network interface name returned from <code>/dev/ip</code>:</p> <pre>example# /usr/sbin/in.rarpd -a</pre> <p>EXAMPLE 2 Starting an <code>in.rarpd</code> daemon on the device <code>/dev/le</code> with the device instance number 0</p> <p>The following command starts one <code>in.rarpd</code> on the device <code>/dev/le</code> with the device instance number 0.</p> <pre>example# /usr/sbin/in.rarpd le 0</pre>
SUMMARY OF TRUSTED SOLARIS CHANGES	<p><code>in.rarpd</code> should be started from the trusted path with a UID0 and sensitivity label of <code>ADMIN_LOW</code>. It must inherit the <code>sys_net_config</code> and <code>net_broadcast</code> privileges.</p>
FILES	<pre>/etc/ethers File or other source, as specified by nsswitch.conf(4).</pre> <pre>/etc/hosts File or other source, as specified by nsswitch.conf(4).</pre>

in.rarpd(1M)

/tftpboot Directory for remote boot scripts.
/dev/ip List of available network interfaces.
/dev/arp Address resolution protocol list.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

**Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual**

ifconfig(1M), nsswitch.conf(4)
boot(1M), ethers(4), hosts(4), netconfig(4), attributes(5), dlpi(7P)
RFC-903, *A Reverse Address Resolution Protocol*, Network Information Center, SRI
International.
Unix International, *Data Link Provider Interface*, Version 2, May 7, 1991, Sun
Microsystems, 800-6915-01.

in.rdisc(1M)

NAME	in.rdisc, rdisc – Network router discovery daemon
SYNOPSIS	<pre>/usr/sbin/in.rdisc [-a] [-f] [-s] [send-address] [receive-address] /usr/sbin/in.rdisc -r [-p preference] [-T interval] [send-address] [receive-address]</pre>
DESCRIPTION	<p><code>in.rdisc</code> implements the ICMP router discovery protocol. The first form of the command is used on hosts and the second form is used on routers. On a host, <code>in.rdisc</code> is invoked at boot time to populate the network routing tables with default routes. On a router, it is also invoked at boot time in order to start advertising the router to all the hosts.</p>
Host (First Form)	<p>On a host, <code>in.rdisc</code> listens on the <code>ALL_HOSTS</code> (224.0.0.1) multicast address for <code>ROUTER_ADVERTISE</code> messages from routers. The received messages are handled by first ignoring those listed router addresses with which the host does not share a network. Among the remaining addresses, the ones with the highest preference are selected as default routers and a default route is entered in the kernel routing table for each one of them.</p> <p>Optionally, <code>in.rdisc</code> can avoid waiting for routers to announce themselves by sending out a few <code>ROUTER_SOLICITATION</code> messages to the <code>ALL_ROUTERS</code> (224.0.0.2) multicast address when it is started.</p> <p>A timer is associated with each router address. The address will no longer be considered for inclusion in the routing tables if the timer expires before a new <i>advertise</i> message is received from the router. The address will also be excluded from consideration if the host receives an <i>advertise</i> message with the preference being maximally negative.</p>
Router (Second Form)	<p>When <code>in.rdisc</code> is started on a router, it uses the <code>SIOCGIFCONF</code> <code>ioctl</code>(2) to find the interfaces configured into the system and it starts listening on the <code>ALL_ROUTERS</code> multicast address on all the interfaces that support multicast. It sends out <i>advertise</i> messages to the <code>ALL_HOSTS</code> multicast address advertising all its IP addresses. A few initial <i>advertise</i> messages are sent out during the first 30 seconds and after that it will transmit <i>advertise</i> messages approximately every 600 seconds.</p> <p>When <code>in.rdisc</code> receives a <i>solicitation</i> message, it sends an <i>advertise</i> message to the host that sent the <i>solicitation</i> message.</p> <p>When <code>in.rdisc</code> is terminated by a signal, it sends out an <i>advertise</i> message with the preference being maximally negative.</p>
OPTIONS	<p>-a Accept all routers independent of the preference they have in their <i>advertise</i> messages. Normally, <code>in.rdisc</code> only accepts (and enters in the kernel routing tables) the router or routers with the highest preference.</p> <p>-f Run <code>in.rdisc</code> forever even if no routers are found. Normally, <code>in.rdisc</code> gives up if it has not received any <i>advertise</i> message</p>

in.rdisc(1M)

after soliciting three times, in which case it exits with a non-zero exit code. If -f is not specified in the first form then -s must be specified.

-r

Act as a router, rather than a host.

-s

Send three *solicitation* messages initially to quickly discover the routers when the system is booted. When -s is specified, in.rdisc exits with a non-zero exit code if it can not find any routers. This can be overridden with the -f option.

-p preference

Set the preference transmitted in the *solicitation* messages. The default is zero.

-T interval

Set the interval between transmitting the *advertise* messages. The default time is 600 seconds.

SUMMARY OF TRUSTED SOLARIS CHANGES

in.rdisc must be started from the trusted path. To modify kernel routing tables, it must inherit the sys_net_config privilege. To open a raw socket, it needs the net_rawaccess privilege. To send multicast or broadcast packets, it needs the net_broadcast privilege.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8 4/01 Reference Manual

in.routed(1M)

ioctl(2), attributes(5), icmp(7P), inet(7P)

Deering, S.E., editor, *ICMP Router Discovery Messages*, RFC 1256, Network Information Center, SRI International, Menlo Park, California, September 1991.

in.rexecd(1M)

NAME	in.rexecd, rexecd – Remote execution server
SYNOPSIS	in.rexecd
DESCRIPTION	<p>in.rexecd is the server for the <code>rexec(3SOCKET)</code> routine. The server provides remote execution facilities with authentication based on user names and passwords. It is invoked automatically as needed by <code>inetd(1M)</code>, and then executes the following protocol:</p> <ol style="list-style-type: none">1) The server reads characters from the socket up to a null (<code>\0</code>) byte. The resultant string is interpreted as an ASCII number, base 10.2) If the number received in step 1 is non-zero, it is interpreted as the port number of a secondary stream to be used for the <code>stderr</code>. A second connection is then created to the specified port on the client's machine.3) A null terminated user name of at most 16 characters is retrieved on the initial socket.4) A null terminated password of at most 16 characters is retrieved on the initial socket.5) A null terminated command to be passed to a shell is retrieved on the initial socket. The length of the command is limited by the upper bound on the size of the system's argument list.6) <code>rexecd</code> then validates the user as is done at login time and, if the authentication was successful, changes to the user's home directory, and establishes the user and group protections of the user. Access is denied unless the user has the remote login authorization. If the <code>/etc/nologin</code> file exists, access is denied. If any of these steps fail the connection is aborted and a diagnostic message is returned.7) A null byte is returned on the connection associated with the <code>stderr</code> and the command line is passed to the normal login shell of the user. The shell inherits the network connections established by <code>rexecd</code>.
USAGE	in.rexecd and rexecd are IPv6-enabled. See <code>ip6(7P)</code> .
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES Trusted Solaris 8 4/01 Reference Manual	<p>Login is not allowed unless the user has the <code>remote login</code> authorization. If the <code>/etc/nologin</code> file exists, the user is not allowed to log in.</p> <p><code>inetd(1M)</code>, <code>inetd.conf(4)</code></p> <p><code>rexec(3SOCKET)</code>, <code>attributes(5)</code>, <code>ip6(7P)</code></p>
--	--

DIAGNOSTICS All diagnostic messages are returned on the connection associated with the `stderr`, after which any network connections are closed. An error is indicated by a leading byte with a value of 1 (0 is returned in step 7 above upon successful completion of all the steps prior to the command execution).

<code>username too long</code>	The name is longer than 16 characters.
<code>password too long</code>	The password is longer than 16 characters.
<code>command too long</code>	The command line passed exceeds the size of the argument list (as configured into the system).
<code>Login incorrect</code>	No password file entry for the user name existed.
<code>Password incorrect</code>	The wrong password was supplied.
<code>No remote directory</code>	The <code>chdir</code> command to the home directory failed.
<code>Try again.</code>	A fork by the server failed.
<code>/usr/bin/sh: ...</code>	The user's login shell could not be started.

in.rlogind(1M)

NAME	in.rlogind, rlogind – Remote login server
SYNOPSIS	/usr/sbin/in.rlogind -U -T
DESCRIPTION	<p>in.rlogind is the server for the rlogin(1) program. The server provides a remote login facility with authentication based on privileged port numbers.</p> <p>in.rlogind is invoked by inetd(1M) when a remote login connection is established, and executes the following protocol:</p> <ul style="list-style-type: none">■ The server checks the client's source port. If the port is not in the range 0-1023, the server aborts the connection.■ The server checks the client's source address. If an entry for the client exists in both /etc/hosts and /etc/hosts.equiv, a user logging in from the client is not prompted for a password. If the address is associated with a host for which no corresponding entry exists in /etc/hosts, the user is prompted for a password, regardless of whether an entry for the client is present in /etc/hosts.equiv. See hosts(4) and hosts.equiv(4). <p>Once the source port and address have been checked, in.rlogind allocates a pseudo-terminal and manipulates file descriptors so that the slave half of the pseudo-terminal becomes the stdin, stdout, and stderr for a login process. The login process is an instance of the login(1) program, invoked with the -r.</p> <p>The login process then proceeds with the in.rshd(1M) authentication process.</p> <p>The parent of the login process manipulates the master side of the pseudo-terminal, operating as an intermediary between the login process and the client instance of the rlogin program. In normal operation, a packet protocol is invoked to provide Ctrl-S and Ctrl-Q type facilities and propagate interrupt signals to the remote programs. The login process propagates the client terminal's baud rate and terminal type, as found in the environment variable, TERM; see environ(4).</p> <p>The -U option is used to pass the UID of the client to login(1). The -T option is used if the client has the trusted path attribute.</p>

USAGE rlogind and in.rlogind are IPv6-enabled. See ip6(7P).

SUMMARY OF TRUSTED ATTRIBUTES CHANGES Two new options (-U and -T) are used in the call to login(1).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8 4/01 Reference Manual login(1), in.rshd(1M), inetd(1M), inetd.conf(4)

Trusted Solaris 8 4/01 Reference Manual rlogin(1), environ(4), hosts(4), hosts.equiv(4), attributes(5), ip6(7P)

in.rlogind(1M)

DIAGNOSTICS	<p>All diagnostic messages are returned on the connection associated with the <code>stderr</code>, after which any network connections are closed. An error is indicated by a leading byte with a value of 1.</p> <p>Hostname for your address unknown. No entry in the host name database existed for the client's machine.</p> <p>Try again. A <i>fork</i> by the server failed.</p> <p>/usr/bin/sh: ... The user's login shell could not be started.</p>
NOTES	<p>The authentication procedure used here assumes the integrity of each client machine and the connecting medium. This is insecure, but it is useful in an "open" environment.</p> <p>A facility to allow all data exchanges to be encrypted should be present.</p>

in.routed(1M)

NAME	in.routed, routed – Network routing daemon
SYNOPSIS	/usr/sbin/in.routed [-s] [-q] [-t] [-g] [-S] [-v] [<i>logfile</i>]
DESCRIPTION	<p><code>in.routed</code> is invoked at boot time to manage the network routing tables. The routing daemon uses a variant of the Xerox NS Routing Information Protocol in maintaining up-to-date kernel routing table entries.</p> <p>In normal operation, <code>in.routed</code> listens on udp(7P) socket 520 (decimal) for routing information packets. If the host is an internetwork router, it periodically supplies copies of its routing tables to any directly connected hosts and networks.</p> <p>When <code>in.routed</code> is started, it uses the <code>SIOCGIFCONF</code> <code>ioctl(2)</code> to find those directly connected interfaces configured into the system and marked “up” (the software loopback interface is ignored). If multiple interfaces are present, it is assumed the host will forward packets between networks. <code>in.routed</code> then transmits a <i>request</i> packet on each interface (using a broadcast packet if the interface supports it) and enters a loop, listening for <i>request</i> and <i>response</i> packets from other hosts.</p> <p>For trusted routing, extended security attributes must be associated with a route along with the simple metric that indicates the number of hops to the destination. The additional security routing information (SRI) includes a sensitivity label range, and can include a CIPSO domain of interpretation, a RIPSO label, and a RIPSO error, and some additional keywords: <code>ripso_only</code>, <code>cipso_only</code>, and <code>msix_only</code>. The SRI combined with the simple metric is called the extended metric, or <code>emetric</code>.</p> <p>For Trusted Solaris 7 and later systems, two additional types of packets are exchanged. The first one is <i>sec_response</i>, which is like the <i>response</i> packet but also carries the SRI for the routes. Similar to the <i>response</i> packet, the <i>sec_response</i> packet propagates a route while adjusting its metric and SRI one hop at a time. The SRI that is carried in <i>sec_response</i> packets cannot be propagated through non-Trusted Solaris gateways.</p> <p>The second additional packet type is <i>sec_t_response</i>, which has the exact format as <i>sec_response</i> but with a different command number. The <i>sec_t_response</i> packets are used for tunneling. Every time a <i>response</i> is sent, a <i>sec_response</i> and a <i>sec_t_response</i> packet are also sent.</p> <p>Tunneling can be set up for trusted routing between Trusted Solaris 7 and later gateways when non-Trusted Solaris gateways exist between the Trusted Solaris 7 and later gateways. For tunneling to work, all Trusted Solaris gateways must be running Trusted Solaris 2.5.1 or 7 or 8, and they must be using the extended <code>in.routed(1M)</code> for dynamic routing. Also, the non-Trusted Solaris gateways must be using the standard <code>in.routed(1M)</code> for dynamic routing. All gateways must be in the same Intranet. To forward SRIs through non-Trusted Solaris gateways to a target (sub)network, a Trusted Solaris system sends an unlabeled <i>sec_t_response</i> packet in a (sub)network directed broadcast to the target (sub)network on behalf of the non-Trusted Solaris gateway connected to that (sub)network. Trusted Solaris systems on the (sub)network can use the SRI to configure their routing tables correctly, and Trusted Solaris 7 gateways on that (sub)network can propagate the SRI to other</p>

in.routed(1M)

(sub)networks. A machine that does tunneling is called the forwarding machine; any Trusted Solaris gateway can be a forwarding machine.

Tunneling is enabled by the existence of the file `/etc/security/tsol/tunnel`, and the target (sub)network addresses are obtained from this file. A Trusted Solaris gateway can be responsible for tunneling to more than one (sub)network. The file is composed of a series of lines, each in the following format:

broadcast_addr

A Trusted Solaris gateway can be responsible for tunneling to more than one (sub)network.

A Trusted Solaris system ignores a *response* packet if it is sent by another Trusted Solaris gateway, because in this case, *sec_response* packets should be used in place of *response* packets. A Trusted Solaris system processes a *response* packet if it is sent by a non-Trusted Solaris gateway. If tunneling is done on behalf of that non-Trusted Solaris gateway, it will process both the *response* packets sent by the non-Trusted Solaris gateway and the *sec_response* packets sent by a remote Trusted Solaris gateway on behalf of the non-Trusted Solaris gateway.

When a *request* packet is received, `in.routed` formulates a reply based on the information maintained in its internal tables. The *response* packet contains a list of known routes, each marked with a “hop count” metric (a count of 16, or greater, is considered “infinite”). The metric associated with each route returned, provides a metric relative to the sender.

sec_response and *sec_t_response* packets are formulated by ANDing the emetric of the route with the emetric derived from the outgoing interface. Before the *response* packet is sent, a *sec_response* and a *sec_t_response* packet are sent to the same destination with the same metric and additional SRI.

response, *sec_response*, and *request* packets received by `in.routed` are used to update the routing tables if one of the following conditions is satisfied:

- No routing table entry exists for the destination network or host, and the metric indicates the destination is “reachable” (that is, the hop count is not infinite). For *sec_response* and *sec-t_response* packets, a destination is also unreachable if its SRI restricts all possible packets.
- The source host of the packet is the same as the router in the existing routing table entry. That is, updated information is being received from the very internetwork router through which packets for the destination are being routed. The only exception occurs when `in.routed` is supposed to process both the *response* packet from a non-Trusted Solaris gateway and the *sec_response* packet tunneled on behalf of that non-Trusted Solaris gateway. In this situation, if both packets carry routing information for the same route, the SRI from the tunneled *sec_response* packet and the metric from the *response* packet are used.

`in.routed(1M)`

- The existing entry in the routing table has not been updated for some time (defined to be 90 seconds) and the route is at least as cost effective as the current route.
- The new route describes a shorter route to the destination than the one currently stored in the routing tables; the metric of the new route is compared against the one stored in the table to decide this.

For *sec_response* and *sec_t_response* packets, the last rule above is changed to compare the SRIs as well as the metrics. One route is better than another if (a) its metric is smaller; and (b) its SRI is more relaxed than or equal to that of the other. Note that when comparing the SRIs of two routes, one route cannot always serve as a substitute for the other. For example, if the SRIs of two routes have different sensitivity labels, one SRI cannot be said to be more restrictive, because they restrict different sensitivity label ranges.

If two routes cannot be compared, both routes are kept in the routing table, because they represent two routes to the same destination although with different security characteristics; and both routes are needed.

When an update is applied, `in.routed` records the change in its internal tables and generates a *sec_response* packet and a *response* packet to all directly connected hosts and networks. `in.routed` waits a short period of time (no more than 30 seconds) before modifying the kernel's routing tables to allow possible unstable situations to settle.

In addition to processing incoming packets, `in.routed` also periodically checks the routing table entries. If an entry has not been updated for 3 minutes, the entry's metric is set to infinity and marked for deletion. Deletions are delayed an additional 60 seconds to insure the invalidation is propagated throughout the internet.

Hosts acting as internetwork routers gratuitously supply their routing tables every 30 seconds to all directly connected hosts and networks.

In addition to the facilities described above, `in.routed` supports the notion of "distant" passive and active gateways. When `in.routed` is started up, it reads the file `gateways` to find gateways which may not be identified using the `SIOCGIFCONF` ioctl. Gateways specified in this manner should be marked passive if they are not expected to exchange routing information, while gateways marked active should be willing to exchange routing information (that is, they should have a `in.routed` process running on the machine). Passive gateways are maintained in the routing tables forever. Information regarding their existence is not included in any routing information transmitted. Active gateways are treated equally to network interfaces. Routing information is distributed to the gateway and if no routing information is received for a period of time, the associated route is deleted.

The gateways is comprised of a series of lines, each in the following format:

```
< net | host> filename1 gateway filename2 metric value < passive | active >
```

in.routed(1M)

The `net` or `host` keyword indicates if the route is to a network or specific host.

filename1 is the name of the destination network or host. This may be a symbolic name located in `networks` or `hosts`, or an Internet address specified in “dot” notation; see `inet(3SOCKET)`.

filename2 is the name or address of the gateway to which messages should be forwarded.

value is a metric indicating the hop count to the destination host or network.

The keyword `passive` or `active` indicates if the gateway should be treated as passive or active (as described above).

For both the passive and active gateway, the SRIs of their routes are obtained initially from their remote host template. For an active gateway, further routing information will be exchanged with this machine. If later a `sec_response` packet is received from the active gateway or a `sec_t_response` tunneled on its behalf is received, the initial SRI will be updated. If no `sec_response` packet is ever received for this active gateway, use of the initial SRI is continued. For a passive gateway, no further routing information will be exchanged; therefore, the initial SRI is continuously used.

`in.routed` must be started from the Trusted path at `ADMIN_HIGH`. It must inherit the `net_mac_read`, `net_privaddr`, `net_broadcast`, and `sys_net_config` privileges. If a log file is specified, `in.routed` must also inherit the `file_mac_write` privilege.

OPTIONS

- g Is used on internetwork routers to offer a route to the “default” destination. This is typically used on a gateway to the Internet, or on a gateway that uses another routing protocol whose routes are not reported to other local routers.
- q Is the opposite of the `-s` option.
- s Forces `in.routed` to supply routing information whether it is acting as an internetwork router or not.
- S If `in.routed` is not acting as an internetwork router it will, instead of entering the whole routing table in the kernel, only enter a default route for each internetwork router. This reduces the the memory requirements without losing any routing reliability.
- t All packets sent or received are printed on standard output. In addition, `in.routed` will not divorce itself from the controlling terminal so that interrupts from the keyboard will kill the process. Any other argument supplied is interpreted as the name of the file in which `in.routed`’s actions should be logged. This log contains information about any changes to the routing tables and a history of recent messages sent and received which are related to the changed route.

in.routed(1M)

	-v	Allows a logfile (whose name must be supplied) to be created showing the changes made to the routing tables with a timestamp.				
FILES	/etc/gateways	For distant gateways				
	/etc/networks	Associations of Internet Protocol network numbers with network names				
	/etc/hosts	Internet host table				
	/etc/security/tsolgateways	For trusted routing through listed gateways				
	/etc/security/tsol/tunnel	Tunneling information table for Trusted Solaris hosts				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:					
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWcsu</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE					
Availability	SUNWcsu					
SUMMARY OF TRUSTED SOLARIS CHANGES	in.routed should be started at ADMIN_HIGH. It must inherit the net_mac_read, net_privaddr, net_broadcast, and sys_net_config privileges. If a log file is specified, in.routed must also inherit the file_mac_write privilege. Because trusted routing considers the security of the route along with the route's metric when making routing decisions, in.routed sends two additional types of response packets containing security information for routes: sec_response packets for communications with connected Trusted Solaris gateways, and sec_t_response packets for tunneling to Trusted Solaris gateways on the other side of non-Trusted Solaris gateways.					
Trusted Solaris 8 4/01 Reference Manual NOTES	route(1M)					
	ioctl(2), inet(3SOCKET), attributes(5), inet(7P), udp(7P)					
	The kernel's routing tables may not correspond to those of in.routed for short periods of time while processes that utilize existing routes exit; the only remedy for this is to place the routing process in the kernel.					
	in.routed should listen to intelligent interfaces, such as an IMP, and to error protocols, such as ICMP, to gather more information.					
	in.routed initially obtains a routing table by examining the interfaces configured on a machine and the gateways file. It then sends a request on all directly connected networks for more routing information. in.routed does not recognize or use any routing information already established on the machine prior to startup. With the exception of interface changes, in.routed does not see any routing table changes that have been done by other programs on the machine, for example, routes added, deleted or flushed by way of the route(1M) command. Therefore, these types of changes should not be done while in.routed is running. Rather, shut down in.routed, make the changes required, and then restart in.routed.					

NAME	in.rshd, rshd – Remote shell server
SYNOPSIS	in.rshd <i>host.port</i>
DESCRIPTION	<p>in.rshd is the server for the rsh(1) program. The server provides remote execution facilities with authentication based on privileged port numbers.</p> <p>in.rshd is invoked by inetd(1M) each time a shell service is requested, and executes the following protocol:</p> <ol style="list-style-type: none"> 1. The server checks the client's source port. If the port is not in the range 0-1023, the server aborts the connection. The client's host address (in hex) and port number (in decimal) are the arguments passed to in.rshd. 2. The server reads characters from the socket up to a null (0) byte. The resultant string is interpreted as an ASCII number, base 10. 3. If the number received in step 1 is non-zero, it is interpreted as the port number of a secondary stream to be used for the stderr. A second connection is then created to the specified port on the client's machine. The source port of this second connection is also in the range 0-1023. 4. The server checks the client's source address. If the address is associated with a host for which no corresponding entry exists in the host name data base (see hosts(4)), the server aborts the connection. Please refer to the SECURITY section below for more details. 5. A null terminated user name of at most 16 characters is retrieved on the initial socket. This user name is interpreted as a user identity to use on the <i>server's</i> machine. 6. A null terminated user name of at most 16 characters is retrieved on the initial socket. This user name is interpreted as the user identity on the <i>client's</i> machine. 7. A null terminated command to be passed to a shell is retrieved on the initial socket. The length of the command is limited by the upper bound on the size of the system's argument list. 8. in.rshd checks whether logins are currently allowed by looking for an /etc/nologin file. If the file exists, the connection is terminated. If logins are allowed, the user is validated according to the following steps. The remote user name is looked up in the password file and a chdir is performed to the user's home directory. If the lookup fails, the connection is terminated. If the chdir fails, it does a chdir to / (root). If the user is not the superuser, (user ID 0), and if the pam_rhosts_auth PAM module is configured for authentication, the file /etc/hosts.equiv is consulted for a list of hosts considered "equivalent". If the client's host name is present in this file, the authentication is considered successful. See the SECURITY section below for a discussion of PAM authentication. <p>If the lookup fails, or the user is root, then the file .rhosts in the home directory of the remote user is checked for the machine name and identity of the user on the client's machine. If this lookup fails, the connection is terminated</p>

in.rshd(1M)

9. A null byte is returned on the connection associated with the `stderr` and the command line is passed to the normal login shell of the user. (The `PATH` variable is set to `/usr/bin`.) The shell inherits the network connections established by `in.rshd`.

USAGE `rshd` and `in.rshd` are IPv6-enabled. See `ip6(7P)`.

SECURITY `in.rshd` uses `pam(3PAM)` for authentication, account management, and session management. The PAM configuration policy, listed through `/etc/pam.conf`, specifies the modules to be used for `in.rshd`. Here is a partial `pam.conf` file with entries for the `rsh` command using `rhhosts` authentication, UNIX account management, and session management module.

rsh	auth	required	/usr/lib/security/pam_rhosts_auth.so.1
rsh	account	required	/usr/lib/security/pam_unix.so.1
rsh	session	required	/usr/lib/security/pam_unix.so.1

If there are no entries for the `rsh` service, then the entries for the "other" service will be used. To maintain the authentication requirement for `in.rshd`, the `rsh` entry must always be configured with the `pam_rhosts_auth.so.1` module.

SUMMARY OF TRUSTED SOLARIS CHANGES If the `/etc/nologin` file exists, the server will not allow connections. The values of the trusted path, label view, and label-translation process attributes from the client process are propagated to the remote shell.

FILES `/etc/hosts.equiv`

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8 4/01 Reference Manual `inetd(1M)`, `inetd.conf(4)`, `pam_unix(5)`

Usage `rsh(1)`, `pam(3PAM)`, `hosts(4)`, `pam.conf(4)`, `attributes(5)`, `pam_rhosts_auth(5)`, `ip6(7P)`

DIAGNOSTICS The following diagnostic messages are returned on the connection associated with `stderr`, after which any network connections are closed. An error is indicated by a leading byte with a value of 1 in step 9 above (0 is returned above upon successful completion of all the steps prior to the command execution).

`locuser too long`

The name of the user on the client's machine is longer than 16 characters.

in.rshd(1M)

remuser too long

The name of the user on the remote machine is longer than 16 characters.

command too long

The command line passed exceeds the size of the argument list (as configured into the system).

Hostname for your address unknown.

No entry in the host name database existed for the client's machine.

Login incorrect.

No password file entry for the user name existed.

Permission denied.

The authentication procedure described above failed.

Can't make pipe.

The pipe needed for the stderr was not created.

Try again.

A fork by the server failed.

NOTES The authentication procedure used here assumes the integrity of each client machine and the connecting medium. This is insecure, but it is useful in an "open" environment.

A facility to allow all data exchanges to be encrypted should be present.

install(1M)

NAME	install – Install commands				
SYNOPSIS	<pre> /usr/sbin/install -c <i>dira</i> [-m <i>mode</i>] [-u <i>user</i>] [-g <i>group</i>] [-o] [-s] <i>file</i> /usr/sbin/install -f <i>dirb</i> [-m <i>mode</i>] [-u <i>user</i>] [-g <i>group</i>] [-o] [-s] <i>file</i> /usr/sbin/install -n <i>dirc</i> [-m <i>mode</i>] [-u <i>user</i>] [-g <i>group</i>] [-o] [-s] <i>file</i> /usr/sbin/install -d -i [-m <i>mode</i>] [-u <i>user</i>] [-g <i>group</i>] [-o] [-s] <i>dirx...</i> /usr/sbin/install [-m <i>mode</i>] [-u <i>user</i>] [-g <i>group</i>] [-o] [-s] <i>file</i> [<i>dirx...</i>] </pre>				
DESCRIPTION	<p>install is most commonly used in “makefiles” (see make(1S)) to install a <i>file</i> in specific locations, or to create directories within a file system. Each <i>file</i> is installed by copying it into the appropriate directory.</p> <p>install uses no special privileges to copy files from one place to another. The implications of this are:</p> <ul style="list-style-type: none"> ■ You must have permission to read the files to be installed. ■ You must have permission to copy into the destination directory. ■ You must have permission to change the modes on the final copy of the file if you want to use the -m option. ■ You must assume an administrative role if you want to specify the ownership of the installed file with the -u or -g options. If you are not in an administrative role, the installed file will be owned by you, regardless of who owns the original. <p>install prints messages telling the user exactly what files it is replacing or creating and where they are going.</p> <p>If no options or directories (<i>dirx...</i>) are given, install searches a set of default directories (/bin, /usr/bin, /etc, /lib, and /usr/lib, in that order) for a file with the same name as <i>file</i>. When the first occurrence is found, install issues a message saying that it is overwriting that file with <i>file</i>, and proceeds to do so. If the file is not found, the program states this and exits.</p> <p>If one or more directories (<i>dirx...</i>) are specified after <i>file</i>, those directories are searched before the default directories.</p>				
OPTIONS	<table> <tr> <td>-c <i>dira</i></td><td>Install <i>file</i> in the directory specified by <i>dira</i>, if <i>file</i> does not yet exist. If it is found, install issues a message saying that the file already exists, and exits without overwriting it.</td></tr> <tr> <td>-f <i>dirb</i></td><td>Force <i>file</i> to be installed in given directory, even if the file already exists. If the file being installed does not already exist, the mode and owner of the new file will be set to 755 and bin, respectively. If the file already exists, the mode and owner will be that of the already existing file.</td></tr> </table>	-c <i>dira</i>	Install <i>file</i> in the directory specified by <i>dira</i> , if <i>file</i> does not yet exist. If it is found, install issues a message saying that the file already exists, and exits without overwriting it.	-f <i>dirb</i>	Force <i>file</i> to be installed in given directory, even if the file already exists. If the file being installed does not already exist, the mode and owner of the new file will be set to 755 and bin, respectively. If the file already exists, the mode and owner will be that of the already existing file.
-c <i>dira</i>	Install <i>file</i> in the directory specified by <i>dira</i> , if <i>file</i> does not yet exist. If it is found, install issues a message saying that the file already exists, and exits without overwriting it.				
-f <i>dirb</i>	Force <i>file</i> to be installed in given directory, even if the file already exists. If the file being installed does not already exist, the mode and owner of the new file will be set to 755 and bin, respectively. If the file already exists, the mode and owner will be that of the already existing file.				

- n *dir* If *file* is not found in any of the searched directories, it is put in the directory specified in *dir*. The mode and owner of the new file will be set to 755 and *bin*, respectively.
- d Create a directory. Missing parent directories are created as required as in `mkdir -p`. If the directory already exists, the owner, group and mode will be set to the values given on the command line.
- i Ignore default directory list, searching only through the given directories (*dirx...*).
- m *mode* The mode of the new file is set to *mode*. Set to 0755 by default.
- u *user* The owner of the new file is set to *user*. Only available to administrative roles. Set to *bin* by default.
- g *group* The group id of the new file is set to *group*. Only available to administrative roles. Set to *bin* by default.
- o If *file* is found, save the “found” file by copying it to *OLDfile* in the directory in which it was found. This option is useful when installing a frequently used file such as */bin/sh* or */lib/saf/ttymon*, where the existing file cannot be removed.
- s Suppress printing of messages other than error messages.

USAGE See `largefile(5)` for the description of the behavior of `install` when encountering files greater than or equal to 2 Gbyte (2³¹ bytes).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES
Trusted Solaris 8 4/01 Reference Manual
Summary of Changes

To specify the ownership of an installed file with the `-u` or `-g` options, you must assume an administrative role.

`chgrp(1)`, `chmod(1)`, `chown(1)`, `mkdir(1)`
`cp(1)`, `make(1S)`, `chown(1M)`, `attributes(5)`, `largefile(5)`

in.tftpd(1M)

NAME	in.tftpd, tftpd – Internet Trivial File Transfer Protocol server
SYNOPSIS	in.tftpd [-s] [<i>homedir</i>]
DESCRIPTION	<p>tftpd is a server that supports the Internet Trivial File Transfer Protocol (TFTP). This server is normally started by inetd(1M) and operates at the port indicated in the tftp Internet service description in the /etc/inetd.conf file. By default, the entry for in.tftpd in etc/inetd.conf is commented out. To make in.tftpd operational, the comment character(s) must be deleted from the file. See inetd.conf(4).</p> <p>Before responding to a request, the server attempts to change its current directory to <i>homedir</i>; the default directory is /tftpboot.</p> <p>The use of tftp does not require an account or password on the remote system. Due to the lack of authentication information, in.tftpd will allow only publicly readable files to be accessed. Files may be written only if they already exist and are publicly writable. Note that this extends the concept of “public” to include all users on all hosts that can be reached through the network; this may not be appropriate on all systems, and its implications should be considered before enabling this service.</p> <p>in.tftpd runs with the user ID and group ID set to [GU] ID_NOBODY under the assumption that no files exist with that owner or group. However, nothing checks this assumption or enforces this restriction.</p>
OPTIONS	-s Secure. When specified, the directory change to <i>homedir</i> must succeed. The daemon also changes its root directory to <i>homedir</i> .
SUMMARY OF TRUSTED SOLARIS CHANGES	in.tftpd should be started from the trusted path with a UID of 0; it must inherit the proc_chroot, proc_owner, and proc_setid privileges.
	/etc/inetd.conf Configuration file for inetd.
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

inetd(1M)
tftp(1), inetd.conf(4), netconfig(4), attributes(5), ip6(7P)
Sollins, K.R., *The TFTP Protocol (Revision 2)*, RFC 783, Network Information Center, SRI International, Menlo Park, California, June 1981.

NAME	ipsecconf – configure system wide IPsec policy
SYNOPSIS	<pre> /usr/sbin/ipsecconf /usr/sbin/ipsecconf -a file [-q] /usr/sbin/ipsecconf -d index /usr/sbin/ipsecconf -f /usr/sbin/ipsecconf -l [-n] </pre>
DESCRIPTION	<p>The <code>ipsecconf</code> utility configures the IPsec policy for a host. Once the policy is configured, all outbound and inbound datagrams are subject to policy checks as they exit and enter the host. If no entry is found, no policy checks will be completed, and all the traffic will pass through. Datagrams that are being forwarded will not be subjected to policy checks that are added using this command. See <code>ifconfig(1M)</code> and <code>tun(7M)</code> for information on how to protect forwarded packets. Depending upon the match of the policy entry, a specific action will be taken.</p> <p>This command requires the <code>sys_net_config</code> privilege. The label in the <code>exec_attr(4)</code> profile needs to match the label assigned to files read by this command. Each entry protects traffic only in one direction, that is, either outbound or inbound. Thus to protect traffic in both directions, you need to have one entry in each direction.</p> <p>When the command is issued without any arguments, the policies configured in the system are shown. Each entry is displayed with an <i>index</i> followed by a number. You can use the <code>-d</code> option with the <i>index</i> to delete a given policy in the system. The entries are displayed in the order that they were added, which is not necessarily the order that the traffic match will take place. To view the order in which the traffic match will take place, use the <code>-l</code> option.</p> <p>Policy entries are not preserved across reboot. Thus the policy needs to be added everytime the machine reboots. To configure policies early in the boot, one can setup policies in the <code>/etc/inet/ipsecinit.conf</code> file, which are then read from the <code>inetinit</code> startup script.</p> <p>See SECURITY CONSIDERATIONS.</p>
OPTIONS	<p><code>ipsecconf</code> supports the following options:</p> <p><i>-a file</i> Add the IPsec policy to the system as specified by each entry in the <i>file</i>. An IPsec configuration file contains one or more entries that specify the configuration. Once the policy is added, all outbound and inbound datagrams are subject to policy checks.</p> <p>Entries in the files are described in the OPERANDS section below. Examples can be found in the EXAMPLES section below.</p> <p>Policy is latched for TCP/UDP sockets on which a <code>connect(3SOCKET)</code> or <code>accept(3SOCKET)</code> is issued. So, the</p>

ipseconf(1M)

	<p>addition of new policy entries may not affect such endpoints or sockets. However, the policy will be latched for a socket with an existing non-null policy. Thus, make sure that there are no preexisting connections that will be subject to checks by the new policy entries.</p> <p>The feature of policy latching explained above may change in the future. It is not advisable to depend upon this feature.</p>
-d <i>index</i>	Delete the policy denoted by the <i>index</i> . The <i>index</i> is obtained by viewing the policy configured in the system. Once the entry is deleted, all outbound and inbound datagrams affected by this policy entry will not be subjected to policy checks. Be advised that with connections for which the policy has been latched, packets will continue to go out with the same policy, even if it has been deleted.
-f	Flush all the policies in the system. Constraints are similar to the -d option with respect to latching.
-l	Listing of the internal system policy table. When ipseconf is invoked without any arguments, a complete list of policy entries added by the user since boot is displayed. The current table can differ from the previous one if, for example, a multi-homed entry was added or policy reordering occurred. In the case of a multi-homed entry, all the addresses are listed explicitly. If a mask was not specified earlier but was instead inferred from the address, it will be explicitly listed here. This option is used to view policy entries in the correct order. The outbound and inbound policy entries are listed separately.
-n	Show network addresses, ports, protocols in numbers. The -n option may only be used with the -l option.
-q	Quiet mode. Suppresses the warning message generated when adding policies.
OPERANDS	<p>Each policy entry contains 3 parts specified as follows :</p> <p>{pattern} action {properties} Every policy entry begins on a new line and can span multiple lines. "pattern" specifies the traffic pattern that should be matched against the outbound and inbound datagrams. If there is a match, a specific "action" determined by the second argument will be taken, depending upon the "properties" of the policy entry. Pattern and properties are name-value pairs where name and value are separated by space, tab or newline. Multiple name-value pairs should be separated by space, tab or newline. The beginning and end of the pattern and properties are marked by "{" and "}" respectively.</p>

Files can contain multiple policy entries. An unspecified name-value pair in the "pattern" will be considered as a wildcard. Wildcard entries match any corresponding entry in the datagram.

File can be commented by using "#" as the first character. Comments may be inserted either at the beginning or the end of a line.

The complete syntax of a policy entry is:

```

policy ::= {pattern} action {properties}

pattern ::= <pattern_name_value_pair>|
            <pattern_name_value_pair>, <pattern>

action ::= apply | permit | bypass

properties ::= <prop_name_value_pair>|
              <prop_name_value_pair>, <properties>

pattern_name_value_pair ::=
    <saddr/prefix address>|
    <smask mask>|
    <sport part>|
    <daddr/prefix address>|
    <dmask mask>|
    <dport port>|
    <ulp protocol>

address ::= <Internet dot notation> | <String recognized by gethostbyname> |
           <String recognized by getnetbyname>

prefix ::= <number>

mask ::= <0xhexdigit[hexdigit]> | <0Xhexdigit[hexdigit]> |
        <Internet dot notation>

port ::= <number> | <String recognized by getservbyname>

protocol ::= <number> | <String recognized by getprotobyname>

prop_name_value_pair ::=
    <auth_algs auth_alg>|
    <encr_algs encr_alg>|
    <encr_auth_algs auth_alg>|
    <sa sa_val>|
    <dir dir_val>

auth_alg ::= <md5 | hmac-md5 | sha | sha1 | hmac-sha | hmac-sha1 | number>

encr_alg ::= <des | des-cbc | 3des | 3des-cbc | number>

sa_val ::= shared | unique

dir_val ::= out | in

number ::= < 0 | 1 | 2 ... 9> <number>

```

ipsecconf(1M)

Policy entries may contain the following (name value) pairs in the pattern field. Each (name value) pair may appear only once in given policy entry.

saddr/plen

The value that follows is the source address of the datagram with the prefix length. Only *plen* leading bits of the source address of the packet will be matched. *plen* is optional.

The source address value can be a hostname as described in `gethostbyname(3XNET)` or a network name as described in `getnetbyname(3XNET)` or a host address or network address in the Internet standard dot notation. See `inet_addr(3XNET)`.

If a hostname is given and `gethostbyname(3XNET)` returns multiple addresses for the host, then policy will be added for each of the addresses with other entries remaining the same.

daddr/plen

The value that follows is the destination address of the datagram with the prefix length. Only *plen* leading bits of the destination address of the packet will be matched. *plen* is optional.

See *saddr* for valid values that can be given. If multiple source and destination addresses are found, then a policy entry that covers each source address-destination address pair will be added to the system.

smask

The value that follows is the source mask. If prefix length is given with *saddr*, this should not be given. This can be represented either in hexadecimal number with a leading 0x or 0X, for example, 0xffff0000, 0Xffff0000 or in the Internet decimal dot notation, for example, 255.255.0.0 and 255.255.255.0. The mask should be contiguous and the behavior is not defined for non-contiguous masks.

smask is considered only when *saddr* is given.

dmask

The value that follows is the destination mask. If prefix length is given with *daddr*, this should not be given. This can be represented either in hexadecimal number with a leading 0x or 0X, for example, 0xffff0000, 0Xffff0000 or in the Internet decimal dot notation, for example, 255.255.0.0 and 255.255.255.0. The mask should be contiguous and the behavior is not defined for non-contiguous masks.

	<i>dmask</i> is considered only when <i>daddr</i> is given.
<i>sport</i>	The value that follows is the source port of the datagram. This can be either a port number or a string searched with a <code>NULL</code> proto argument, as described in <code>getservbyname(3XNET)</code>
<i>dport</i>	The value that follows is the destination port of the datagram. This can be either a port number or a string as described in <code>getservbyname(3XNET)</code> searched with <code>NULL</code> proto argument.
<i>ulp</i>	The value that follows is the Upper Layer Protocol that this entry should be matched against. It could be a number or a string as described in <code>getprotobyname(3XNET)</code>

If any component of the entry is not given, it will be considered as a wildcard entry. Thus, if the pattern is null, all packets will match the policy entry. If neither the prefix length nor the mask is given for the address, a mask will be inferred. For example, if `a.b.c.d` is the address and

- `b, c` and `d` are zeroes, the mask is `0xff000000`.
- only `c` and `d` are zeroes, the mask is `0xffff0000`.
- only `d` is zero, the mask is `0xffffffff00`.
- neither `a, b, c`, nor `d` are zeroes, the mask is `0xffffffff`.

To avoid ambiguities, it is advisable to explicitly give either the prefix length or the mask.

Policy entries may contain the following (name value) pairs in the properties field. Each (name value) pair may appear only once in a given policy entry.

<code>auth_algs</code>	<p>An acceptable value following this implies that IPsec AH header will be present in the outbound datagram. Values following this describe the authentication algorithms that will be used while applying the IPsec AH on outbound datagrams and verified to be present on inbound datagrams. See <i>RFC 2402</i>.</p> <p>This entry can contain either a string or a decimal number.</p> <table> <tr> <td><code>string</code></td><td>This should be either MD5 or HMAC-MD5 denoting the HMAC-MD5 algorithm as described in <i>RFC 2403</i>, and SHA1, or HMAC-SHA1 or SHA or HMAC-SHA denoting the HMAC-SHA algorithm described in <i>RFC 2404</i>. The string can also be</td></tr> </table>	<code>string</code>	This should be either MD5 or HMAC-MD5 denoting the HMAC-MD5 algorithm as described in <i>RFC 2403</i> , and SHA1, or HMAC-SHA1 or SHA or HMAC-SHA denoting the HMAC-SHA algorithm described in <i>RFC 2404</i> . The string can also be
<code>string</code>	This should be either MD5 or HMAC-MD5 denoting the HMAC-MD5 algorithm as described in <i>RFC 2403</i> , and SHA1, or HMAC-SHA1 or SHA or HMAC-SHA denoting the HMAC-SHA algorithm described in <i>RFC 2404</i> . The string can also be		

ipsecconf(1M)

		<p>ANY, which denotes no-preference for the algorithm. Default algorithms will be chosen based upon the SAs available at this time for manual SAs and the key negotiating daemon for automatic SAs. Strings are not case-sensitive.</p>
	number	<p>A number in the range 1-255. This is useful when new algorithms can be dynamically loaded.</p>
		<p>If <i>auth_algs</i> is not present, the AH header will not be present in the outbound datagram, and the same will be verified for the inbound datagram.</p>
encr_algs		<p>An acceptable value following this implies that IPsec ESP header will be present in the outbound datagram. The value following this describes the encryption algorithms that will be used to apply the IPsec ESP protocol to outbound datagrams and verify it to be present on inbound datagrams. See <i>RFC 2406</i>.</p> <p>This entry can contain either a string or a decimal number. Strings are not case-sensitive.</p>
	string	<p>This should be either DES or DES-CBC, to denote the algorithm described in <i>RFC 2405</i> or 3DES or 3DES-CBC, to denote the used of 3DES in a manner consistent with <i>RFC 2451</i>. The value can be NULL which implies a NULL encryption pursuant to <i>RFC 2410</i>. This means that the payload will not be encrypted. The string can be ANY, which denotes no preference for the algorithm. Default algorithms will be chosen depending upon the SAs available at this time for manual SAs and upon the key negotiating daemon for automatic SAs.</p>
	number	<p>A decimal number in the range 1-255. This is useful when new algorithms can be dynamically loaded.</p>
encr_auth_algs		<p>An acceptable value following <i>encr_auth_algs</i> implies that the IPsec ESP header will be present in the</p>

outbound datagram. The values following `encr_auth_algs` describe the authentication algorithms that will be used while applying the IPsec ESP protocol on outbound datagrams and verified to be present on inbound datagrams. See *RFC 2406*. This entry can contain either a string or a number. Strings are case-insensitive.

`string` Valid values are the same as the ones described for `auth_algs` above.

`number` This should be a decimal number in the range 1-255. This is useful when new algorithms can be dynamically loaded. If `encr_algs` is present and `encr_auth_algs` is not present in a policy entry, the system will use an ESP SA regardless of whether the SA has an authentication algorithm or not.

If `encr_algs` is not present and `encr_auth_algs` is present in a policy entry, null encryption will be provided, which is equivalent to `encr_algs` with `NULL`, for outbound and inbound datagrams.

If both `encr_algs` and `encr_auth_algs` are not present in a policy entry, ESP header will not be present for outbound datagrams and the same will be verified for inbound datagrams.

If both `encr_algs` and `encr_auth_algs` are present in a policy entry, ESP header with integrity checksum will be present on outbound datagrams and the same will be verified for inbound datagrams.

`dir`

Values following this decides whether this entry is for outbound or inbound datagram. Valid values are strings that should be one of the following:

`out` This means that this policy entry should be considered only for outbound datagrams.

`in` This means that this policy entry should be considered only for inbound datagrams.

This entry is not needed when the action is "apply" or "permit". But if it is given while the action is "apply" or

ipsecconf(1M)

sa	<p>"permit", it should be "out" or "in" respectively. This is mandatory when the action is "bypass".</p>		
	<p>Values following this decide the attribute of the security association. Value indicates whether a unique security association should be used or any existing SA can be used. If there is a policy requirement, SAs are created dynamically on the first outbound datagram using the key management daemon. Static SAs can be created using <code>ipseckey(1M)</code>. The values used here determine whether a new SA will be used/obtained. Valid values are strings that could be one of the following:</p>		
	<table> <tr> <td data-bbox="792 716 883 737">unique</td><td data-bbox="1013 716 1398 1192"> <p>Unique Association. A new/unused association will be obtained/used for packets matching this policy entry. If an SA that was previously used by the same 5 tuples, that is, {Source address, Destination address, Source port, Destination Port, Protocol (for example, TCP/UDP)} exists, it will be reused. Thus uniqueness is expressed by the 5 tuples given above. The security association used by the above 5 tuples will not be used by any other socket. For inbound datagrams, uniqueness will not be verified.</p> </td></tr> </table>	unique	<p>Unique Association. A new/unused association will be obtained/used for packets matching this policy entry. If an SA that was previously used by the same 5 tuples, that is, {Source address, Destination address, Source port, Destination Port, Protocol (for example, TCP/UDP)} exists, it will be reused. Thus uniqueness is expressed by the 5 tuples given above. The security association used by the above 5 tuples will not be used by any other socket. For inbound datagrams, uniqueness will not be verified.</p>
unique	<p>Unique Association. A new/unused association will be obtained/used for packets matching this policy entry. If an SA that was previously used by the same 5 tuples, that is, {Source address, Destination address, Source port, Destination Port, Protocol (for example, TCP/UDP)} exists, it will be reused. Thus uniqueness is expressed by the 5 tuples given above. The security association used by the above 5 tuples will not be used by any other socket. For inbound datagrams, uniqueness will not be verified.</p>		
	<table> <tr> <td data-bbox="792 1213 883 1234">shared</td><td data-bbox="1013 1213 1398 1486"> <p>Shared association. If an SA exists already for this source-destination pair, it will be used. Otherwise a new SA will be obtained. This is mandatory only for outbound policy entries and should not be given for entries whose action is "bypass". If this entry is not given for inbound entries, for example, when "dir" is in or "action" is permit, it will be assumed to be shared.</p> </td></tr> </table>	shared	<p>Shared association. If an SA exists already for this source-destination pair, it will be used. Otherwise a new SA will be obtained. This is mandatory only for outbound policy entries and should not be given for entries whose action is "bypass". If this entry is not given for inbound entries, for example, when "dir" is in or "action" is permit, it will be assumed to be shared.</p>
shared	<p>Shared association. If an SA exists already for this source-destination pair, it will be used. Otherwise a new SA will be obtained. This is mandatory only for outbound policy entries and should not be given for entries whose action is "bypass". If this entry is not given for inbound entries, for example, when "dir" is in or "action" is permit, it will be assumed to be shared.</p>		
	<p>Action follows the pattern and should be given before properties. It should be one of the following and this field is mandatory.</p>		
apply	<p>Apply IPsec to the datagram as described by the properties, if the pattern matches the datagram. If <code>apply</code> is given, the pattern is matched only on the outbound datagram.</p>		
permit	<p>Permit the datagram if the pattern matches the incoming datagram and satisfies the constraints described by the properties. If it does</p>		

not satisfy the properties, discard the datagram. If `permit` is given, the pattern is matched only for inbound datagrams.

`bypass` Bypass any policy checks if the pattern matches the datagram. `dir` in the properties decides whether the check is done on outbound or inbound datagrams. All the `bypass` entries are checked before checking with any other policy entry in the system. This has the highest precedence over any other entries. `dir` is the only field that should be present when action is `bypass`.

If the file contains multiple policy entries, for example, they are assumed to be listed in the order in which they are to be applied. In cases of multiple entries matching the outbound and inbound datagram, the first match will be taken. The system will reorder the policy entry, that is, add the new entry before the old entry, only when:

- The level of protection is "stronger" than the old level of protection. Currently, strength is defined as:

AH and ESP > ESP > AH The standard uses of AH and ESP were what drove this ranking of "stronger". There are flaws with this. ESP can be used either without authentication, which will allow cut-and-paste or replay attacks, or without encryption, which makes it equivalent or slightly weaker than AH. An administrator should take care to use ESP properly. See `ipsecesp(7P)` for more details.

- If the new entry has `bypass` as action.

`bypass` has the highest precedence. It can be added in any order, and the system will still match all the `bypass` entries before matching any other entries. This is useful for key management daemons which can use this feature to bypass IPsec as it protects its own traffic.

Entries with both AH (`auth_algs` present in the policy entry) and ESP (`encr_auth_algs` or `encr_auth_algs` present in the policy entry) protection are ordered after all the entries with AH and ESP and before any AH-only and ESP-only entries. In all other cases the order specified by the user is not modified, that is, newer entries are added at the end of all the old entries. See EXAMPLES.

A new entry is considered duplicate of the old entry if an old entry matches the same traffic pattern as the new entry. See EXAMPLES for information on duplicates.

SECURITY CONSIDERATIONS

If, for example, the policy file comes over the wire from an NFS mounted file system, an adversary can modify the data contained in the file, thus changing the policy configured on the machine to suit his needs. Administrators should be cautious about transmitting a copy of the policy file over a network.

Policy is latched for TCP/UDP sockets on which a `connect(3SOCKET)` or `accept(3SOCKET)` has been issued. Adding new policy entries will not have any effect on them. This feature of latching may change in the future. It is not advisable to depend upon this feature.

Make sure to set up the policies before starting any communications, as existing connections may be affected by the addition of new policy entries. Similarly, do not change policies in the middle of a communication.

If your source address is a host that can be looked up over the network, and your naming system itself is compromised, then any names used will no longer be trustworthy.

EXAMPLES

EXAMPLE 1 Protecting Outbound TCP Traffic With ESP and the DES Algorithm

```
#
# Protect the outbound TCP traffic between hosts spiderweb
# and arachnid with ESP and use DES algorithm.
#
{
    saddr spiderweb
    daddr arachnid
    ulp tcp          #only TCP datagrams.
} apply {
    encr_algs DES
}
```

This entry specifies that any TCP packet from spiderweb to arachnid should be encrypted with DES, and the SA could be a shared one. As no prefix length or mask is given, a mask will be inferred. To look at the mask, use the `ipsecconf` command with the `-l` option. Note that `dir` is not given in properties as `apply` implies that only outbound packets will be matched with the pattern.

EXAMPLE 2 Verifying Whether or Not Inbound Traffic is Encrypted

The above entry will not verify whether or not the inbound traffic is encrypted. Thus you need the following entry to protect inbound traffic:

```
#
# Protect the TCP traffic on inbound with ESP/DES from arachnid
# to spiderweb
#
{
    saddr arachnid
    daddr spiderweb
    ulp tcp
} permit {
    encr_algs DES
}
```

"sa" can be absent for inbound policy entries as it implies that it can be a shared one. Uniqueness is not verified on inbound. Note that in both the above entries, authentication was never specified. This can lead to cut and paste attacks. As

EXAMPLE 2 Verifying Whether or Not Inbound Traffic is Encrypted (Continued)

mentioned previously, though the authentication is not specified, the system will still use an ESP SA with `encr_auth_alg` specified, if it was found in the SA tables.

EXAMPLE 3 Authenticating All Inbound Traffic to the Telnet Port

```
#
# All the inbound traffic to the telnet port should be
# authenticated.
#
{
    dport telnet          # telnet is 23
} permit {
    auth_algs SHA1
    dir in
}
```

This entry specifies that any inbound datagram to telnet port should come in authenticated with the SHA1 algorithm. Otherwise the datagram should not be permitted. Without this entry, traffic destined to port number 23 can come in clear. Note that `dir` as given is optional, as `permit` implies that this policy entry will be checked only on inbound. "sa" is not specified, which implies that it is shared. This can be done only for inbound entries. You need to have an equivalent entry to protect outbound traffic so that the outbound traffic is authenticated as well.

EXAMPLE 4 Verifying Inbound Traffic is Null-Encrypted

```
#
# Make sure that all inbound traffic from network-B is NULL
# encrypted, but bypass for host-B alone from that network.
# Add the bypass first.
{
    saddr host-B
} bypass {
    dir in
}
# Now add for network-B.
{
    saddr network-B/16
} permit {
    encr_algs NULL
    encr_auth_algs md5
}
```

The first entry specifies that any packet with address host-B should not be checked against any policies. The second entry specifies that all inbound traffic from network-B should be encrypted with a NULL encryption algorithm and the MD5 authentication algorithm. NULL encryption implies that ESP header will be used without encrypting the datagram. As the first entry is bypass it need not be given first in order, as bypass entries have the highest precedence. Thus any inbound traffic will be matched against all bypass entries before any other policy entries.

EXAMPLE 5 Encrypting a Packet with 3DES and SHA1

The following entry on hostB specifies that any packet from hostA to hostB should be encrypted with 3DES and SHA1.

```
{
    saddr hostA
    daddr hostB
} permit {
    encr_algs 3DES
    encr_auth_algs SHA1
}
```

If you try to add an entry

```
{
    saddr hostA
    daddr hostB
    dport 23
} permit {
    encr_algs DES
}
```

it will fail with "ioctl: File exists". But if you change the order, that is, give the second entry first, and first entry second, it will succeed. This is because traffic to port number 23 from hostB to hostA will be protected with DES and the remainder will be protected with 3DES and SHA1.

If you modify the second entry as follows,

```
{
    saddr hostA
    daddr hostB
    dport 23
} permit {
    encr_algs DES
    auth_algs SHA1
}
```

it will not fail. This entry gets ordered first in the list, as the entry is protected with AH and ESP, which has precedence before the prior entry that has only ESP. You can add a bypass entry in any order and it will always have the highest precedence. But, all other entries are subject to the check as explained above.

The following entry

```
{
    daddr 134.56.0.0      # Network address
    dmask 0xffff0000
} permit { auth_algs any}
```

expects any traffic originating from 134.56.0.0 to be authenticated. You cannot add the following entry after the above entry has been added,

```
{
    daddr 134.56.123.0
    dmask 0xffffffff00
} permit { encr_algs any}
```

EXAMPLE 5 Encrypting a Packet with 3DES and SHA1 (Continued)

as the previous entry would match the traffic from 134.56.0.0. But you can add this entry before adding the previous entry, or you can add it with AH and ESP protection. It will be reordered and considered before the previous one.

EXAMPLE 6 Entries to Bypass Traffic from IPsec

The first two entries provide that any datagram leaving the machine with source port 500 or coming into port number 500 should not be subjected to IPsec policy checks, irrespective of any other policy entry in the system. Thus the latter two entries will be considered only for ports other than port number 500.

```
#
# Bypass traffic for port no 500
#
{sport 500} bypass {dir out}
{dport 500} bypass {dir in}
{saddr spiderweb} apply { encr_algs any sa unique}
{daddr spiderweb} permit { encr_algs any}
```

EXAMPLE 7 Protecting Outbound Traffic

```
#
# Protect the outbound traffic from all interfaces.
#
{ saddr spiderweb} apply {auth_algs any sa unique}
```

If the `gethostbyname()` call for `spiderweb` yields multiple addresses, multiple policy entries will be added for all the source address with the same properties.

```
{
    saddr spiderweb
    daddr arachnid
} apply { auth_algs any sa unique}
```

If the `gethostbyname` call for `spiderweb` and the `gethostbyname` call for `arachnid` yield multiple addresses, multiple policy entries will be added for each (saddr daddr) pair with the same properties. Use `ipseccnf -l` to view all the policy entries added here.

EXAMPLE 8 Bypassing Unauthenticated Traffic

```
#
# Protect all the outbound traffic with ESP except any traffic
# to network-b which should be authenticated and bypass anything
# to network-c
#
{daddr network-b/16} apply { auth_algs any }
{ } apply { encr_algs any sa shared} # NULL pattern
{daddr network-c/16} bypass {dir out}
```

Note that `bypass` can be given anywhere and it will take precedence over all other entries. `NULL` pattern matches all the traffic.

ipseconf(1M)

FILES	/var/run/ipsecpolicy.conf	Cache of IPsec policies currently configured for the system, maintained by ipseconf command. Do not edit this file.
	/etc/inet/ipsecinit.conf	File containing IPsec policies to be installed at the time the system transitions from run-level 2 or 3. If present, these policies are loaded after /usr is mounted but before any non-boot-time routing information is processed and before any Internet services are started, including naming services.
	/etc/inet/ipsecinit.sample	Sample input file for ipseconf.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu
Interface Stability	Evolving

SUMMARY OF TRUSTED SOLARIS CHANGES

Trusted Solaris 8
4/01 Reference Manual
SunOS 5.8
Reference Manual

This command requires the sys_net_config privilege. The label in the exec_attr(4) profile needs to match the label assigned to files read by this command.

init(1M), ifconfig(1M), ipseckey(1M), accept(3SOCKET)

connect(3SOCKET), gethostbyname(3XNET), getnetbyname(3XNET), getprotobyname(3XNET), getservbyname(3XNET), socket(3SOCKET), attributes(5), ipsecah(7P), ipsecesp(7P), tun(7M)

Glenn, R. and Kent, S., *RFC 2410, The NULL Encryption Algorithm and Its Use With IPsec*, The Internet Society, 1998.

Kent, S. and Atkinson, R., *RFC 2402, IP Authentication Header*, The Internet Society, 1998.

Kent, S. and Atkinson, R., *RFC 2406, IP Encapsulating Security Payload (ESP)*, The Internet Society, 1998.

Madsen, C. and Glenn, R., *RFC 2403, The Use of HMAC-MD5-96 within ESP and AH*, The Internet Society, 1998.

Madsen, C. and Glenn, R., *RFC 2404, The Use of HMAC-SHA-1-96 within ESP and AH*, The Internet Society, 1998.

Madsen, C. and Doraswamy, N., *RFC 2405, The ESP DES-CBC Cipher Algorithm With Explicit IV*, The Internet Society, 1998.

DIAGNOSTICS

Pereira, R. and Adams, R., *RFC 2451, The ESP CBC-Mode Cipher Algorithms*, The Internet Society, 1998.

Bad "string" on line N.

Duplicate "string" on line N.

"string" refers to one of the names in pattern or properties. A Bad string indicates that an argument is malformed; a Duplicate string indicates that there are multiple arguments of a similar type, for example, multiple Source Address arguments..

Error before or at line N.

Indicates parsing error before or at line N.

Non-existent index

Reported when the *index* for delete is not a valid one.

ioctl: File exists

Reported when there is already a policy entry that matches the traffic of this new entry.

ipseckey(1M)

NAME	ipseckey – manually manipulate an IPsec Security Association Database (SADB)	
SYNOPSIS	<pre>ipseckey [-nvp] ipseckey [-nvp] -f <i>filename</i> ipseckey [-nvp] [delete get] SA_TYPE {EXTENSION <i>value...</i>} ipseckey [-np] [monitor passive_monitor pmonitor] ipseckey [-nvp] flush {SA_TYPE} ipseckey [-nvp] dump {SA_TYPE} ipseckey [-nvp] save SA_TYPE {<i>filename</i>} ipseckey [-nvp] -s <i>filename</i></pre>	
DESCRIPTION	<p>The ipseckey command is used to manually manipulate the security association databases of the network security services, ipsecak(7P) and ipsecesp(7P). You can use the ipseckey command to set up security associations between communicating parties when automated key management is not available.</p> <p>While the ipseckey utility has only a limited number of general options, it supports a rich command language. The user may specify requests to be delivered by means of a programmatic interface specific for manual keying. See pf_key(7P). When ipseckey is invoked with no arguments, it will enter an interactive mode which prints a prompt to the standard output and accepts commands from the standard input until the end-of-file is reached. Some commands require an explicit security association (“SA”) type, while others permit the SA type to be unspecified and act on all SA types.</p> <p>ipseckey uses a PF_KEY socket and the message types SADB_ADD, SADB_DELETE, SADB_GET, SADB_UPDATE, SADB_FLUSH, and SADB_X_PROMISE. This command requires the sys_net_config privilege. The label in the exec_attr(4) profile needs to match the label assigned to files read by this command.</p> <p>ipseckey handles sensitive cryptographic keying information. Please read the SECURITY CONSIDERATIONS section for details on how to use this command securely.</p>	
OPTIONS	<pre>-f [<i>filename</i>]</pre>	<p>Read commands from an input file, <i>filename</i>. The lines of the input file are identical to the command line language. The load command provides similar functionality. The -s option or the save command can generate files readable by the -f argument.</p>
	<pre>-n</pre>	<p>Prevent attempts to print host and network names symbolically when reporting actions. This is useful, for example, when all name servers are down or are otherwise unreachable.</p>

COMMANDS

-p	Paranoid. Do not print any keying material, even if saving SAs. Instead of an actual hexadecimal digit, print an X when this flag is turned on.
-s <i>[filename]</i>	The opposite of the -f option. If '-' is given for a <i>filename</i> , then the output goes to the standard output. A snapshot of all current SA tables will be output in a form readable by the -f option. The output will be a series of add commands.
-v	Verbose. Print the messages being sent into the PF_KEY socket, and print raw seconds values for lifetimes.
add	Add an SA. Because it involves the transfer of keying material, it cannot be invoked from the command line. The add command accepts all extension-value pairs described below.
update	Update SA lifetime, and in the cases of larval SAs (leftover from faulty automated key management), keying material and other extensions. Like add, this command cannot be invoked from the command line because keying material could be seen by the ps(1) command. The update command accepts all extension-value pairs, but normally is only used for SA lifetime updates.
delete	Delete a specific SA from a specific SADB. This command requires the spi extension, and the dest extension for IPsec SAs. Other extension-value pairs are superfluous for a delete message.
get	Lookup and display a security association from a specific SADB. Like delete, this command only requires spi and dest for IPsec.
flush	Remove all SA for a given SA_TYPE, or all SA for all types.
monitor	Continuously report on any PF_KEY messages. This uses the SADB_X_PROMISC message to enable messages that a normal PF_KEY socket would not receive to be received. See pf_key(7P).
passive_monitor	Like monitor, except that it does not use the SADB_X_PROMISC message.
pmonitor	Synonym for passive_monitor.
dump	Will display all SAs for a given SA type, or will display all SAs. Because of the large amount of data generated by this command, there is no guarantee that all SA

ipseckey(1M)

SECURITY ASSOCIATION TYPES

	information will be successfully delivered, or that this command will even complete.
save	Is the command analog of the <code>-s</code> option. It is included as a command to provide a way to snapshot a particular SA type, for example, <code>esp</code> or <code>ah</code> .
help	Prints a brief summary of commands.
all	Specifies all known SA types. This type is only used for the <code>flush</code> and <code>dump</code> commands. This is equivalent to having no SA type for these commands.
ah	Specifies the IPsec Authentication Header ("AH") SA.
esp	Specifies the IPsec Encapsulating Security Payload ("ESP") SA.

EXTENSION VALUE TYPES

Commands like `add`, `delete`, `get`, and `update` require that certain extensions and associated values be specified. The extensions will be listed here, followed by the commands that use them, and the commands that require them. Requirements are currently documented based upon the IPsec definitions of an SA. Required extensions may change in the future. `<number>` can be in either hex (`0xnnn`), decimal (`nnn`) or octal (`0nnn`). `<string>` is a text string. `<hexstr>` is a long hexadecimal number with a bit-length. Extensions are usually paired with values; however, some extensions require two values after them.

`spi <number>`

Specifies the security parameters index of the SA. This extension is required for the `add`, `delete`, `get` and `update` commands.

`replay <number>`

Specifies the replay window size. If not specified, the replay window size is assumed to be zero. It is not recommended that manually added SAs have a replay window. This extension is used by the `add` and `update` commands.

`state <string> | <number>`

Specifies the SA state, either by numeric value or by the strings "larval", "mature", "dying" or "dead". If not specified, the value defaults to mature. This extension is used by the `add` and `update` commands.

`auth_alg <string> | <number>`

`authalg <string> | <number>`

Specifies the authentication algorithm for an SA, either by numeric value, or by strings indicating an algorithm name. Current authentication algorithms include:

HMAC-MD5 md5, hmac-md5

HMAC-SH-1 sha, sha-1, hmac-sha1, hmac-sha

Often, algorithm names will have several synonyms. This extension is required by the `add` command for certain SA types. It is also used by the `update` command.

encr_alg <string> | <number>

encr_alg <string> | <number>

Specifies the encryption algorithm for an SA, either by numeric value, or by strings indicating an algorithm name. Current encryption algorithms include DES ("des") and Triple-DES ("3des"). This extension is required by the add command for certain SA types. It is also used by the update command.

The next six extensions are lifetime extensions. There are two varieties, "hard" and "soft". If a hard lifetime expires, the SA will be deleted automatically by the system. If a soft lifetime expires, an SADB_EXPIRE message will be transmitted by the system, and its state will be downgraded to dying from mature. See pf_key(7P). The monitor command to key allows you to view SADB_EXPIRE messages.

soft_bytes <number>

hard_bytes <number>

Specifies the number of bytes that this SA can protect. If <number> is not specified, the default value is zero, which means that the SA will not expire based on the number of bytes protected. This extension is used by the add and update commands.

soft_addtime <number>

hard_addtime <number>

Specifies the number of seconds that this SA can exist after being added or updated from a larval SA. An update of a mature SA does not reset the initial time that it was added. If <number> is not specified, the default value is zero, which means the SA will not expire based on how long it has been since it was added. This extension is used by the add and update commands.

soft_usetime <number>

hard_usetime <number>

Specifies the number of seconds this SA can exist after first being used. If <number> is not specified, the default value is zero, which means the SA will not expire based on how long it has been since it was added. This extension is used by the add and update commands.

srcaddr <address>

src <address>

srcaddr <address> and src <address> are synonyms that indicate the source address of the SA. If unspecified, the source address will either remain unset, or it will be set to a wildcard address if a destination address was supplied. This is valid for IPsec SAs. Future SA types may alter this assumption. This extension is used by the add, update, get and delete commands.

dstaddr <addr>

dst <addr>

dstaddr <addr> and dst <addr> are synonyms that indicate the destination address of the SA. If unspecified, the destination address will remain unset. Because IPsec SAs require a specified destination address and spi for identification, this extension, with a specific value, is required for the add, update, get and delete commands.

ipseckey(1M)

proxyaddr <address>

proxy <address>

proxyaddr <address> and proxy <address> are synonyms that indicate the proxy address for the SA. A proxy address is used for an SA that is protecting an inner protocol header. The proxy address is the source address of the inner protocol's header. This extension is used by the add and update commands.

authkey <hexstring>

Specifies the authentication key for this SA. The key is expressed as a string of hexadecimal digits, with an optional / at the end, for example, 123/12. Bits are counted from the most-significant bits down. For example, to express three '1' bits, the proper syntax is the string "e/3". For multi-key algorithms, the string is the concatenation of the multiple keys. This extension is used by the add and update commands.

encrkey <hexstring>

Specifies the encryption key for this SA. The syntax of the key is the same as authkey. A concrete example of a multi-key encryption algorithm is 3des, which would express itself as a 192-bit key, which is three 64-bit parity-included DES keys. This extension is used by the add and update commands.

Keying material is very sensitive and should be generated as randomly as possible. Some algorithms have known weak keys. IPsec algorithms have built-in weak key checks, so that if a weak key is in a newly added SA, the add command will fail with an invalid value.

Certificate identities are very useful in the context of automated key management, as they tie the SA to the public key certificates used in most automated key management protocols. They are less useful for manually added SAs. Unlike other extensions, srcidtype takes two values, a type, and an actual value. The type can be one of the following:

prefix	An address prefix.
fqdn	A fully-qualified domain name.
domain	Domain name, synonym for fqdn.
user_fqdn	User identity of the form user@fqdn.
mailbox	Synonym for user_fqdn.

The value is an arbitrary text string, which should identify the certificate.

srcidtype <type, value>

Specifies a source certificate identity for this SA. This extension is used by the add and update commands.

dstidtype <type, value>

Specifies a destination certificate identity for this SA. This extension is used by the add and update commands

**SECURITY
CONSIDERATIONS**

The ipseckey command allows a privileged user to enter cryptographic keying information. If an adversary gains access to such information, the security of IPsec traffic is compromised. The following issues should be taken into account when using the ipseckey command.

1. Is the TTY going over a network (interactive mode)?
 - If it is, then the security of the keying material is the security of the network path for this TTY's traffic. Using ipseckey over a clear-text telnet or rlogin session is risky.
 - Even local windows may be vulnerable to attacks where a concealed program that reads window events is present.
2. Is the file accessed over the network or readable to the world (-f option)?
 - A network-mounted file can be sniffed by an adversary as it is being read. A world-readable file with keying material in it is also risky.

If your source address is a host that can be looked up over the network, and your naming system itself is compromised, then any names used will no longer be trustworthy.

Security weaknesses often lie in misapplication of tools, not the tools themselves. Administrators are urged to be cautious when using ipseckey. The safest mode of operation is probably on a console, or other hard-connected TTY.

For further thoughts on this subject, see the afterward by Matt Blaze in Bruce Schneier's *Applied Cryptography: Protocols, Algorithms, and Source Code in C*.

EXAMPLES**EXAMPLE 1** Emptying Out All SAs

To empty out all SA:

```
example# ipseckey flush
```

EXAMPLE 2 Flushing Out IPsec AH SAs Only

To flush out only IPsec AH SAs:

```
example# ipseckey flush ah
```

EXAMPLE 3 Saving All SAs To Standard Output

To save all SAs to the standard output:

```
example# ipseckey save all
```

EXAMPLE 4 Saving ESP SAs To The File /tmp/snapshot

To save ESP SAs to the file /tmp/snapshot::

```
example# ipseckey save esp /tmp/snapshot
```

EXAMPLE 4 Saving ESP SAs To The File /tmp/snapshot (Continued)

EXAMPLE 5 Deleting an IPsec SA

To delete an IPsec SA, only the SPI and the destination address are needed:

```
example# ipseckey delete esp spi 0x2112 dst 224.0.0.1
```

EXAMPLE 6 Getting Information on an IPsec SA

Likewise, getting information on a SA only requires the destination address and SPI:

```
example# ipseckey get ah spi 0x5150 dst mypeer
```

EXAMPLE 7 Adding or Updating IPsec SAs

Adding or updating SAs requires entering interactive mode:

```
example# ipseckey
ipseckey> add ah spi 0x90125 src me.domain.com dst you.domain.com \
    authalg md5 authkey 1234567890abcdef1234567890abcdef
ipseckey> update ah spi 0x90125 dst you.domain.com hard_bytes \
    16000000
ipseckey> exit
```

EXAMPLE 8 Adding an SA in the Opposite Direction

In the case of IPsec, SAs are unidirectional. To communicate securely, a second SA needs to be added in the opposite direction. The peer machine also needs to add both SAs.

```
example# ipseckey
ipseckey> add ah spi 0x2112 src you.domain.com dst me.domain.com \
    authalg md5 authkey bde359723576fdea08e56cbe876e24ad \
    hard_bytes 16000000
ipseckey> exit
```

EXAMPLE 9 Monitoring PF_KEY Messages

Monitoring for PF_KEY messages is straightforward:

```
example# ipseckey monitor
```

EXAMPLE 10 Using Commands in a File

Commands can be placed in a file that can be parsed with the -f option. This file may contain comment lines that begin with the “#” symbol. For example:

```
# This is a sample file for flushing out the ESP table and
# adding a pair of SAs.

flush esp
```


EXAMPLE 10 Using Commands in a File (Continued)

```
### Watch out! I have keying material in this file. See the
### SECURITY CONSIDERATIONS section in this manual page for why this can be
### dangerous .

add esp spi 0x2112 src me.domain.com dst you.domain.com \
    authalg md5 authkey bde359723576fdea08e56cbe876e24ad \
    encralg des encrkey be02938e7def2839 hard_usetime 28800
add esp spi 0x5150 src you.domain.com dst me.domain.com \
    authalg md5 authkey 930987dbe09743ade09d92b4097d9e93 \
    encralg des encrkey 8bd4a52e10127deb hard_usetime 28800

## End of file - This is a gratuitous comment
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu
Interface Stability	Evolving

SUMMARY OF
TRUSTED
SOLARIS
CHANGES

Trusted Solaris 8
4/01 Reference
Manual

SUNWcsu
Reference Manual

This command requires the `sys_net_config` privilege. The label in the `exec_attr(4)` profile needs to match the label assigned to files read by this command.

ipsecconf(1M), route(1M)

ps(1), attributes(5), ipsec(7P), ipsecah(7P), ipsecesp(7P), pf_key(7P)

Schneier, B., *Applied Cryptography: Protocols, Algorithms, and Source Code in C*. Second ed. New York, New York: John Wiley & Sons, 1996.

DIAGNOSTICS

Parse error on line N.

If an interactive use of ipseckey would print usage information, this would print instead. Usually proceeded by another diagnostic.

Unexpected end of command line.

An additional argument was expected on the command line.

Unknown

A value for a specific extension was unknown.

Address type N not supported.

A name-to-address lookup returned an unsupported address family.

is not a bit specifier

bit length N is too big for

string is not a hex string

Keying material was not entered appropriately.

ipseckey(1M)

Can only specify single
A duplicate extension was entered.

Don't use extension for <string> for <command>.
An extension not used by a command was used.

NOTES In spite of its IPsec-specific name, `ipseckey` is analogous to `route(1M)`, in that it is a command-line interface to a socket-based administration engine, in this case, `PF_KEY`. `PF_KEY` was originally developed at the United States Naval Research Laboratory.

To have machines communicate securely with manual keying, SAs need to be added by all communicating parties. If two nodes wish to communicate securely, both nodes need the appropriate SAs added.

If the `-n` flag is not used when saving SAs, the resulting name for an address may not directly map to the address of an SA. In the future `ipseckey` may be invoked under additional names as other security protocols become available to `PF_KEY`.

NAME	lockd – Network lock daemon						
SYNOPSIS	<code>/usr/lib/nfs/lockd [-g <i>graceperiod</i>] [-t <i>timeout</i>] [<i>nthreads</i>]</code>						
DESCRIPTION	<p>The <code>lockd</code> utility is part of the NFS lock manager, which supports record locking operations on NFS files. See <code>fcntl(2)</code> and <code>lockf(3C)</code>. The lock manager provides two functions:</p> <ul style="list-style-type: none">■ it forwards <code>fcntl(2)</code> locking requests for NFS mounted file systems to the lock manager on the NFS server■ it generates local file locking operations in response to requests forwarded from lock managers running on NFS client machines. <p>State information kept by the lock manager about these locking requests can be lost if the <code>lockd</code> is killed or the operating system is rebooted. Some of this information can be recovered as follows. When the server lock manager restarts, it waits for a grace period for all client-site lock managers to submit reclaim requests. Client-site lock managers, on the other hand, are notified by the status monitor daemon, <code>statd(1M)</code>, of the restart and promptly resubmit previously granted lock requests. If the lock daemon fails to secure a previously granted lock at the server site, then it sends <code>SIGLOST</code> to a process.</p>						
OPTIONS	<table><tr><td><code>-g <i>graceperiod</i></code></td><td>Specify the number of seconds that clients have to reclaim locks after the server reboots. The default is 45 seconds.</td></tr><tr><td><code>-t <i>timeout</i></code></td><td>Specify the number of seconds to wait before retransmitting a lock request to the remote server. The default value is 15 seconds.</td></tr><tr><td><code><i>nthreads</i></code></td><td>Specify the maximum number of concurrent threads that the server can handle. This concurrency is achieved by up to <code><i>nthreads</i></code> threads created as needed in the kernel. <code><i>nthreads</i></code> should be based on the load expected on this server. If <code><i>nthreads</i></code> is not specified, the maximum number of concurrent threads will default to 20.</td></tr></table>	<code>-g <i>graceperiod</i></code>	Specify the number of seconds that clients have to reclaim locks after the server reboots. The default is 45 seconds.	<code>-t <i>timeout</i></code>	Specify the number of seconds to wait before retransmitting a lock request to the remote server. The default value is 15 seconds.	<code><i>nthreads</i></code>	Specify the maximum number of concurrent threads that the server can handle. This concurrency is achieved by up to <code><i>nthreads</i></code> threads created as needed in the kernel. <code><i>nthreads</i></code> should be based on the load expected on this server. If <code><i>nthreads</i></code> is not specified, the maximum number of concurrent threads will default to 20.
<code>-g <i>graceperiod</i></code>	Specify the number of seconds that clients have to reclaim locks after the server reboots. The default is 45 seconds.						
<code>-t <i>timeout</i></code>	Specify the number of seconds to wait before retransmitting a lock request to the remote server. The default value is 15 seconds.						
<code><i>nthreads</i></code>	Specify the maximum number of concurrent threads that the server can handle. This concurrency is achieved by up to <code><i>nthreads</i></code> threads created as needed in the kernel. <code><i>nthreads</i></code> should be based on the load expected on this server. If <code><i>nthreads</i></code> is not specified, the maximum number of concurrent threads will default to 20.						
ATTRIBUTES	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWcsu</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu		
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWcsu						
SUMMARY OF TRUSTED SOLARIS CHANGES	<p><code>lockd</code> must be started with a UID of 0, a sensitivity label of <code>ADMIN_LOW</code>, and a clearance of <code>ADMIN_HIGH</code>. It must be started from the trusted path and must have these privileges: <code>net_mac_read</code>, <code>net_privaddr</code>, <code>net_upgrade_sl</code>, and <code>sys_nfs</code>.</p>						
Trusted Solaris 8 4/01 Reference Manual	<p><code>statd(1M)</code>, <code>fcntl(2)</code></p> <p><code>lockf(3C)</code>, <code>attributes(5)</code></p>						

lpadmin(1M)

NAME	lpadmin – Configure the LP print service
SYNOPSIS	<p>lpadmin -p <i>printer options</i></p> <p>lpadmin -x <i>dest</i></p> <p>lpadmin -d [<i>dest</i>]</p> <p>lpadmin -S <i>print-wheel</i> -A <i>alert-type</i> [-W <i>minutes</i>] [-Q <i>requests</i>]</p> <p>lpadmin -M -f <i>form-name</i> [-a [-o <i>filebreak</i>] [-t <i>tray-number</i>]]</p>
DESCRIPTION	<p>lpadmin configures the LP print service by defining printers and devices. It is used to add and change printers, to remove printers from service, to set or change the system default destination, to define alerts for printer faults, and to mount print wheels.</p>
Adding or Changing a Printer	<p>The first form of the lpadmin command (lpadmin -p <i>printer options</i>) is used to configure a new printer or to change the configuration of an existing printer. When creating a new printer, one of three options (-v, -U, or -s) must be supplied. In addition, only one of the following may be supplied: -e, -i, or -m; if none of these three options is supplied, the model standard is used. The -h and -l options are mutually exclusive. Printer and class names may be no longer than 14 characters and must consist entirely of the characters A-Z, a-z, 0-9, dash (-) and underscore (_). If -s is specified, the following options are invalid: -A, -e, -F, -h, -i, -l, -M, -m, -o, -U, -v, and -W.</p> <p>The following printer options may appear in any order.</p> <p>-A <i>alert-type</i> [-W <i>minutes</i>]</p> <p>The -A option is used to define an alert that informs the administrator when a printer fault is detected, and periodically thereafter, until the printer fault is cleared by the administrator. The <i>alert-types</i> are:</p> <p>mail</p> <p>Send the alert message using mail (see mail(1)) to the administrator.</p> <p>write</p> <p>Write the message to the terminal on which the administrator is logged in. If the administrator is logged in on several terminals, one is chosen arbitrarily.</p> <p>quiet</p> <p>Do not send messages for the current condition. An administrator can use this option to temporarily stop receiving further messages about a known problem. Once the fault has been cleared and printing resumes, messages will again be sent when another fault occurs with the printer.</p> <p>showfault</p> <p>Attempt to execute a fault handler on each system that has a print job in the queue. The fault handler is /etc/lp/alerts/printer. It is invoked with three parameters: <i>printer_name</i>, <i>date</i>, and <i>file_name</i>. The <i>file_name</i> is the name of a file containing the fault message.</p>

none

Do not send messages; any existing alert definition for the printer will be removed. No alert will be sent when the printer faults until a different alert-type (except quiet) is used.

shell-command

Run the *shell-command* each time the alert needs to be sent. The shell command should expect the message in standard input. If there are blank spaces embedded in the command, enclose the command in quotes. Note that the `mail` and `write` values for this option are equivalent to the values `mail user-name` and `write user-name` respectively, where *user-name* is the current name for the administrator. This will be the login name of the person submitting this command unless he or she has used the `su` command to change to another user ID. If the `su` command has been used to change the user ID, then the *user-name* for the new ID is used.

list

Display the type of the alert for the printer fault. No change is made to the alert.

The message sent appears as follows:

```
The printer printer has stopped printing for the reason given below.
Fix the problem and bring the printer back on line.
Printing has stopped, but will be restarted in a few minutes;
issue an enable command if you want to restart sooner.
```

Unless someone issues the change request:

```
lp -i request-id -P ...to change the page list to print, the current request will be
reprinted from the beginning. The reason(s) it stopped (multiple reasons indicate
reprinted attempts):reason
```

The LP print service can detect printer faults only through an adequate fast filter and only when the standard interface program or a suitable customized interface program is used. Furthermore, the level of recovery after a fault depends on the capabilities of the filter.

If the *printer* is `all`, the alerting defined in this command applies to all existing printers.

If the `-W` option is not used to arrange fault alerting for *printer*, the default procedure is to mail one message to the administrator of *printer* per fault. This is equivalent to specifying `-W once` or `-W 0`. If *minutes* is a number greater than zero, an alert will be sent at intervals specified by *minutes*.

-c class

Insert *printer* into the specified *class*. *class* will be created if it does not already exist.

lpadmin(1M)

-D *comment*

Save this *comment* for display whenever a user asks for a full description of *printer* `lpstat(1)`. The LP print service does not interpret this comment.

-e *printer*

Copy the interface program of an existing *printer* to be the interface program for *printer*. (Options `-i` and `-m` may not be specified with this option.)

-F *fault-recovery*

This option specifies the recovery to be used for any print request that is stopped because of a printer fault, according to the value of *fault-recovery*:

continue	Continue printing on the top of the page where printing stopped. This requires a filter to wait for the fault to clear before automatically continuing.
----------	---

beginning	Start printing the request again from the beginning.
-----------	--

wait	Disable printing on <i>printer</i> and wait for the administrator or a user to enable printing again.
------	---

During the wait, the administrator or the user who submitted the stopped print request can issue a change request that specifies where printing should resume. (See the `-i` option of the `lp` command.) If no change request is made before printing is enabled, printing resumes at the top of the page where stopped, if the filter allows; otherwise, the request is printed from the beginning.

-f *allow:form-list*

-f *deny:form-list*

Allow or deny the forms in *form-list* to be printed on *printer*. By default no forms are allowed on a new printer.

For each printer, the LP print service keeps two lists of forms: an “allow-list” of forms that may be used with the printer, and a “deny-list” of forms that may not be used with the printer. With the `-f allow` option, the forms listed are added to the allow-list and removed from the deny-list. With the `-f deny` option, the forms listed are added to the deny-list and removed from the allow-list.

If the allow-list is not empty, only the forms in the list may be used on the printer, regardless of the contents of the deny-list. If the allow-list is empty, but the deny-list is not, the forms in the deny-list may not be used with the printer. All forms can be excluded from a printer by specifying `-f deny:all`. All forms can be used on a printer (provided the printer can handle all the characteristics of each form) by specifying `-f allow:all`.

The LP print service uses this information as a set of guidelines for determining where a form can be mounted. Administrators, however, are not restricted from mounting a form on any printer. If mounting a form on a particular printer is in

disagreement with the information in the allow-list or deny-list, the administrator is warned but the mount is accepted. Nonetheless, if a user attempts to issue a print or change request for a form and printer combination that is in disagreement with the information, the request is accepted only if the form is currently mounted on the printer. If the form is later unmounted before the request can print, the request is canceled and the user is notified by mail.

If the administrator tries to specify a form as acceptable for use on a printer that doesn't have the capabilities needed by the form, the command is rejected.

Note the other use of `-f`, with the `-M` option, below.

The `-T` option must be invoked first with `lpadmin` to identify the printer type before the `-f` option can be used.

`-h`

Indicate that the device associated with the printer is hardwired. If neither of the mutually exclusive options, `-h` and `-l`, is specified, `-h` is assumed.

`-I content-type-list`

Allow *printer* to handle print requests with the content types listed in a *content-type-list*. If the list includes names of more than one type, the names must be separated by commas or blank spaces. (If they are separated by blank spaces, the entire list must be enclosed in double quotes.)

The type `simple` is recognized as the default content type for files in the UNIX system. A `simple` type of file is a data stream containing only printable ASCII characters and the following control characters.

Control Character		Octal Value	Meaning
BACKSPACE	10		move back one character, except at beginning of line
TAB	11		move to next tab stop
LINEFEED (newline)	12		move to beginning of next line
FORMFEED	14		move to beginning of next page
RETURN	15		move to beginning of current line

To prevent the print service from considering `simple` a valid type for the printer, specify either an explicit value (such as the printer type) in the *content-type-list*, or an empty list. If you do want `simple` included along with other types, you must include `simple` in the *content-type-list*.

Except for `simple`, each *content-type* name is freely determined by the administrator. If the printer type is specified by the `-T` option, then the printer type is implicitly considered to be also a valid content type.

`-i interface`

Establish a new interface program for *printer*. *interface* is the pathname of the new program. (The `-e` and `-m` options may not be specified with this option.)

-l

Indicate that the device associated with *printer* is a login terminal. The LP scheduler (lpsched) disables all login terminals automatically each time it is started. (The *-h* option may not be specified with this option.)

-M -f *form-name* [-a [-o *filebreak*]] [-t *tray-number*]

Mount the form *form-name* on *printer*. Print requests that need the pre-printed form *form-name* will be printed on *printer*. If more than one printer has the form mounted and the user has specified any (with the *-d* option of the *lp* command) as the printer destination, then the print request will be printed on the one printer that also meets the other needs of the request.

The page length and width, and character and line pitches needed by the form are compared with those allowed for the printer, by checking the capabilities in the *terminfo* database for the type of printer. If the form requires attributes that are not available with the printer, the administrator is warned but the mount is accepted. If the form lists a print wheel as mandatory, but the print wheel mounted on the printer is different, the administrator is also warned but the mount is accepted.

If the *-a* option is given, an alignment pattern is printed, preceded by the same initialization of the physical printer that precedes a normal print request, with one exception: no banner page is printed. Printing is assumed to start at the top of the first page of the form. After the pattern is printed, the administrator can adjust the mounted form in the printer and press return for another alignment pattern (no initialization this time), and can continue printing as many alignment patterns as desired. The administrator can quit the printing of alignment patterns by typing *q*.

If the *-o filebreak* option is given, a formfeed is inserted between each copy of the alignment pattern. By default, the alignment pattern is assumed to correctly fill a form, so no formfeed is added.

If the *-t tray-number* option is specified, printer tray *tray-number* will be used.

A form is "unmounted" either by mounting a new form in its place or by using the *-f none* option. By default, a new printer has no form mounted.

Note the other use of *-f* without the *-M* option above.

-M -S *print-wheel*

Mount the *print-wheel* on *printer*. Print requests that need the *print-wheel* will be printed on *printer*. If more than one printer has *print-wheel* mounted and the user has specified any (with the *-d* option of the *lp* command) as the printer destination, then the print request will be printed on the one printer that also meets the other needs of the request.

If the *print-wheel* is not listed as acceptable for the printer, the administrator is warned but the mount is accepted. If the printer does not take print wheels, the command is rejected.

A print wheel is “unmounted” either by mounting a new print wheel in its place or by using the option `-S none`. By default, a new printer has no print wheel mounted.

Note the other uses of the `-S` option without the `-M` option described below.

`-m model`

Select *model* interface program, provided with the LP print service, for the printer. (Options `-e` and `-i` may not be specified with this option.)

`-o option`

The `-o` option defines default printer configuration values given to an interface program. The default may be explicitly overwritten for individual requests by the user (see `lp(1)`), or taken from a preprinted form description (see `lpforms(1M)` and `lp(1)`).

There are several options which are pre-defined by the system. In addition, any number of key-value pairs may be defined. Each of the predefined and undefined options are described.

The Predefined Options

The following options are predefined: adjusting printer capabilities, adjusting printer port characteristics, configuring network printers, and controlling the use of banner.

Adjusting Printer Capabilities

```
length=scaled-decimal-number
width=scaled-decimal-number
cpi=scaled-decimal-number
lpi=scaled-decimal-number
```

The term *scaled-decimal-number* refers to a non-negative number used to indicate a unit of size. The type of unit is shown by a “trailing” letter attached to the number. Three types of *scaled-decimal-numbers* can be used with the LP print service: numbers that show sizes in centimeters (marked with a trailing `c`); numbers that show sizes in inches (marked with a trailing `i`); and numbers that show sizes in units appropriate to use (without a trailing letter), that is, lines, characters, lines per inch, or characters per inch.

The option values must agree with the capabilities of the type of physical printer, as defined in the terminfo database for the printer type. If they do not, the command is rejected.

The defaults are defined in the `terminfo` entry for the specified printer type. The defaults may be reset by:

```
lpadmin -p printername -o length=
lpadmin -p printername o width=
lpadmin -p printername o cpi=
lpadmin -p printername o lpi=
```

Adjusting Printer Port Characteristics

```
stty=" ' stty-option-list' "
```

The *stty-option-list* is not checked for allowed values, but is passed directly to the *stty* program by the standard interface program. Any error messages produced by *stty* when a request is processed (by the standard interface program) are mailed to the user submitting the request.

The default for *stty* is:

```
stty="'9600 cs8 -cstopb -p arenb ixon
-i xany opost -o lcuc onlcr
-o crnl -o nocr
-o nlret -o fill nl0 cr0 tab0 bs0 vt0 ff0'"
```

The default may be reset by:

```
lpadmin -p printername -o stty=
```

Configuring Network Printers

```
dest=string
protocol=string
bsdctrl=string
timeout=non-negative-integer-seconds
```

These four options are provided to support network printing. Each option is passed directly to the interface program; any checking for allowed values is done there.

The value of *dest* is the name of the destination for the network printer; the semantics for value *dest* are dependent on the printer and the configuration. There is no default.

The value of option *protocol* sets the over-the-wire protocol to the printer. The default for option *protocol* is *bsd*. The value of option *bsdctrl* sets the print order of control and data files (BSD protocol only); the default for this option is *control file first*. The value of option *timeout* sets the seed value for backoff time when the printer is busy. The default value for the *timeout* option is 10 seconds. The defaults may be reset by:

```
lpadmin -p printername -o protocol=
lpadmin -p printername -o bsdctrl=
lpadmin -p printername -o timeout=
```

Controlling the Use of the Banner Page

nobanner

Allow a user to submit a print request specifying that no banner page be printed.

banner

Force a banner page to be printed with every print request, even when a user asks for no banner page. This is the default. Specify *-o nobanner* to allow

users to specify `-o nobanner` with the `lp` command. Undefined Options

key=value

Each *key=value* is passed directly to the interface program. Any checking for allowed values is done in the interface program.

Any default values for a given *key=value* option are defined in the interface program. If a default is provided, it may be reset by typing the key without any value:

```
lpadmin -p printername -o key=
```

`-P paper-name`

Specify a paper type list that the printer supports.

`-r class`

Remove *printer* from the specified *class*. If *printer* is the last member of *class*, then *class* will be removed.

`-S list`

Allow either the print wheels or aliases for character sets named in *list* to be used on the printer.

If the printer is a type that takes print wheels, then *list* is a comma or space separated list of print wheel names. (Enclose the list with quotes if it contains blank spaces.) These will be the only print wheels considered mountable on the printer. (You can always force a different print wheel to be mounted.) Until the option is used to specify a list, no print wheels will be considered mountable on the printer, and print requests that ask for a particular print wheel with this printer will be rejected.

If the printer is a type that has selectable character sets, then *list* is a comma or blank separated list of character set name "mappings" or aliases. (Enclose the list with quotes if it contains blank spaces.) Each "mapping" is of the form *known-name=alias*. The *known-name* is a character set number preceded by *cs* (such as *cs3* for character set three) or a character set name from the `terminfo` database entry *csnm*. See `terminfo(4)`. If this option is not used to specify a list, only the names already known from the `terminfo` database or numbers with a prefix of *cs* will be acceptable for the printer. If *list* is the word *none*, any existing print wheel lists or character set aliases will be removed.

Note the other uses of the `-S` with the `-M` option described above.

The `-T` option must be invoked first with `lpadmin` to identify the printer type before the `-S` option can be used.

`-s system-name[!printer-name]`

Make a remote printer (one that must be accessed through another system) accessible to users on your system. *system-name* is the name of the remote system on which the remote printer is located it. *printer-name* is the name used on the remote

lpadmin(1M)

system for that printer. For example, if you want to access *printer1* on *system1* and you want it called *printer2* on your system:

```
-p printer2 -s system1!printer1
```

-T *printer-type-list*

Identify the printer as being of one or more *printer-types*. Each *printer-type* is used to extract data from the `terminfo` database; this information is used to initialize the printer before printing each user's request. Some filters may also use a *printer-type* to convert content for the printer. If this option is not used, the default *printer-type* will be unknown; no information will be extracted from `terminfo` so each user request will be printed without first initializing the printer. Also, this option must be used if the following are to work: `-o cpi`, `-o lpi`, `-o width`, and `-o length` options of the `lpadmin` and `lp` commands, and the `-S` and `-f` options of the `lpadmin` command.

If the *printer-type-list* contains more than one type, then the *content-type-list* of the `-I` option must either be specified as `simple`, as empty (`-I ""`), or not specified at all.

-t *number-of-trays*

Specify the number of trays when creating the printer.

-u *allow:login-ID-list*

-u *deny:login-ID-list*

Allow or deny the users in *login-ID-list* access to the printer. By default all users are allowed on a new printer. The *login-ID-list* argument may include any or all of the following constructs:

<i>login-ID</i>	a user on any system
<i>system-name</i> ! <i>login-ID</i>	a user on system <i>system-name</i>
<i>system-name</i> !all	all users on system <i>system-name</i>
all! <i>login-ID</i>	a user on all systems
all	all users on all systems

For each printer, the LP print service keeps two lists of users: an "allow-list" of people allowed to use the printer, and a "deny-list" of people denied access to the printer. With the `-u allow` option, the users listed are added to the allow-list and removed from the deny-list. With the `-u deny` option, the users listed are added to the deny-list and removed from the allow-list.

If the allow-list is not empty, only the users in the list may use the printer, regardless of the contents of the deny-list. If the allow-list is empty, but the deny-list is not, the users in the deny-list may not use the printer. All users can be denied access to the printer by specifying `-u deny:all`. All users may use the printer by specifying `-u allow:all`.

**Removing a
Printer Destination****Setting/Changing
the System Default
Destination****Setting an Alert
for a Print Wheel****-U *dial-info***

The -U option allows your print service to access a remote printer. (It does not enable your print service to access a remote printer service.) Specifically, -U assigns the “dialing” information *dial-info* to the printer. *dial-info* is used with the `dial` routine to call the printer. Any network connection supported by the Basic Networking Utilities will work. *dial-info* can be either a phone number for a modem connection, or a system name for other kinds of connections. Or, if -U *direct* is given, no dialing will take place, because the name *direct* is reserved for a printer that is directly connected. If a system name is given, it is used to search for connection details from the file `/etc/uucp/Systems` or related files. The Basic Networking Utilities are required to support this option. By default, -U *direct* is assumed.

-v *device*

Associate a *device* with *printer*. *device* is the path name of a file that is writable by lp. Note that the same *device* can be associated with more than one printer.

The -x *dest* option removes the destination *dest* (a printer or a class), from the LP print service. If *dest* is a printer and is the only member of a class, then the class will be deleted, too. If *dest* is `all`, all printers and classes are removed. No other *options* are allowed with -x.

The -d [*dest*] option makes *dest* (an existing printer or class) the new system default destination. If *dest* is not supplied, then there is no system default destination. No other *options* are allowed with -d.

-S *print-wheel* -A *alert-type* [-W *minutes*] [-Q *requests*]

The -S *print-wheel* option is used with the -A *alert-type* option to define an alert to mount the print wheel when there are jobs queued for it. If this command is not used to arrange alerting for a print wheel, no alert will be sent for the print wheel. Note the other use of -A, with the -p option, above.

The *alert-types* are:

mail	Send the alert message using the <code>mail</code> command to the administrator.
write	Write the message, using the <code>write</code> command, to the terminal on which the administrator is logged in. If the administrator is logged in on several terminals, one is arbitrarily chosen.
quiet	Do not send messages for the current condition. An administrator can use this option to temporarily stop receiving further messages about a known problem. Once the <i>print-wheel</i> has been mounted and subsequently unmounted, messages will again be sent when the number of print requests reaches the threshold specified by the -Q option.
none	Do not send messages until the -A option is given again with a different <i>alert-type</i> (other than <code>quiet</code>).

lpadmin(1M)

	<i>shell-command</i>	Run the <i>shell-command</i> each time the alert needs to be sent. The shell command should expect the message in standard input. If there are blanks embedded in the command, enclose the command in quotes. Note that the <code>mail</code> and <code>write</code> values for this option are equivalent to the values <code>mail user-name</code> and <code>write user-name</code> respectively, where <i>user-name</i> is the current name for the administrator. This will be the login name of the person submitting this command unless he or she has used the <code>su</code> command to change to another user ID. If the <code>su</code> command has been used to change the user ID, then the <i>user-name</i> for the new ID is used.				
	<i>list</i>	Display the type of the alert for the print wheel on standard output. No change is made to the alert.				
	<p>The message sent appears as follows:</p> <pre>The print wheel <i>print-wheel</i> needs to be mounted on the printer(s): printer (<i>integer1</i>requests) <i>integer2</i> print requests await this print wheel.</pre> <p>The printers listed are those that the administrator had earlier specified were candidates for this print wheel. The number <i>integer1</i> listed next to each printer is the number of requests eligible for the printer. The number <i>integer2</i> shown after the printer list is the total number of requests awaiting the print wheel. It will be less than the sum of the other numbers if some requests can be handled by more than one printer.</p> <p>If the <i>print-wheel</i> is <code>all</code>, the alerting defined in this command applies to all print wheels already defined to have an alert.</p> <p>If the <code>-W</code> option is not given, the default procedure is that only one message will be sent per need to mount the print wheel. Not specifying the <code>-W</code> option is equivalent to specifying <code>-W once</code> or <code>-W 0</code>. If <i>minutes</i> is a number greater than zero, an alert will be sent at intervals specified by <i>minutes</i>.</p> <p>If the <code>-Q</code> option is also given, the alert will be sent when a certain number (specified by the argument <i>requests</i>) of print requests that need the print wheel are waiting. If the <code>-Q</code> option is not given, or <i>requests</i> is 1 or any (which are both the default), a message is sent as soon as anyone submits a print request for the print wheel when it is not mounted.</p>					
EXIT STATUS	<p>The following exit values are returned:</p> <table><tr><td>0</td><td>Successful completion.</td></tr><tr><td>non-zero</td><td>An error occurred.</td></tr></table>		0	Successful completion.	non-zero	An error occurred.
0	Successful completion.					
non-zero	An error occurred.					
FILES	<table><tr><td>/var/spool/lp/*</td><td>LP print queue.</td></tr><tr><td>/etc/lp</td><td>Printing system control files.</td></tr></table>	/var/spool/lp/*	LP print queue.	/etc/lp	Printing system control files.	
/var/spool/lp/*	LP print queue.					
/etc/lp	Printing system control files.					

lpadmin(1M)

/etc/lp/alerts/printer Fault handler for lpadmin.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWpcu

SUMMARY OF TRUSTED SQUARES CHANGES
Trusted Solaris 4/01 Reference Manual
SunOS 5.8 Reference Manual

Use of the lpadmin command requires the administer printing authorization.

enable(1), lp(1), lpstat(1), accept(1M), lpforms(1M), lpsched(1M),
lpsystem(1M)

Trusted Solaris Administrator's Procedures

mail(1), stty(1), dial(3NSL), terminfo(4), attributes(5)

System Administration Guide, Volume 1

lpfilter(1M)

NAME	lpfilter – Administer filters used with the LP print service												
SYNOPSIS	<code>/usr/sbin/lpfilter -f filter-name { - -i -l -x -F pathname }</code>												
DESCRIPTION	The <code>lpfilter</code> command is used to add, change, delete, or list a filter used with the LP print service. These filters convert the content of a file to have a content type acceptable to a printer.												
OPTIONS	<p>Arguments consist of the <code>-f filter-name</code> option and exactly one of the arguments appearing within braces ({ }) in the SYNOPSIS.</p> <table> <tr> <td><code>-f filter-name</code></td><td>Specifies the <i>filter-name</i> of the filter to be added, changed, reset, deleted, or listed. The filter name <code>all</code> is a special filter name defined below. The <code>-f</code> option is required.</td></tr> <tr> <td><code>-</code></td><td>Adds or changes a filter as specified from standard input. The format of the input is specified below. If <code>-f all</code> is specified with the <code>-</code> option, the specified change is made to all existing filters. This is not useful.</td></tr> <tr> <td><code>-F pathname</code></td><td>Adds or changes a filter as specified by the contents of the file <i>pathname</i>. The format of the file's contents is specified below. If <code>-f all</code> is specified with the <code>-F</code> option, the specified change is made to all existing filters. This is not useful.</td></tr> <tr> <td><code>-i</code></td><td>Resets a filter to its default settings. Using <code>-f all</code> with the <code>-i</code> option restores all filters for which predefined settings are available to their original settings.</td></tr> <tr> <td><code>-x</code></td><td>Deletes a filter. Using <code>-f all</code> with the <code>-x</code> option results in all filters being deleted.</td></tr> <tr> <td><code>-l</code></td><td>Lists a filter description. Using <code>-f all</code> with the <code>-l</code> option produces a list of all filters.</td></tr> </table>	<code>-f filter-name</code>	Specifies the <i>filter-name</i> of the filter to be added, changed, reset, deleted, or listed. The filter name <code>all</code> is a special filter name defined below. The <code>-f</code> option is required.	<code>-</code>	Adds or changes a filter as specified from standard input. The format of the input is specified below. If <code>-f all</code> is specified with the <code>-</code> option, the specified change is made to all existing filters. This is not useful.	<code>-F pathname</code>	Adds or changes a filter as specified by the contents of the file <i>pathname</i> . The format of the file's contents is specified below. If <code>-f all</code> is specified with the <code>-F</code> option, the specified change is made to all existing filters. This is not useful.	<code>-i</code>	Resets a filter to its default settings. Using <code>-f all</code> with the <code>-i</code> option restores all filters for which predefined settings are available to their original settings.	<code>-x</code>	Deletes a filter. Using <code>-f all</code> with the <code>-x</code> option results in all filters being deleted.	<code>-l</code>	Lists a filter description. Using <code>-f all</code> with the <code>-l</code> option produces a list of all filters.
<code>-f filter-name</code>	Specifies the <i>filter-name</i> of the filter to be added, changed, reset, deleted, or listed. The filter name <code>all</code> is a special filter name defined below. The <code>-f</code> option is required.												
<code>-</code>	Adds or changes a filter as specified from standard input. The format of the input is specified below. If <code>-f all</code> is specified with the <code>-</code> option, the specified change is made to all existing filters. This is not useful.												
<code>-F pathname</code>	Adds or changes a filter as specified by the contents of the file <i>pathname</i> . The format of the file's contents is specified below. If <code>-f all</code> is specified with the <code>-F</code> option, the specified change is made to all existing filters. This is not useful.												
<code>-i</code>	Resets a filter to its default settings. Using <code>-f all</code> with the <code>-i</code> option restores all filters for which predefined settings are available to their original settings.												
<code>-x</code>	Deletes a filter. Using <code>-f all</code> with the <code>-x</code> option results in all filters being deleted.												
<code>-l</code>	Lists a filter description. Using <code>-f all</code> with the <code>-l</code> option produces a list of all filters.												
Adding or Changing a Filter	<p>The filter named in the <code>-f</code> option is added to the filter table. If the filter already exists, its description is changed to reflect the new information in the input.</p> <p>When <code>-</code> is specified, standard input supplies the filter description. When <code>-F</code> is specified, the file <i>pathname</i> supplies the filter description. One of these two options must be specified to add or change a filter.</p> <p>When an existing filter is changed with the <code>-F</code> or <code>-</code> option, lines in the filter description that are not specified in the new information are not changed. When a new filter is added with this command, unspecified lines receive default values. See below.</p> <p>Filters are used to convert the content of a request from its initial type into a type acceptable to a printer. For a given print request, the LP print service knows the following:</p> <ul style="list-style-type: none"> ■ The content type of the request (specified by <code>lp -T</code> or determined implicitly) ■ The name of the printer (specified by <code>lp -d</code>) 												

- The printer type (specified by `lpadmin -T`)
The printer type is intended to be a printer model, but some people specify it with a content type even though `lpadmin -I` is intended for this purpose.
- The content types acceptable to the printer (specified by `lpadmin -I`)
The values specified by the `lpadmin -T` are treated as if they were specified by the `-I` option as well.
- The modes of printing asked for by the originator of the request (specified by various options to `lp`)

The system uses the above information to construct a list of one or more filters that converts the document's content type into a content type acceptable to the printer and consumes all `lp` arguments that invoke filters (`-Y` and `-P`).

The contents of the file (specified by the `-F` option) and the input stream from standard input (specified by `-`) must consist of a series of lines, such that each line conforms to the syntax specified by one of the seven lines below. All lists are comma or space separated. Each item contains a description.

Input types: *content-type-list*
 Output types: *content-type-list*
 Printer types: *printer-type-list*
 Printers: *printer-list*
 Filter type: *filter-type*
 Command: *shell-command*
 Options: *template-list*

Input types	This gives the content types that can be accepted by the filter. The default is any. The document content type must be a member of this list for the initial filter in the sequence.
Output types	This gives the content types that the filter can produce from any of the input (content) types. The default is any. The intersection of the output types of this list and the content types acceptable to the printer (from <code>lpadmin -I</code> and <code>lpadmin -T</code>) must be non-null for the last filter in the sequence. For adjacent filters in the sequence, the intersection of output types of one and the input types of the next must be non-null.
Printer types	This gives the printer types for which this printer can be used. The LP print service will restrict the use of the filter to these printer types (from <code>lpadmin -T</code>). The default is any.
Printers	This gives the names of the printers for which the filter can be used. The LP print service will restrict the use of the filter to just the printers named. The default is any.
Filter type	This marks the filter as a <code>slow</code> filter or a <code>fast</code> filter. Slow filters are generally those that take a long time to convert their input (that is, minutes or hours). They are run before the job is scheduled

lpfilter(1M)

	<p>for a printer, to keep the printers from being tied up while the filter is running. If a listed printer is on a remote system, the filter type for it must have the value <code>slow</code>. That is, if a client defines a filter, it must be a slow filter. Fast filters are generally those that convert their input quickly (that is, faster than the printer can process the data), or those that must be connected to the printer when run. Fast filters will be given to the interface program to run while connected to the physical printer.</p>
Command	<p>This specifies which program to run to invoke the filter. The full program pathname as well as fixed options must be included in the <i>shell-command</i>; additional options are constructed, based on the characteristics of each print request and on the <i>Options</i> field. A command must be given for each filter. The command must accept a data stream as standard input and produce the converted data stream on its standard output. This allows filter pipelines to be constructed to convert data not handled by a single filter.</p>
Options	<p>This is a comma-separated list of templates used by the LP print service to construct options to the filter from the characteristics of each print request listed in the table later. The <code>-y</code> and <code>-P</code> arguments to the <code>lp</code> command cause a filter sequence to be built even if there is no need for a conversion of content types.</p> <p>In general, each template is of the following form:</p> <p><i>keyword pattern = replacement</i></p> <p>The <i>keyword</i> names the characteristic that the template attempts to map into a filter-specific option; each valid <i>keyword</i> is listed in the table below.</p> <p>A <i>pattern</i> is one of the following: a literal pattern of one of the forms listed in the table, a single asterisk (*), or a regular expression. If <i>pattern</i> matches the value of the characteristic, the template fits and is used to generate a filter-specific option. The <i>replacement</i> is what will be used as the option.</p> <p>Regular expressions are the same as those found on the <code>regex(5)</code> manual page. This includes the <code>\(. . . \)</code> and <code>\n</code> constructions, which can be used to extract portions of the <i>pattern</i> for copying into the <i>replacement</i>, and the <code>&</code>, which can be used to copy the entire <i>pattern</i> into the <i>replacement</i>.</p> <p>The <i>replacement</i> can also contain a <code>*</code>; it too, is replaced with the entire <i>pattern</i>, just like the <code>&</code> of <code>regex(5)</code>.</p> <p>The keywords are:</p>

lp Option	Characteristic	keyword	Possible patterns
-T	Content type (input)	INPUT	content-type
not applicable	Content type (output)	OUTPUT	content-type
not applicable	Printer type	TERM	printer-type
-d	Printer name	PRINTER	<i>printer-name</i>
-f, -o cpi=	Character pitch	CPI	integer
-f, -o lpi=	Line pitch	LPI	integer
-f, -o length=	Page length	LENGTH	integer
-f, -o width=	Page width	WIDTH	integer
-P	Pages to print	PAGES	page-list
-S	Character set	CHARSET	character-set-name
-S	Print wheel	CHARSET	print-wheel-name
-f	Form name	FORM	form-name
-y	Modes	MODES	mode
-n	Number of copies	COPIES	<i>integer</i>

Large File Behavior

See `largefile(5)` for the description of the behavior of `lpfilter` when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).

EXAMPLES

EXAMPLE 1 Printing examples.

For example, the template

```
MODES landscape = -l
```

shows that if a print request is submitted with the `-y landscape` option, the filter will be given the option `-l`.

As another example, the template

```
TERM * = -T *
```

shows that the filter will be given the option `-T printer-type` for whichever *printer-type* is associated with a print request using the filter.

As a last example, consider the template

```
MODES prwidth=\(.*\) = -w\1
```

Suppose a user gives the command

```
lp -y prwidth=10
```

lpfilter(1M)

EXAMPLE 1 Printing examples. (Continued)

From the table above, the LP print service determines that the `-y` option is handled by a MODES template. The MODES template here works because the pattern `prwidth=)` matches the `prwidth=10` given by the user. The replacement `-w1` causes the LP print service to generate the filter option `-w10`. If necessary, the LP print service will construct a filter pipeline by concatenating several filters to handle the user's file and all the print options. See `sh(1)` for a description of a pipeline. If the print service constructs a filter pipeline, the INPUT and OUTPUT values used for each filter in the pipeline are the types of input and output for that filter, not for the entire pipeline.

Resetting a Filter to Defaults

If the filter named is one originally delivered with the LP print service, the `-i` option restores the original filter description.

Deleting a Filter

The `-x` option is used to delete the filter specified in `filter-name` from the LP filter table.

Listing a Filter Description

The `-l` option is used to list the description of the filter named in `filter-name`. If the command is successful, the following message is sent to standard output:

Input types: *content-type-list* Output types: *content-type-list*

Printer types: *printer-type-list* Printers: *printer-list*

Filter type: *filter-type* Command: *shell-command* Options: *template-list*

If the command fails, an error message is sent to standard error.

EXIT STATUS

The following exit values are returned:

0 Successful completion.

non-zero An error occurred.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWpsu

SUMMARY OF TRUSTED

Trusted Solaris 4/01 Reference Manual

Use of the `lpfilter` command requires the administer printing authorization.

`lp(1)`, `lpadmin(1M)`

Trusted Solaris Administrator's Procedures

SunOS 5.8 Reference Manual

`sh(1)`, `attributes(5)`, `largefile(5)`, `regexp(5)`

NOTES If the `lp` command specifies more than one document, the filtering chain is determined by the first document. Other documents may have a different format, but they will print correctly only if the filter chain is able to handle their format.

lpforms(1M)

NAME	lpforms – Administer forms used with the LP print service
SYNOPSIS	<p>lpforms -f <i>form-name</i> <i>option</i></p> <p>lpforms -f <i>form-name</i> -A <i>alert-type</i> [-P <i>paper-name</i> [-d]] [-Q <i>requests</i>] [-W <i>minutes</i>]</p>
DESCRIPTION	<p>The lpforms command administers the use of preprinted forms, such as company letterhead paper, with the LP print service. A form is specified by its <i>form-name</i>. Users may specify a form when submitting a print request (see lp(1)). The argument all can be used instead of <i>form-name</i> with either of the command lines shown above. The first command line allows the administrator to add, change, and delete forms, to list the attributes of an existing form, and to allow and deny users access to particular forms. The second command line is used to establish the method by which the administrator is alerted that the form must be mounted on a printer.</p>
OPTIONS	<p>The following options are supported:</p> <p>-f <i>form-name</i> Specify a form.</p> <p>The first form of lpforms requires that one of the following options (-, -l, -F, -x) must be used:</p> <p>-F <i>pathname</i> To add or change form <i>form-name</i>, as specified by the information in <i>pathname</i>.</p> <p>- To add or change form <i>form-name</i>, as specified by the information from standard input.</p> <p>-l To list the attributes of form <i>form-name</i>.</p> <p>-x To delete form <i>form-name</i> (this option must be used separately; it may not be used with any other option).</p> <p>The second form of the lpforms command requires the -A <i>alert-type</i> option. The other options are optional.</p> <p>-A <i>alert-type</i> Defines an alert to mount the form when there are queued jobs which need it.</p> <p>-P <i>paper-name</i> [-d] Specify the paper name when creating the form. If -d is specified, this paper is the default.</p> <p>-Q <i>requests</i> An alert will be sent when a certain number of print requests that need the form are waiting.</p> <p>-W <i>minutes</i> An alert will be sent at intervals specified by <i>minutes</i>.</p>
Adding or Changing a Form	<p>The -F <i>pathname</i> option is used to add a new form, <i>form-name</i>, to the LP print service, or to change the attributes of an existing form. The form description is taken from <i>pathname</i> if the -F option is given, or from the standard input if the - option is used. One of these two options must be used to define or change a form.</p>

pathname is the path name of a file that contains all or any subset of the following information about the form.

```

Page length: scaled-decimal-number1
Page width: scaled-decimal-number2
Number of pages: integer
Line pitch: scaled-decimal-number3
Character pitch: scaled-decimal-number4
Character set choice: character-set/print-wheel [mandatory]
Ribbon color: ribbon-color
Comment:
comment
Alignment pattern: [content-type]
content

```

The term “scaled-decimal-number” refers to a non-negative number used to indicate a unit of size. The type of unit is shown by a “trailing” letter attached to the number. Three types of scaled decimal numbers can be used with the LP print service: numbers that show sizes in centimeters (marked with a trailing *c*); numbers that show sizes in inches (marked with a trailing *i*); and numbers that show sizes in units appropriate to use (without a trailing letter); lines, characters, lines per inch, or characters per inch.

Except for the last two lines, the above lines may appear in any order. The *Comment :* and *comment* items must appear in consecutive order but may appear before the other items, and the *Alignment pattern:* and the *content* items must appear in consecutive order at the end of the file. Also, the *comment* item may not contain a line that begins with any of the key phrases above, unless the key phrase is preceded with a *>* sign. Any leading *>* sign found in the *comment* will be removed when the comment is displayed. There is no case distinction among the key phrases.

When this command is issued, the form specified by *form-name* is added to the list of forms. If the form already exists, its description is changed to reflect the new information. Once added, a form is available for use in a print request, except where access to the form has been restricted, as described under the *-u* option. A form may also be allowed to be used on certain printers only.

A description of each form attribute is below:

Page length and Page Width

Before printing the content of a print request needing this form, the generic interface program provided with the LP print service will initialize the physical printer to handle pages *scaled-decimal-number1* long, and *scaled-decimal-number2* wide using the printer type as a key into the `terminfo(4)` database. The page length and page width will also be passed, if possible, to each filter used in a request needing this form.

Number of pages

Each time the alignment pattern is printed, the LP print service will attempt to truncate the *content* to a single form by, if possible, passing to each filter the page subset of 1-*integer*.

Line pitch and Character pitch

Before printing the content of a print request needing this form, the interface program provided with the LP print service will initialize the physical printer to handle these pitches, using the printer type as a key into the `terminfo(4)` database. Also, the pitches will be passed, if possible, to each filter used in a request needing this form. *scaled-decimal-number3* is in lines-per-centimeter if a *c* is appended, and lines-per-inch otherwise; similarly, *scaled-decimal-number4* is in characters-per-centimeter if a *c* is appended, and characters-per-inch otherwise. The character pitch can also be given as *elite* (12 characters-per-inch), *pica* (10 characters-per-inch), or *compressed* (as many characters-per-inch as possible).

Character set choice

When the LP print service alerts an administrator to mount this form, it will also mention that the print wheel *print-wheel* should be used on those printers that take print wheels. If printing with this form is to be done on a printer that has selectable or loadable character sets instead of print wheels, the interface programs provided with the LP print service will automatically select or load the correct character set. If *mandatory* is appended, a user is not allowed to select a different character set for use with the form; otherwise, the character set or print wheel named is a suggestion and a default only.

Ribbon color

When the LP print service alerts an administrator to mount this form, it will also mention that the color of the ribbon should be *ribbon-color*.

Comment

The LP print service will display the *comment* unaltered when a user asks about this form (see `lpstat(1)`).

Alignment pattern

When mounting this form, an administrator can ask for the *content* to be printed repeatedly, as an aid in correctly positioning the preprinted form. The optional *content-type* defines the type of printer for which *content* had been generated. If *content-type* is not given, *simple* is assumed. Note that the *content* is stored as given, and will be readable only by the user *lp*.

When an existing form is changed with this command, items missing in the new information are left as they were. When a new form is added with this command, missing items will get the following defaults:

```
Page Length: 66
Page Width: 80
Number of Pages: 1
Line Pitch: 6
Character Pitch: 10
Character Set Choice: any
Ribbon Color: any
```

Deleting a Form

The *-x* option is used to delete the form *form-name* from the LP print service.

Listing Form Attributes	<p>The <code>-l</code> option is used to list the attributes of the existing form <i>form-name</i>. The attributes listed are those described under Adding and Changing a Form, above. Because of the potentially sensitive nature of the alignment pattern, only the administrator can examine the form with this command. Other people may use the <code>lpstat(1)</code> command to examine the non-sensitive part of the form description.</p>										
Allowing and Denying Access to a Form	<p>The <code>-u</code> option, followed by the argument <code>allow:login-ID-list</code> or <code>-u deny:login-ID-list</code> lets you determine which users will be allowed to specify a particular form with a print request. This option can be used with the <code>-F</code> or <code>-</code> option, each of which is described above under Adding or Changing a Form.</p> <p>The <i>login-ID-list</i> argument may include any or all of the following constructs:</p> <table> <tr> <td><i>login-ID</i></td><td>A user on any system</td></tr> <tr> <td><i>system_name</i>!<i>login-ID</i></td><td>A user on system <i>system_name</i></td></tr> <tr> <td><i>system_name</i>!all</td><td>All users on system <i>system_name</i></td></tr> <tr> <td>all!<i>login-ID</i></td><td>A user on all systems</td></tr> <tr> <td>all</td><td>All users on all systems</td></tr> </table> <p>The LP print service keeps two lists of users for each form: an “allow-list” of people allowed to use the form, and a “deny-list” of people that may not use the form. With the <code>-u allow</code> option, the users listed are added to the allow-list and removed from the deny-list. With the <code>-u deny</code> option, the users listed are added to the deny-list and removed from the allow-list. (Both forms of the <code>-u</code> option can be run together with the <code>-F</code> or the <code>-</code> option.)</p> <p>If the allow-list is not empty, only the users in the list are allowed access to the form, regardless of the content of the deny-list. If the allow-list is empty but the deny-list is not, the users in the deny-list may not use the form, (but all others may use it). All users can be denied access to a form by specifying <code>-f deny:all</code>. All users can be allowed access to a form by specifying <code>-f allow:all</code>. (This is the default.)</p>	<i>login-ID</i>	A user on any system	<i>system_name</i> ! <i>login-ID</i>	A user on system <i>system_name</i>	<i>system_name</i> !all	All users on system <i>system_name</i>	all! <i>login-ID</i>	A user on all systems	all	All users on all systems
<i>login-ID</i>	A user on any system										
<i>system_name</i> ! <i>login-ID</i>	A user on system <i>system_name</i>										
<i>system_name</i> !all	All users on system <i>system_name</i>										
all! <i>login-ID</i>	A user on all systems										
all	All users on all systems										
Setting an Alert to Mount a Form	<p>The <code>-f form-name</code> option is used with the <code>-A alert-type</code> option to define an alert to mount the form when there are queued jobs which need it. If this option is not used to arrange alerting for a form, no alert will be sent for that form.</p> <p>The method by which the alert is sent depends on the value of the <i>alert-type</i> argument specified with the <code>-A</code> option. The <i>alert-types</i> are:</p> <table> <tr> <td>mail</td><td>Send the alert message using the <code>mail</code> command to the administrator.</td></tr> <tr> <td>write</td><td>Write the message, using the <code>write</code> command, to the terminal on which the administrator is logged in. If the administrator is logged in on several terminals, one is arbitrarily chosen.</td></tr> <tr> <td>quiet</td><td>Do not send messages for the current condition. An administrator can use this option to temporarily stop receiving further messages</td></tr> </table>	mail	Send the alert message using the <code>mail</code> command to the administrator.	write	Write the message, using the <code>write</code> command, to the terminal on which the administrator is logged in. If the administrator is logged in on several terminals, one is arbitrarily chosen.	quiet	Do not send messages for the current condition. An administrator can use this option to temporarily stop receiving further messages				
mail	Send the alert message using the <code>mail</code> command to the administrator.										
write	Write the message, using the <code>write</code> command, to the terminal on which the administrator is logged in. If the administrator is logged in on several terminals, one is arbitrarily chosen.										
quiet	Do not send messages for the current condition. An administrator can use this option to temporarily stop receiving further messages										

lpforms(1M)

	about a known problem. Once the form <i>form-name</i> has been mounted and subsequently unmounted, messages will again be sent when the number of print requests reaches the threshold specified by the <i>-Q</i> option.
<i>showfault</i>	Attempt to execute a form alert handler on each system that has a print job for that form in the queue. The fault handler is <i>/etc/lp/alerts/form</i> . It is invoked with three parameters: <i>form_name</i> , <i>date</i> , <i>file_name</i> . <i>file_name</i> is the name of a file containing the form alert message.
<i>none</i>	Do not send messages until the <i>-A</i> option is given again with a different <i>alert-type</i> (other than <i>quiet</i>).
<i>shell-command</i>	Run the <i>shell-command</i> each time the alert needs to be sent. The shell command should expect the message in standard input. If there are blank spaces embedded in the command, enclose the command in quotes. Note that the <i>mail</i> and <i>write</i> values for this option are equivalent to the values <i>mail login-ID</i> and <i>write login-ID</i> respectively, where <i>login-ID</i> is the current name for the administrator. This will be the login name of the person submitting this command unless he or she has used the <i>su</i> command to change to another login-ID. If the <i>su</i> command has been used to change the user ID, then the <i>user-name</i> for the new ID is used.
<i>list</i>	Display the type of the alert for the form on standard output. No change is made to the alert.

The message sent appears as follows:

```
The form form-name needs to be mounted
on the printer(s):printer (integer1 requests).
integer2 print requests await this form.
Use the ribbon-color ribbon.
Use the print-wheel print wheel, if appropriate.
```

The printers listed are those that the administrator has specified as candidates for this form. The number *integer1* listed next to each printer is the number of requests eligible for the printer. The number *integer2* shown after the list of printers is the total number of requests awaiting the form. It will be less than the sum of the other numbers if some requests can be handled by more than one printer. The *ribbon-color* and *print-wheel* are those specified in the form description. The last line in the message is always sent, even if none of the printers listed use print wheels, because the administrator may choose to mount the form on a printer that does use a print wheel.

Where any color ribbon or any print wheel can be used, the statements above will read:

```
Use any ribbon. Use any print-wheel.
```

	<p>If <i>form-name</i> is any, the <i>alert-type</i> defined in this command applies to any form for which an alert has not yet been defined. If <i>form-name</i> is <code>all</code>, the <i>alert-type</i> defined in this command applies to all forms.</p> <p>If the <code>-W minutes</code> option is not given, the default procedure is that only one message will be sent per need to mount the form. Not specifying the <code>-W</code> option is equivalent to specifying <code>-W once</code> or <code>-W 0</code>. If <i>minutes</i> is a number greater than 0, an alert will be sent at intervals specified by <i>minutes</i>.</p> <p>If the <code>-Q requests</code> option is also given, the alert will be sent when a certain number (specified by the argument <i>requests</i>) of print requests that need the form are waiting. If the <code>-Q</code> option is not given, or the value of <i>requests</i> is 1 or any (which are both the default), a message is sent as soon as anyone submits a print request for the form when it is not mounted.</p>				
Listing the Current Alert	<p>The <code>-f</code> option, followed by the <code>-A</code> option and the argument <i>list</i> is used to list the <i>alert-type</i> that has been defined for the specified form <i>form-name</i>. No change is made to the alert. If <i>form-name</i> is recognized by the LP print service, one of the following lines is sent to the standard output, depending on the type of alert for the form.</p> <ul style="list-style-type: none"> – When <i>requests</i> requests are queued: alert with <i>shell-command</i> every <i>minutes</i> minutes – When <i>requests</i> requests are queued: write to <i>user-name</i> every <i>minutes</i> minutes – When <i>requests</i> requests are queued: mail to <i>user-name</i> every <i>minutes</i> minutes – No alert <p>The phrase every <i>minutes</i> minutes is replaced with <code>once</code> if <i>minutes</i> (<code>-W minutes</code>) is 0.</p>				
Terminating an Active Alert	<p>The <code>-A quiet</code> option is used to stop messages for the current condition. An administrator can use this option to temporarily stop receiving further messages about a known problem. Once the form has been mounted and then unmounted, messages will again be sent when the number of print requests reaches the threshold <i>requests</i>.</p>				
Removing an Alert Definition	<p>No messages will be sent after the <code>-A none</code> option is used until the <code>-A</code> option is given again with a different <i>alert-type</i>. This can be used to permanently stop further messages from being sent as any existing alert definition for the form will be removed.</p>				
Large File Behavior	<p>See <code>largefile(5)</code> for the description of the behavior of <code>lpforms</code> when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).</p>				
EXIT STATUS	<p>The following exit values are returned:</p> <table> <tr> <td>0</td><td>Successful completion.</td></tr> <tr> <td>non-zero</td><td>An error occurred.</td></tr> </table>	0	Successful completion.	non-zero	An error occurred.
0	Successful completion.				
non-zero	An error occurred.				

lpforms(1M)

FILES /etc/lp/alerts/form Fault handler for lpform.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWpsu

SUMMARY OF TRUSTED SQUARES CHANGES Use of the lpforms command requires the administer printing authorization.
lp(1), lpstat(1), lpadmin(1M)
Trusted Solaris 4/01 Reference Manual *Trusted Solaris Administrator's Procedures*
SunOS 5.8 Reference Manual terminfo(4), attributes(5), largefile(5)
System Administration Guide, Volume 1

NAME	lpmove – Move print requests								
SYNOPSIS	lpmove <i>request-ID destination</i> lpmove <i>destination1 destination2</i>								
DESCRIPTION	<p>The lpmove command moves print requests queued by lp(1) or lpr(1) between destinations. Only use lpmove to move jobs on the local system.</p> <p>lpmove requires the administer printing authorization.</p> <p>The first form of lpmove moves specific print requests (<i>request-ID</i>) to a specific (<i>destination</i>).</p> <p>The second form of the lpmove command moves all print requests from one destination (<i>destination1</i>) to another (<i>destination2</i>). This form of lpmove also rejects new print requests for <i>destination1</i>.</p> <p>When moving requests, lpmove does not check the acceptance status of the destination to which the print requests are being moved (see accept(1M)). lpmove does not move requests that have options (for example, content type or requiring a special form) that cannot be handled by the new destination.</p>								
OPERANDS	<p>The following operands are supported.</p> <table> <tr> <td><i>destination</i></td><td>The name of the printer or class of printers (see lpadmin(1M)) to which lpmove moves a <i>specified</i> print request. Specify <i>destination</i> using atomic, POSIX-style (<i>server:destination</i>), or Federated Naming Service (FNS) () names. See printers.conf(4) for information regarding the naming conventions for atomic and FNS names.</td></tr> <tr> <td><i>destination1</i></td><td>The name of the destination from which lpmove moves <i>all</i> print requests. Specify <i>destination</i> using atomic, POSIX-style (<i>server:destination</i>), or Federated Naming Service (FNS) (.../service/printer/...) names. See printers.conf(4) for information regarding the naming conventions for atomic and FNS names, and standards(5) for information regarding POSIX.</td></tr> <tr> <td><i>destination2</i></td><td>The name of the destination to which lpmove moves all print requests. Specify <i>destination</i> using atomic, POSIX-style (<i>server:destination</i>), or Federated Naming Service (FNS) (.../service/printer/...) names. See printers.conf(4) for information regarding the naming conventions for atomic and FNS names.</td></tr> <tr> <td><i>request-ID</i></td><td>The specific print request to be moved. Specify <i>request-ID</i> as the identifier associated with a print request as reported by lpstat. See lpstat(1).</td></tr> </table>	<i>destination</i>	The name of the printer or class of printers (see lpadmin(1M)) to which lpmove moves a <i>specified</i> print request. Specify <i>destination</i> using atomic, POSIX-style (<i>server:destination</i>), or Federated Naming Service (FNS) () names. See printers.conf(4) for information regarding the naming conventions for atomic and FNS names.	<i>destination1</i>	The name of the destination from which lpmove moves <i>all</i> print requests. Specify <i>destination</i> using atomic, POSIX-style (<i>server:destination</i>), or Federated Naming Service (FNS) (.../service/printer/...) names. See printers.conf(4) for information regarding the naming conventions for atomic and FNS names, and standards(5) for information regarding POSIX.	<i>destination2</i>	The name of the destination to which lpmove moves all print requests. Specify <i>destination</i> using atomic, POSIX-style (<i>server:destination</i>), or Federated Naming Service (FNS) (.../service/printer/...) names. See printers.conf(4) for information regarding the naming conventions for atomic and FNS names.	<i>request-ID</i>	The specific print request to be moved. Specify <i>request-ID</i> as the identifier associated with a print request as reported by lpstat . See lpstat(1) .
<i>destination</i>	The name of the printer or class of printers (see lpadmin(1M)) to which lpmove moves a <i>specified</i> print request. Specify <i>destination</i> using atomic, POSIX-style (<i>server:destination</i>), or Federated Naming Service (FNS) () names. See printers.conf(4) for information regarding the naming conventions for atomic and FNS names.								
<i>destination1</i>	The name of the destination from which lpmove moves <i>all</i> print requests. Specify <i>destination</i> using atomic, POSIX-style (<i>server:destination</i>), or Federated Naming Service (FNS) (.../service/printer/...) names. See printers.conf(4) for information regarding the naming conventions for atomic and FNS names, and standards(5) for information regarding POSIX.								
<i>destination2</i>	The name of the destination to which lpmove moves all print requests. Specify <i>destination</i> using atomic, POSIX-style (<i>server:destination</i>), or Federated Naming Service (FNS) (.../service/printer/...) names. See printers.conf(4) for information regarding the naming conventions for atomic and FNS names.								
<i>request-ID</i>	The specific print request to be moved. Specify <i>request-ID</i> as the identifier associated with a print request as reported by lpstat . See lpstat(1) .								
EXIT STATUS	The following exit values are returned:								

lpmove(1M)

0 Successful completion.

non-zero An error occurred.

FILES /var/spool/print/* LP print queue.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWpcu

SUMMARY OF TRUSTED SOLARIS CHANGES
Trusted Solaris 4/01 Reference Manual
SunOS 5.8 Reference Manual

lpmove requires the administer printing authorization.

lp(1), lpstat(1), lpr(1), accept(1M), lpadmin(1M), lpsched(1M)

Trusted Solaris Administrator's Procedures

printers.conf(4), attributes(5), standards(5)

System Administration Guide, Volume 1

NAME	lpsched – Start the LP print service									
SYNOPSIS	lpsched [-f num_filters] [-n num_notifiers] [-p fd_limit] [-r reserved_fds]									
DESCRIPTION	<p>The lpsched command starts or restarts the LP print service. lpsched must inherit these privileges: file_chown, file_dac_read, file_dac_write, file_dac_search, file_downgrade_sl, file_mac_read, file_mac_search, file_mac_write, file_owner, file_setdac, file_setid, file_upgrade_sl, net_downgrade_sl, net_mac_read, net_setpriv, net_setid, proc_setclr, proc_setsl, proc_setid, proc_audit_tcb, proc_owner, proc_mac_write, and sys_trans_label.</p> <p>The lpshut(1M) command stops the LP print service. Printers that are restarted using lpsched reprint (in their entirety) print requests that were stopped by lpshut.</p>									
OPTIONS	<p>The following options are supported:</p> <table><tr><td>-f num_filters</td><td>Specifies the number of concurrent slow filters that may be run on a print server. A default value of 1 is used if none is specified. Depending on server configuration, a value of 1 may cause printers to remain idle while there are jobs queued to them.</td></tr><tr><td>-n num_notifiers</td><td>Specifies the number of concurrent notification processes that can run on a print server. A default value of 1 is used when none is specified.</td></tr><tr><td>-p fd_limit</td><td>Specifies the file descriptor resource limit for the lpsched process. A default value of 4096 is used if none is specified. On extremely large and active print servers, it may be necessary to increase this value.</td></tr><tr><td>-r reserved_fds</td><td>Specifies the number of file descriptors that the scheduler reserves for internal communications under heavy load. A default value of 2 is used when none is specified. It should not be necessary to modify this value unless instructed to do so when troubleshooting problems under high load.</td></tr></table>		-f num_filters	Specifies the number of concurrent slow filters that may be run on a print server. A default value of 1 is used if none is specified. Depending on server configuration, a value of 1 may cause printers to remain idle while there are jobs queued to them.	-n num_notifiers	Specifies the number of concurrent notification processes that can run on a print server. A default value of 1 is used when none is specified.	-p fd_limit	Specifies the file descriptor resource limit for the lpsched process. A default value of 4096 is used if none is specified. On extremely large and active print servers, it may be necessary to increase this value.	-r reserved_fds	Specifies the number of file descriptors that the scheduler reserves for internal communications under heavy load. A default value of 2 is used when none is specified. It should not be necessary to modify this value unless instructed to do so when troubleshooting problems under high load.
-f num_filters	Specifies the number of concurrent slow filters that may be run on a print server. A default value of 1 is used if none is specified. Depending on server configuration, a value of 1 may cause printers to remain idle while there are jobs queued to them.									
-n num_notifiers	Specifies the number of concurrent notification processes that can run on a print server. A default value of 1 is used when none is specified.									
-p fd_limit	Specifies the file descriptor resource limit for the lpsched process. A default value of 4096 is used if none is specified. On extremely large and active print servers, it may be necessary to increase this value.									
-r reserved_fds	Specifies the number of file descriptors that the scheduler reserves for internal communications under heavy load. A default value of 2 is used when none is specified. It should not be necessary to modify this value unless instructed to do so when troubleshooting problems under high load.									
EXIT STATUS	<p>The following exit values are returned:</p> <table><tr><td>0</td><td>Successful completion.</td></tr><tr><td>non-zero</td><td>An error occurred.</td></tr></table>		0	Successful completion.	non-zero	An error occurred.				
0	Successful completion.									
non-zero	An error occurred.									
FILES	/var/spool/lp/*	LP print queue.								
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:									

lpsched(1M)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWpsu

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**
**Trusted Solaris 8
4/01 Reference
Manual**

**SunOS 5.8
Reference Manual**

lpsched must be started from the Trusted Path. It must be started as lp or root at the label admin_high and must inherit appropriate privileges.

enable(1), lp(1), lpstat(1), lpmove(1M), lpshut(1M), lpadmin(1M)

Trusted Solaris Administrator's Procedures

attributes(5)

System Administration Guide, Volume 1

NAME	lpshut – Stop the LP print service				
SYNOPSIS	lpshut				
DESCRIPTION	<p>The lpshut command stops the LP print service.</p> <p>lpshut requires the administer printing authorization.</p> <p>Printers that are printing when lpshut is invoked stop printing. Start or restart printers using lpsched(1M).</p>				
EXIT STATUS	<p>The following exit values are returned:</p> <p>0 Successful completion.</p> <p>non-zero An error occurred.</p>				
FILES	/var/spool/lp/* LP print queue.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWpsu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWpsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWpsu				
SUMMARY OF TRUSTED SOLARIS 4/01 CHANGES	<p>lpshut requires the administer printing authorization.</p> <p>lp(1), lpstat(1), lpadmin(1M), lpmove(1M), lpsched(1M)</p> <p><i>Trusted Solaris Administrator's Procedures</i></p>				
SunOS 5.8 Reference Manual	<p>attributes(5)</p> <p><i>System Administration Guide, Volume 1</i></p>				

lpsystem(1M)

NAME lpsystem – Register remote systems with the print service

DESCRIPTION The lpsystem command is obsolete. The print system no longer uses the information generated by lpsystem. See lpadmin(1M), lpusers(1M) or printers.conf(4) for equivalent functionality.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWpcu
Stability Level	Obsolete*

* This command could be removed at any time.

SUMMARY OF TRUSTED SOLARIS CHANGES Use of the lpsystem command requires the administer printing authorization

Trusted Solaris 4/01 Reference Manual lpadmin(1M), lpusers(1M)

Unix 4.3 Reference Manual printers.conf(4), attributes(5)

NAME	lpusers – Set printing queue priorities										
SYNOPSIS	<p>lpusers -d <i>priority-level</i></p> <p>lpusers -q <i>priority-level</i> -u <i>login-ID-list</i></p> <p>lpusers -u <i>login-ID-list</i></p> <p>lpusers -q <i>priority-level</i></p> <p>lpusers -l</p>										
DESCRIPTION	<p>The lpusers command sets limits to the queue priority level that can be assigned to jobs submitted by users of the LP print service.</p> <p>The first form of the command (with -d) sets the system-wide priority default to <i>priority-level</i>, where <i>priority-level</i> is a value of 0 to 39, with 0 being the highest priority. If a user does not specify a priority level with a print request (see lp(1)), the default priority level is used. Initially, the default priority level is 20.</p> <p>The second form of the command (with -q and -u) sets the default highest <i>priority-level</i> (0-39) that the users in <i>login-ID-list</i> can request when submitting a print request. The <i>login-ID-list</i> argument may include any or all of the following constructs:</p> <table> <tr> <td><i>login-ID</i></td><td>A user on any system</td></tr> <tr> <td><i>system_name</i>!<i>login-ID</i></td><td>A user on the system <i>system_name</i></td></tr> <tr> <td><i>system_name</i>!all</td><td>All users on system <i>system_name</i></td></tr> <tr> <td>all!<i>login-ID</i></td><td>A user on all systems</td></tr> <tr> <td>all</td><td>All users on all systems</td></tr> </table> <p>Users that have been given a limit cannot submit a print request with a higher priority level than the one assigned, nor can they change a request that has already been submitted to have a higher priority. Any print requests submitted with priority levels higher than allowed will be given the highest priority allowed.</p> <p>The third form of the command (with -u) removes any explicit priority level for the specified users.</p> <p>The fourth form of the command (with -q) sets the default highest priority level for all users not explicitly covered by the use of the second form of this command.</p> <p>The last form of the command (with -l) lists the default priority level and the priority limits assigned to users.</p>	<i>login-ID</i>	A user on any system	<i>system_name</i> ! <i>login-ID</i>	A user on the system <i>system_name</i>	<i>system_name</i> !all	All users on system <i>system_name</i>	all! <i>login-ID</i>	A user on all systems	all	All users on all systems
<i>login-ID</i>	A user on any system										
<i>system_name</i> ! <i>login-ID</i>	A user on the system <i>system_name</i>										
<i>system_name</i> !all	All users on system <i>system_name</i>										
all! <i>login-ID</i>	A user on all systems										
all	All users on all systems										
OPTIONS	<p>The following options are supported:</p> <p>-d <i>priority-level</i> Set the system-wide priority default to <i>priority-level</i>.</p>										

lpusers(1M)

- l
List the default priority level and the priority limits assigned to users.
- q *priority-level*
Set the default highest priority level for all users not explicitly covered.
- q *priority-level* -u *login-ID-list*
Set the default highest *priority-level* that the users in *login-ID-list* can request when submitting a print request.
- u *login-ID-list*
Remove any explicit priority level for the specified users.

EXIT STATUS The following exit values are returned:

- 0 Successful completion.
- non-zero An error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWpsu

SUMMARY OF TRUSTED SOLARIS CHANGES Use of the lpusers command requires the administer printing authorization.

Trusted Solaris 4/01 Reference Manual
lp(1)
attributes(5)

NAME	mkdevdb, mkdevalloc, mkdevmaps – make device_allocate and device_maps entries						
SYNOPSIS	<pre> /usr/sbin/mkdevdb [-a -m] /usr/sbin/mkdevalloc /usr/sbin/mkdevmaps </pre>						
DESCRIPTION	<p>The mkdevdb command writes by default to the /etc/security/device_allocate and /etc/security/device_maps files a set of device_allocate(4) and device_maps(4) entries describing the system's frame buffer, audio and removable media devices.</p> <p>The mkdevdb command is used by the <code>init.d(4)</code> scripts and utilities such as <code>devfsadm(1M)</code> to create or update the /etc/security/device_allocate and /etc/security/device_maps files. If either option is specified, mkdevdb writes to standard out only the specified file. If neither option is specified, mkdevdb writes both files.</p> <p>The mkdevmaps command writes to standard out a set of device_maps(4) entries describing the system's frame buffer, audio and removable media devices. Its use is deprecated.</p> <p>The mkdevalloc command writes to standard out a set of device_allocate(4) entries describing the system's frame buffer, audio and removable media devices. Its use is deprecated.</p> <p>Entries are generated based on the device special files found in /dev:</p> <pre> audio /dev/audio, /dev/audioctl, /dev/sound/... tape /dev/rst*, /dev/nrst*, /dev/rmt/... floppy /dev/diskette, /dev/fd*, /dev/rdiskette, /dev/rfd* removable disk /dev/dsk/c0t?d0s?, /dev/rdsk/c0t?d0s? frame buffer /dev/fb </pre>						
ATTRIBUTES	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> <tr> <td>Interface Stability</td><td> mkdevalloc — Obsolete mkdevmaps — Obsolete </td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu	Interface Stability	mkdevalloc — Obsolete mkdevmaps — Obsolete
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWcsu						
Interface Stability	mkdevalloc — Obsolete mkdevmaps — Obsolete						
OPTIONS	<p>-a Write device_allocate on standard out. For the different categories of devices, the mkdevdb command checks the device special files found in /dev.</p> <p> All entries set the <i>device-minimum</i> and <i>device-maximum</i> fields to the hex representations of ADMIN_LOW and ADMIN_HIGH, respectively. The</p>						

device-authorization field is set to `solaris.device.allocate`, except for the `framebuffer` entry, where it is set to `*`. The *device-name*, *device-type* and *device-clean* fields are set to the following values:

	device-name	device-type	device-clean
audio	audio	audio	audio_clean_wrapper
tape	mag_tape_0,...	st	st_clean
floppy	floppy_0,...	fd	disk_clean
cd	cdrom_0,...	sr	disk_clean
removable disk	rdisk_0,...	sr	disk_clean
frame buffer	framebuffer	fb	/bin/true

-m Write `device_maps` on standard out. For the different categories of devices, the `mkdevmaps` command checks the device special files found in `/dev`.

FILES The following files are used by the `mkdevdb` command:

`/etc/security/device_allocate`

Administrative file defining parameters for device allocation

`/etc/security/device_maps`

Administrative file defining the mapping of device special files to allocatable device names

`/dev/*`

Device special files

SUMMARY OF TRUSTED SOLARIS CHANGES

The `fb` device is added. The names of some devices are changed:

Device	Solaris Name	Trusted Solaris Name
tape	st0,st1,...	mag_tape_0,mag_tape_1,...
floppy	fd0,fd1,...	floppy_0,floppy_1,...
removable disk	sr0,sr1,...	cdrom_0,cdrom_1,...

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

`allocate(1)`, `deallocate(1)`, `device_allocate(4)`, `device_maps(4)`
`attributes(5)`

NAME	mkdevdb, mkdevalloc, mkdevmaps – make device_allocate and device_maps entries						
SYNOPSIS	<pre> /usr/sbin/mkdevdb [-a -m] /usr/sbin/mkdevalloc /usr/sbin/mkdevmaps </pre>						
DESCRIPTION	<p>The mkdevdb command writes by default to the /etc/security/device_allocate and /etc/security/device_maps files a set of device_allocate(4) and device_maps(4) entries describing the system's frame buffer, audio and removable media devices.</p> <p>The mkdevdb command is used by the <code>init.d(4)</code> scripts and utilities such as <code>devfsadm(1M)</code> to create or update the /etc/security/device_allocate and /etc/security/device_maps files. If either option is specified, mkdevdb writes to standard out only the specified file. If neither option is specified, mkdevdb writes both files.</p> <p>The mkdevmaps command writes to standard out a set of device_maps(4) entries describing the system's frame buffer, audio and removable media devices. Its use is deprecated.</p> <p>The mkdevalloc command writes to standard out a set of device_allocate(4) entries describing the system's frame buffer, audio and removable media devices. Its use is deprecated.</p> <p>Entries are generated based on the device special files found in /dev:</p> <pre> audio /dev/audio, /dev/audioctl, /dev/sound/... tape /dev/rst*, /dev/nrst*, /dev/rmt/... floppy /dev/diskette, /dev/fd*, /dev/rdiskette, /dev/rfd* removable disk /dev/dsk/c0t?d0s?, /dev/rdsk/c0t?d0s? frame buffer /dev/fb </pre>						
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> <tr> <td>Interface Stability</td><td> mkdevalloc — Obsolete mkdevmaps — Obsolete </td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu	Interface Stability	mkdevalloc — Obsolete mkdevmaps — Obsolete
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWcsu						
Interface Stability	mkdevalloc — Obsolete mkdevmaps — Obsolete						
OPTIONS	<p>-a Write device_allocate on standard out. For the different categories of devices, the mkdevdb command checks the device special files found in /dev.</p> <p> All entries set the <i>device-minimum</i> and <i>device-maximum</i> fields to the hex representations of ADMIN_LOW and ADMIN_HIGH, respectively. The</p>						

device-authorization field is set to `solaris.device.allocate`, except for the `framebuffer` entry, where it is set to `*`. The *device-name*, *device-type* and *device-clean* fields are set to the following values:

	device-name	device-type	device-clean
audio	audio	audio	audio_clean_wrapper
tape	mag_tape_0,...	st	st_clean
floppy	floppy_0,...	fd	disk_clean
cd	cdrom_0,...	sr	disk_clean
removable disk	rdisk_0,...	sr	disk_clean
frame buffer	framebuffer	fb	/bin/true

-m Write `device_maps` on standard out. For the different categories of devices, the `mkdevmaps` command checks the device special files found in `/dev`.

FILES The following files are used by the `mkdevdb` command:

`/etc/security/device_allocate`

Administrative file defining parameters for device allocation

`/etc/security/device_maps`

Administrative file defining the mapping of device special files to allocatable device names

`/dev/*`

Device special files

SUMMARY OF TRUSTED SOLARIS CHANGES

The `fb` device is added. The names of some devices are changed:

Device	Solaris Name	Trusted Solaris Name
tape	st0,st1,...	mag_tape_0,mag_tape_1,...
floppy	fd0,fd1,...	floppy_0,floppy_1,...
removable disk	sr0,sr1,...	cdrom_0,cdrom_1,...

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

`allocate(1)`, `deallocate(1)`, `device_allocate(4)`, `device_maps(4)`
`attributes(5)`

NAME	mkdevdb, mkdevalloc, mkdevmaps – make device_allocate and device_maps entries						
SYNOPSIS	<pre> /usr/sbin/mkdevdb [-a -m] /usr/sbin/mkdevalloc /usr/sbin/mkdevmaps </pre>						
DESCRIPTION	<p>The mkdevdb command writes by default to the /etc/security/device_allocate and /etc/security/device_maps files a set of device_allocate(4) and device_maps(4) entries describing the system's frame buffer, audio and removable media devices.</p> <p>The mkdevdb command is used by the <code>init.d(4)</code> scripts and utilities such as <code>devfsadm(1M)</code> to create or update the /etc/security/device_allocate and /etc/security/device_maps files. If either option is specified, mkdevdb writes to standard out only the specified file. If neither option is specified, mkdevdb writes both files.</p> <p>The mkdevmaps command writes to standard out a set of device_maps(4) entries describing the system's frame buffer, audio and removable media devices. Its use is deprecated.</p> <p>The mkdevalloc command writes to standard out a set of device_allocate(4) entries describing the system's frame buffer, audio and removable media devices. Its use is deprecated.</p> <p>Entries are generated based on the device special files found in /dev:</p> <pre> audio /dev/audio, /dev/audioctl, /dev/sound/... tape /dev/rst*, /dev/nrst*, /dev/rmt/... floppy /dev/diskette, /dev/fd*, /dev/rdiskette, /dev/rfd* removable disk /dev/dsk/c0t?d0s?, /dev/rdsk/c0t?d0s? frame buffer /dev/fb </pre> <p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> <tr> <td>Interface Stability</td><td>mkdevalloc — Obsolete mkdevmaps — Obsolete</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu	Interface Stability	mkdevalloc — Obsolete mkdevmaps — Obsolete
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWcsu						
Interface Stability	mkdevalloc — Obsolete mkdevmaps — Obsolete						
OPTIONS	<p>-a Write device_allocate on standard out. For the different categories of devices, the mkdevdb command checks the device special files found in /dev.</p> <p>All entries set the <i>device-minimum</i> and <i>device-maximum</i> fields to the hex representations of ADMIN_LOW and ADMIN_HIGH, respectively. The</p>						

mkdevmaps(1M)

device-authorization field is set to `solaris.device.allocate`, except for the `framebuffer` entry, where it is set to `*`. The *device-name*, *device-type* and *device-clean* fields are set to the following values:

	device-name	device-type	device-clean
audio	audio	audio	audio_clean_wrapper
tape	mag_tape_0,...	st	st_clean
floppy	floppy_0,...	fd	disk_clean
cd	cdrom_0,...	sr	disk_clean
removable disk	rdisk_0,...	sr	disk_clean
frame buffer	framebuffer	fb	/bin/true

-m Write `device_maps` on standard out. For the different categories of devices, the `mkdevmaps` command checks the device special files found in `/dev`.

FILES

The following files are used by the `mkdevdb` command:

`/etc/security/device_allocate`

Administrative file defining parameters for device allocation

`/etc/security/device_maps`

Administrative file defining the mapping of device special files to allocatable device names

`/dev/*`

Device special files

SUMMARY OF TRUSTED SOLARIS CHANGES

The `fb` device is added. The names of some devices are changed:

Device	Solaris Name	Trusted Solaris Name
tape	st0,st1,...	mag_tape_0,mag_tape_1,...
floppy	fd0,fd1,...	floppy_0,floppy_1,...
removable disk	sr0,sr1,...	cdrom_0,cdrom_1,...

Trusted Solaris 8 4/01 Reference Manual

`allocate(1)`, `deallocate(1)`, `device_allocate(4)`, `device_maps(4)`
`attributes(5)`

NAME	modload – Load a kernel module				
SYNOPSIS	modload [-p] [-e <i>exec_file</i>] <i>filename</i>				
DESCRIPTION	<p>modload loads the loadable module <i>filename</i> into the running system. <i>filename</i> is an object file produced by <code>ld -r</code>. If <i>filename</i> is an absolute pathname then the file specified by that absolute path is loaded. If <i>filename</i> does not begin with a '/' then the path to load <i>filename</i> is relative to the current directory unless the <code>-p</code> option is specified. The kernel's modpath variable can be set using the <code>/etc/system</code> file. The default value of the kernel's modpath variable is set to the path where the operating system was loaded. Typically this is <code>/kernel /usr/kernel</code>. Hence if you type:</p> <pre>example# modload drv/foo</pre> <p>The kernel will look for <code>./drv/foo</code>.</p> <p>If you type:</p> <pre>example# modload -p drv/foo</pre> <p>The kernel will look for <code>/kernel/drv/foo</code> and then <code>/usr/kernel/drv/foo</code>.</p>				
OPTIONS	<p><code>-p</code> Use the kernel's internal modpath variable as the search path for the module.</p> <p><code>-e <i>exec_file</i></code> Specify the name of a shell script or executable image file that is executed after the module is successfully loaded. The first argument passed is the module ID (in decimal). The other argument is module specific. The module specific information is: the block and character major numbers for drivers, the system call number for system calls, or, for other module types, the index into the appropriate kernel table. See <code>modinfo(1M)</code></p>				
SUMMARY OF TRUSTED SOLARIS ATTRIBUTES CHANGES	<p>To succeed, this command needs the <code>sys_devices</code> privilege.</p> <p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
Trusted Solaris 8 4/01 Reference Manual	<p><code>ld(1)</code>, <code>add_drv(1M)</code>, <code>modunload(1M)</code></p> <p><code>kernel(1M)</code>, <code>modinfo(1M)</code>, <code>system(4)</code>, <code>attributes(5)</code>, <code>modldrv(9S)</code>, <code>modlinkage(9S)</code>, <code>modlstrmod(9S)</code>, <code>module_info(9S)</code></p> <p><i>Writing Device Drivers Solaris Transition Guide</i></p>				
NOTES	Use <code>add_drv(1M)</code> to add device drivers, not <code>modload</code> . See <i>Writing Device Drivers</i> for procedures on adding device drivers.				

modunload(1M)

NAME modunload – Unload a module

SYNOPSIS **modunload** -i *module_id* [-e *exec_file*]

DESCRIPTION modunload unloads a loadable module from the running system. The *module_id* is the ID of the module as shown by modinfo(1M). If ID is 0, all modules that were autoloaded which are unloadable, are unloaded. Modules loaded by modload(1M) are not affected.

OPTIONS

-i <i>module_id</i>	Specify the module to be unloaded.
-e <i>exec_file</i>	Specify the name of a shell script or executable image file to be executed before the module is unloaded. The first argument passed is the module ID (in decimal). There are two additional arguments that are module specific. For loadable drivers, the second and third arguments are the block major and character major numbers respectively. For loadable system calls, the second argument is the system call number. For loadable exec classes, the second argument is the index into the <i>execsw</i> table. For loadable filesystems, the second argument is the index into the <i>vfssw</i> table. For loadable streams modules, the second argument is the index into the <i>fmodsw</i> table. For loadable scheduling classes, the second argument is the index into the class array. Minus one is passed for an argument that does not apply.

SUMMARY OF TRUSTED ATTRIBUTES CHANGES To succeed, this command needs the *sys_devices* privilege.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8
4/01 Reference
Manual
Solaris 8
Modul
Reference Manual

modload(1M)

modinfo(1M), attributes(5)

NAME	mount, umount – mount or unmount file systems and remote resources
SYNOPSIS	<pre> mount [-p -v] mount [-F <i>FSType</i>] [<i>generic_options</i>] [-o <i>specific_options</i>] [-O] <i>special</i> <i>mount_point</i> mount [-F <i>FSType</i>] [<i>generic_options</i>] [-o <i>specific_options</i>] [-O] [-S <i>attribute_list</i>] <i>special mount_point</i> mount -a [-F <i>FSType</i>] [-V] [<i>current_options</i>] [-o <i>specific_options</i>] [-S <i>attribute_list</i>] [<i>mount_point...</i>] umount [-f] [-V] [-o <i>specific_options</i>] <i>special</i> <i>mount_point</i> umount -a [-f] [-V] [-o <i>specific_options</i>] [<i>mount_point...</i>] umount -a [-V] [-o <i>specific_options</i>] [<i>mount_point...</i>] </pre>
DESCRIPTION	<p>mount attaches a file system to the file system hierarchy at the <i>mount_point</i>, which is the pathname of a directory. If <i>mount_point</i> has any contents prior to the mount operation, these are hidden until the file system is unmounted.</p> <p>umount unmounts a currently mounted file system, which may be specified either as a <i>mount_point</i> or as <i>special</i>, the device on which the file system resides.</p> <p>The table of currently mounted file systems can be found by examining the mounted file system information file. This is provided by a file system that is usually mounted on <i>/etc/mnttab</i>. The mounted file system information is described in <i>mnttab(4)</i>. Mounting a file system adds an entry to the mount table; a umount command removes an entry from the table.</p> <p>When invoked with both the <i>special</i> and <i>mount_point</i> arguments and the -F option, mount validates all arguments except for <i>special</i> and invokes the appropriate <i>FSType</i>-specific mount module. If invoked with no arguments, mount lists all the mounted file systems recorded in the mount table, <i>/etc/mnttab</i>. If invoked with a partial argument list (with only one of <i>special</i> or <i>mount_point</i>, or with both <i>special</i> or <i>mount_point</i> specified but not <i>FSType</i>), mount will search <i>/etc/vfstab</i> for an entry that will supply the missing arguments. If no entry is found, and the special argument starts with "/", the default local file system type specified in <i>/etc/default/fs</i> will be used. Otherwise the default remote file system type will be used. The default remote file system type is determined by the first entry in the <i>/etc/dfs/fstypes</i> file. After filling in missing arguments, mount will invoke the <i>FSType</i>-specific mount module.</p> <p>The -o and -S options can be used to assign any or all of the following mount-time security attributes to the named file system when appropriate: a sensitivity label, forced privilege(s), allowed privilege(s), a filesystem label range, or an MLD prefix. If -o or -S options are not used, mount also searches <i>/etc/security/tsol/vfstab_adjunct</i> for any security attributes that may be specified there for the file system being mounted.</p>

mount(1M)

Mount-time security attributes should be specified for file systems whose objects do not support the Trusted Solaris extended security attributes, such as sensitivity labels. When a required attribute is not specified at mount-time, a default value is applied. The defaults are described in the `OPTIONS` section, where the keywords are defined for the `-S` option.

File system types `UFS`, `TMPFS`, and `NFS` (from a Trusted Solaris server) have a full set of Trusted Solaris extended security attributes already defined. (See the `getfsattr(1M)` man page for how to get attributes on mounted file systems). Because the attributes can be changed on these file systems *after* they are mounted, they are called *variable* file systems. For example, the sensitivity label on a file in a variable file system can be changed by an authorized user. The security attributes on a variable file system can be overridden at mount time, but individual objects in the file system retain any attributes that were originally set on the objects.

File systems that do not support the Trusted Solaris extended security attributes are called *fixed* because any attributes assigned to them (either at mount time or by default) cannot be changed. For example, the sensitivity label specified at mount time for a fixed-attribute file system cannot be changed on any of the objects in that file system. An object that is moved or copied from the fixed file system to a variable file system can be changed after the move.

Mount-time security attributes override existing security attributes on a file system. However, mount-time attributes never override security attributes on the files and directories within the file system.

Without privilege, `mount` can be used to list mounted file systems and resources. To be able to mount and unmount, the `mount` command must have the `sys_mount` privilege. The `umount` command must have the `sys_mount` privilege. Because mounting a `UFS` file system enables/disables logging, it requires the `sys_fs_config` privilege. Mandatory and discretionary read access is required both to the mount point and to the device being mounted; otherwise, MAC or DAC override privileges are required as described in `Intro(2)`. To succeed in all cases with no error side effects, the `mount` command needs: `file_mac_read`, `file_dac_read`, `file_mac_write`, `file_dac_write`, `file_mac_search`, `file_dac_search`, `net_privaddr`, `proc_setsl`, `sys_fs_config`, `sys_mount`, and `sys_trans_label`. To succeed in all cases, `umount` needs: `file_mac_read`, `file_dac_read`, `file_mac_write`, `file_dac_write`, `file_mac_search`, and `file_dac_search`.

When mounting a `UFS` file system, `mount` should assert the `sys_fs_config` privilege. Otherwise, the `mount` succeeds, but logging is not enabled/disabled, `errno` is set to `EPERM`, and the user sees an error message.

OPTIONS

`-F FSType`

Used to specify the `FSType` on which to operate. The `FSType` must be specified or must be determinable from `/etc/vfstab`, or by consulting `/etc/default/fs` or `/etc/dfs/fstypes`.

-a [*mount_points*. . .]

Perform mount or umount operations in parallel, when possible.

If mount points are not specified, mount will mount all file systems whose */etc/vfstab* "mount at boot" field is "yes". If mount points are specified, then */etc/vfstab* "mount at boot" field will be ignored.

If mount points are specified, umount will only unmount those mount points. If none is specified, then umount will attempt to unmount all filesystems in */etc/mnttab*, with the exception of certain system required file systems: */*, */usr*, */var*, */proc*, */dev/fd*, and */tmp*.

-f

Forcibly unmount a file system.

Without this option, umount does not allow a file system to be unmounted if a file on the file system is busy. Using this option can cause data loss for open files; programs which access files after the file system has been unmounted will get an error (EIO).

-p

Print the list of mounted file systems in the */etc/vfstab* format. Must be the only option specified.

-v

Print the list of mounted file systems in verbose format. Must be the only option specified.

-V

Echo the complete command line, but do not execute the command. umount generates a command line by using the options and arguments provided by the user and adding to them information derived from */etc/mnttab*. This option should be used to verify and validate the command line.

generic_options

Options that are commonly supported by most *FSType*-specific command modules. The following options are available:

-m

Mount the file system without making an entry in */etc/mnttab*.

-g

Globally mount the file system. On a clustered system, this globally mounts the file system on all nodes of the cluster. On a non-clustered system this has no effect.

-o

Specify *FSType*-specific options in a comma separated (without spaces) list of suboptions and keyword-attribute pairs for interpretation by the *FSType*-specific module of the command. (See *mount_ufs(1M)*)

mount(1M)

-O	Overlay mount. Allow the file system to be mounted over an existing mount point, making the underlying file system inaccessible. If a mount is attempted on a pre-existing mount point without setting this flag, the mount will fail, producing the error “device busy”.												
-r	Mount the file system read-only.												
-S <i>attribute_list</i>	<p>Specify in <i>attribute_list</i> a quoted semicolon-separated list of security attributes to associate with the filesystem mount. Each attribute is specified with a value assigned to a keyword in semicolon-separated fields. All keywords are optional and follow the format:</p> <p>keyword=value where <i>keyword</i> is one of the following:</p> <table><tr><td>slabel</td><td>Sets the sensitivity label for all objects in the file system. Specify the sensitivity label in hexadecimal or text format.</td></tr><tr><td>forced</td><td>Specify one or more forced privileges for all executable files in the file system. Specify symbolic privilege name(s) in a comma-separated list (such as: forced=file_audit, file_chown;) or use all to indicate all privileges. Using none or omitting the keyword results in no forced privileges being applied. See priv_desc(4). Any forced privileges must be a subset of the allowed privileges.</td></tr><tr><td>allowed</td><td>Specify one or more allowed privilege(s) for all executable files in the file system. Specify symbolic privilege names in a comma-separated list (such as: allowed=file_audit, file_chown;) or use all to indicate all privileges. Using none or omitting the keyword results in no allowed privileges being applied. See priv_desc(4) for names of privileges. Any allowed privilege(s) must be a superset of the forced privileges.</td></tr><tr><td>low_range</td><td>Specify the lower bound of the file system label range as a sensitivity label in text format.</td></tr><tr><td>hi_range</td><td>Specify the upper bound of the file system label range as a sensitivity label in text format.</td></tr><tr><td>mld_prefix</td><td>Set a prefix to be used in the adorned names of multilevel directories. (See multilevel directories in the DEFINITIONS in Intro(2) for more about the MLD prefix.) Specify the value in text format (such as: .MLD. or .hidden.). On unlabeled (fixed attribute) file systems, the prefix generally has no useful effect—with the exception that an mld_prefix should be supplied if a variable filesystem is being mounted on the unlabeled filesystem and the root of the variable filesystem</td></tr></table>	slabel	Sets the sensitivity label for all objects in the file system. Specify the sensitivity label in hexadecimal or text format.	forced	Specify one or more forced privileges for all executable files in the file system. Specify symbolic privilege name(s) in a comma-separated list (such as: forced=file_audit, file_chown;) or use all to indicate all privileges. Using none or omitting the keyword results in no forced privileges being applied. See priv_desc(4). Any forced privileges must be a subset of the allowed privileges.	allowed	Specify one or more allowed privilege(s) for all executable files in the file system. Specify symbolic privilege names in a comma-separated list (such as: allowed=file_audit, file_chown;) or use all to indicate all privileges. Using none or omitting the keyword results in no allowed privileges being applied. See priv_desc(4) for names of privileges. Any allowed privilege(s) must be a superset of the forced privileges.	low_range	Specify the lower bound of the file system label range as a sensitivity label in text format.	hi_range	Specify the upper bound of the file system label range as a sensitivity label in text format.	mld_prefix	Set a prefix to be used in the adorned names of multilevel directories. (See multilevel directories in the DEFINITIONS in Intro(2) for more about the MLD prefix.) Specify the value in text format (such as: .MLD. or .hidden.). On unlabeled (fixed attribute) file systems, the prefix generally has no useful effect—with the exception that an mld_prefix should be supplied if a variable filesystem is being mounted on the unlabeled filesystem and the root of the variable filesystem
slabel	Sets the sensitivity label for all objects in the file system. Specify the sensitivity label in hexadecimal or text format.												
forced	Specify one or more forced privileges for all executable files in the file system. Specify symbolic privilege name(s) in a comma-separated list (such as: forced=file_audit, file_chown;) or use all to indicate all privileges. Using none or omitting the keyword results in no forced privileges being applied. See priv_desc(4). Any forced privileges must be a subset of the allowed privileges.												
allowed	Specify one or more allowed privilege(s) for all executable files in the file system. Specify symbolic privilege names in a comma-separated list (such as: allowed=file_audit, file_chown;) or use all to indicate all privileges. Using none or omitting the keyword results in no allowed privileges being applied. See priv_desc(4) for names of privileges. Any allowed privilege(s) must be a superset of the forced privileges.												
low_range	Specify the lower bound of the file system label range as a sensitivity label in text format.												
hi_range	Specify the upper bound of the file system label range as a sensitivity label in text format.												
mld_prefix	Set a prefix to be used in the adorned names of multilevel directories. (See multilevel directories in the DEFINITIONS in Intro(2) for more about the MLD prefix.) Specify the value in text format (such as: .MLD. or .hidden.). On unlabeled (fixed attribute) file systems, the prefix generally has no useful effect—with the exception that an mld_prefix should be supplied if a variable filesystem is being mounted on the unlabeled filesystem and the root of the variable filesystem												

mount(1M)

is an MLD.

Any of the above keywords may be omitted.

Note – The semicolon separators between keyword/value pairs and any brackets used to specify sensitivity labels must be commented out so that the separators and brackets can be interpreted properly by the shell.

When a keyword appears without an attribute value or when a keyword is missing, a default value is assigned to that attribute. The default values for fixed attribute file systems are:

slabel	The default sensitivity label of a fixed file system being mounted from a local device (such as a hard disk, floppy, or CD-ROM) is the sensitivity label of the device. For an allocated device, the file system is assigned the sensitivity label at which the device was allocated.
forced	None
allowed	None
low_range	ADMIN_LOW
hi_range	ADMIN_HIGH
mld_prefix	None

For example, the assignment of `forced=`; results in the default of "none" being applied.

Note – Most of the keyword=value pairs used to specify security attributes with the `-S` option can be entered directly under the `-o` option—with one caveat. Since mount options are comma-separated, any security attribute specified with a keyword followed by multiple values separated by commas is not allowed after `-o`. See Example 2.

USAGE See `largefile(5)` for the description of the behavior of `mount` and `umount` when encountering files greater than or equal to 2 Gbyte (2³¹ bytes).

EXAMPLES **EXAMPLE 1** Assigning Security Attributes with `-o`

In this example, the `-o` is used to assign security attributes.

```
% mount -F tmpfs -o allowed=all,slabel=c swap /mnt
```

EXAMPLE 2 Assigning Security Attributes with `-S`

Trusted Solaris security attributes that are separated with commas cannot be passed to the `-o` option. Therefore, use the `-S` option.

```
% mount -F tmpfs -S "allowed=all;forced=proc_tranquil,proc_dumpcore" \\  
swap /mnt
```

mount(1M)

EXAMPLE 2 Assigning Security Attributes with -S (Continued)

These security attributes cannot be entered with the -o option since the comma separator in the privileges list would be interpreted as the start of a new option.

FILES	/etc/mnttab	Mount table
	/etc/default/fs	Default local file system type. Default values can be set for the following flags in /etc/default/fs. For example: LOCAL=ufs
		LOCAL: The default partition for a command if no <i>FSType</i> is specified.
	/etc/vfstab	List of default parameters for each file system.
	/etc/security/tsol/vfstab_adjunct	Mount-time attributes for file systems.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

Trusted Solaris security policy applies when mounting and unmounting file systems.

Mount-time security attributes may be specified by using mount with the -o or -S option on the command line or by specifying the attributes in the vfstab_adjunct file. Mount-time security attributes override existing security attributes on a file system. However, they never override security attributes on the files and directories within the file system. When access-control decisions are made, security attributes on a file or directory take precedence over security attributes specified either at the filesystem level or at mount time.

Except when merely listing mounted file systems and resources, mount must run with the sys_mount privilege. umount also must run with the sys_mount privilege. To succeed in all cases, mount needs: file_mac_read, file_dac_read, file_mac_write, file_dac_write, file_mac_search, file_dac_search, net_privaddr, proc_setsl, sys_mount, and sys_trans_label.

When mounting a UFS file system, mount should assert the sys_fs_config privilege. Otherwise, the mount succeeds, but logging is not enabled/disabled, errno is set to EPERM, and the user sees an error message.

Trusted Solaris 8 4/01 Reference Manual

getfsattr(1M), getmldadorn(1), mount_hsf(1M), mount_nfs(1M), mount_pcfs(1M), mount_tmpfs(1M), mount_ufs(1M), mountall(1M), setfsattr(1M), mnttab(4), priv_desc(4), vfstab(4), vfstab_adjunct(4)

*Trusted Solaris Administrator's Procedures***SunOS 5.8
Reference Manual**mount_cachefs(1M), default_fs(4), attributes(5), largefile(5), lofs(7FS),
pcfs(7FS)**NOTES**

If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on the directory to which the symbolic link refers, rather than on top of the symbolic link itself.

mountall(1M)

NAME	mountall, umountall – Mount, unmount multiple file systems
SYNOPSIS	<p>mountall [-F <i>FSType</i>] [-l -r] [<i>file_system_table</i>]</p> <p>umountall [-k] [-s] [-F <i>FSType</i>] [-l -r]</p> <p>umountall [-k] [-s] [-h <i>host</i>]</p>
DESCRIPTION	<p>mountall is used to mount file systems specified in a file system table. The file system table must be in <code>vfstab(4)</code> format. If no <i>file_system_table</i> is specified, <code>/etc/vfstab</code> will be used. If '-' is specified as <i>file_system_table</i>, mountall will read the file system table from the standard input. mountall only mounts those file systems with the mount at boot field set to yes in the <i>file_system_table</i>.</p> <p>Each file system which has an <code>fsckdev</code> entry specified in the file system table will be checked using <code>fsck(1M)</code> in order to determine if it may be safely mounted. If the file system does not appear mountable, it is fixed using <code>fsck</code> before the mount is attempted. File systems with a '-' entry in the <code>fsckdev</code> field will be mounted without first being checked.</p> <p>umountall causes all mounted file systems except <code>root</code>, <code>/usr</code>, <code>/var</code>, <code>/var/adm</code>, <code>/var/run</code>, <code>/proc</code>, and <code>/dev/fd</code> to be unmounted. If the <i>FSType</i> is specified, mountall and umountall limit their actions to the <i>FSType</i> specified. There is no guarantee that umountall will unmount <i>busy</i> filesystems, even if the -k option is specified.</p> <p>mountall and umountall must run with the <code>sys_mount</code> privilege.</p> <p>Mandatory and discretionary read access are required to both the mount point and the device being mounted; to override MAC and DAC restrictions requires privilege as described in <code>Intro(2)</code>. To succeed in all cases, the mountall and umountall commands need the privileges: <code>file_mac_read</code>, <code>file_dac_read</code>, <code>file_mac_write</code>, <code>file_dac_write</code>, <code>file_mac_search</code>, <code>file_dac_search</code>, <code>net_privaddr</code>, <code>proc_setsid</code>, <code>sys_mount</code>, and <code>sys_trans_label</code>.</p> <p>When mounting a UFS file system, mount should assert the <code>sys_fs_config</code> privilege. Otherwise, the mount succeeds, but logging is not enabled/disabled, <code>errno</code> is set to <code>EPERM</code>, and the user sees an error message.</p>
OPTIONS	<p>-F Specify the <i>FSType</i> of the file system to be mounted or unmounted.</p> <p>-h <i>host</i> Unmount all file systems listed in <code>/etc/mnttab</code> that are remote-mounted from host.</p> <p>-k Use the <code>fuser -k mount-point</code> command. See the <code>fuser(1M)</code> for details. The -k option sends the <code>SIGKILL</code> signal to each process using the file. As this option spawns kills for each process, the kill messages may not show up immediately. There is no guarantee that umountall will unmount <i>busy</i> filesystems, even if the -k option is specified.</p> <p>-l Limit the action to local file systems.</p>

	-r	Limit the action to remote file system types.				
	-s	Do not perform the umount operation in parallel.				
FILES	/etc/mnttab	mounted file system table				
	/etc/vfstab	table of file system defaults				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:					
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWcsu</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE					
Availability	SUNWcsu					
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>Trusted Solaris security policy applies when mounting and unmounting file systems.</p> <p>mountall and umountall must run with the sys_mount privilege.</p> <p>Mandatory and discretionary read access are required to both the mount point and the device being mounted; to override MAC and DAC restrictions requires privilege as described in Intro(2). To succeed in all cases, the mountall and umountall commands need the privileges: file_mac_read, file_dac_read, file_mac_write, file_dac_write, file_mac_search, file_dac_search, net_privaddr, proc_setsl, sys_mount, and sys_trans_label.</p> <p>When mounting a UFS file system, mount should assert the sys_fs_config privilege. Otherwise, the mount succeeds, but logging is not enabled/disabled, errno is set to EPERM, and the user sees an error message.</p> <p>Mount-time security attributes may be specified in the vfstab_adjunct file.</p>					
Trusted Solaris 8 4/01 Reference Manual	mount(1M), mnttab(4), vfstab(4), vfstab_adjunct(4)					
Solaris 8 4/01 Reference Manual	fsck(1M), fuser(1M), attributes(5)					
DIAGNOSTICS	<p>No messages are printed if the file systems are mountable and clean.</p> <p>Error and warning messages come from fsck(1M) and mount(1M).</p>					

mountd(1M)

NAME	mountd – Server for NFS mount requests and NFS access checks				
SYNOPSIS	<code>/usr/lib/nfs/mountd [-v] [-r]</code>				
DESCRIPTION	<p>mountd is an RPC server that answers requests for NFS access information and file system mount requests. It reads the file <code>/etc/dfs/sharetab</code> to determine which file systems are available for mounting by which remote machines. See <code>sharetab(4)</code>. <code>nfsd</code> running on the local server will contact mountd the first time an NFS client tries to access the file system to determine whether the client should get read-write, read-only, or no access. This access can be dependent on the security mode used in the remotd procedure call from the client. See <code>share_nfs(1M)</code>.</p> <p>The command also provides information as to what file systems are mounted by which clients. This information can be printed using the <code>showmount(1M)</code> command.</p> <p>The mountd daemon is automatically invoked in run level 3.</p> <p>The <code>sys_nfs</code>, <code>sys_devices</code>, <code>sys_net_config</code>, <code>sys_audit</code>, <code>net_mac_read</code>, <code>net_privaddr</code>, <code>file_mac_read</code>, <code>file_mac_write</code>, <code>file_mac_search</code>, <code>file_dac_search</code>, <code>proc_setsl</code>, and <code>proc_setclr</code> privileges are required to run this daemon. This daemon must be started from the Trusted Path; otherwise, the daemon sets its sensitivity label to <code>ADMIN_LOW</code> and clearance to <code>ADMIN_HIGH</code>.</p>				
OPTIONS	<p><code>-v</code> Run the command in verbose mode. Each time mountd determines what access a client should get, it will log the result to the console, as well as how it got that result.</p> <p><code>-r</code> Reject mount requests from clients. Clients that have file systems mounted will not be affected.</p>				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>The <code>sys_nfs</code>, <code>sys_devices</code>, <code>sys_net_config</code>, <code>sys_audit</code>, <code>net_mac_read</code>, <code>net_privaddr</code>, <code>file_mac_read</code>, <code>file_mac_search</code>, <code>file_dac_search</code>, <code>proc_setsl</code>, and <code>proc_setclr</code> privileges are required to run this daemon. This daemon must be started from the Trusted Path. If not started at a sensitivity label of <code>ADMIN_LOW</code> and clearance of <code>ADMIN_HIGH</code>, mountd sets its sensitivity label and clearance to these values.</p> <p>For the mount request to succeed, this daemon requires the client to have the <code>sys_mount</code> privilege. Unless the <code>-n</code> option is specified, the client request must bind to a privileged port.</p>				
FILES	<code>/etc/dfs/sharetab</code> shared file system table				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				

nfstd(1M), share_nfs(1M), showmount(1M), sharetab(4)

attributes(5)

If `nfstd` is running, `mountd` must also be running in order to be assured that the NFS server can respond to requests, otherwise, the NFS service can hang.

Some routines that compare hostnames use case-sensitive string comparisons; some do not. If an incoming request fails, verify that the case of the hostname in the file to be parsed matches the case of the hostname called for, and attempt the request again.

mount_hsf(1M)

NAME	mount_hsf – Mount hsf file systems
SYNOPSIS	mount -F hsf [generic_options] [-o FSType-specific_options] [-O] [-S attribute_list] special mount_point mount -F hsf [generic_options] [-o FSType-specific_options] [-O] [-S attribute_list] special mount_point
DESCRIPTION	<p>mount attaches a High Sierra file system (hsf) to the file system hierarchy at the <i>mount_point</i>, which is the pathname of a directory. If <i>mount_point</i> has any contents prior to the mount operation, these are hidden until the file system is unmounted.</p> <p>If mount is invoked with <i>special</i> or <i>mount_point</i> as the only arguments, mount will search /etc/vfstab to fill in the missing arguments, including the <i>FSType-specific_options</i>; see mount(1M) for more details.</p> <p>If the file system being mounted contains Rock Ridge extensions, by default they will be used, enabling support of features not normally available under High Sierra file systems, such as symbolic links and special files.</p> <p>Security attributes can be specified at mount time, either with the -o or -S option on the mount command line or in the vfstab_adjunct(4) file. See the DESCRIPTION in the mount(1M) man page for more about specifying security attributes.</p> <p>To succeed, the mount command must have the sys_mount privilege. Mandatory and discretionary read access are required to both the mount point and the device being mount; to override MAC and DAC restrictions requires privilege as described in Intro(2). To succeed in all cases, mount -F hsf needs the file_mac_read and file_mac_write privileges.</p>
OPTIONS	<p><i>generic_options</i> See mount(1M) for the list of supported options.</p> <p>-o Specify hsf file system specific options. Most attributes for the -S option may also be specified for the -o option. See the -S option.</p> <p>If invalid options are specified, a warning message is printed and the invalid options are ignored. The following options are available:</p> <p>global noglobal If global is specified and supported on the file system, and the system in question is part of a cluster, the file system will be globally visible on all nodes of the cluster. If noglobal is specified, the mount will not be globally visible. The default behavior is noglobal .</p> <p>ro Mount the file system read-only. This option is required.</p> <p>nrr no Rock Ridge: if Rock Ridge extensions are present in the file system, ignore them; interpret it as a regular High Sierra file system.</p>

notraildot

File names on High Sierra file systems consist of a proper name and an extension separated by a '.' (dot) character. By default, the separating dot is always considered part of the file's name for all file access operations, even if there is no extension present. Specifying **notraildot** makes it optional to specify the trailing dot to access a file whose name lacks an extension.

Exceptions: This option is effective only on file systems for which Rock Ridge extensions are not active, either because they are not present on the CD-ROM, or they are explicitly ignored via the **nrr** option. If Rock Ridge extensions are active, **hfs** quietly ignores this option.

nomaplcse

File names on High Sierra cdroms with no Rock Ridge extensions present should be uppercase characters only. By default, **hfs** maps file names read from a non-Rock Ridge disk to all lowercase characters. **nomaplcse** turns off this mapping. The exceptions for **notraildot** discussed above apply to **nomaplcse**.

nosuid

By default the file system is mounted with **setuid** execution allowed. Specifying **nosuid** causes the file system to be mounted with **setuid** execution disallowed.

devices | nodevices

Allow (disallow) access to character and block devices. The default is **devices**.

Note: In the Trusted Solaris environment, device special files are typically located only in the **/dev** and **/devices** directories in the root file system. All other file systems should be mounted with the **nodevices** option to prevent recognition of devices that may reside in any other directories. The recognition of devices is also affected by the use of the **devices** or **nodevices** options to the **share(1M)** command, either on the command line or in the **dfstab(4)** file.

priv | nopriv

Forced privileges on executables are allowed or disallowed. The default is **priv**. The recognition of forced privileges is also affected by the use of the **priv** or **nopriv** option to the **share(1M)** command, either on the command line or in the **dfstab(4)** file.

-O

Overlay mount. Allow the file system to be mounted over an existing mount point, making the underlying file system inaccessible. If a mount is attempted on a pre-existing mount point without setting this flag, the mount will fail, producing the error device busy.

-S attribute_list

See the **DESCRIPTION** and the attribute list on the **mount(1M)** man page.

SUMMARY OF TRUSTED SOLARIS CHANGES

The **nodevices** and **nopriv** options have been added. Trusted Solaris security policy applies when mounting and unmounting file systems.

mount_hsf(1M)

Except when merely listing mounted file systems and resources, `mount` must run with the `sys_mount` privilege.

Mandatory and discretionary read access is required to both the mount point and the device being mounted; to override MAC and DAC restrictions requires privilege as described in `Intro(2)`. To succeed in all cases, `mount -F hsf` needs the `file_mac_read` and `file_mac_write` privileges.

Mount-time security attributes can be specified for file systems whose objects do not have any attributes (such as user and group IDs) and for file systems that do not have the Trusted Solaris extended security attributes (such as sensitivity labels). Specifying attributes fails for file systems and file system objects (files or directories) that already have a specified attribute. Trusted Solaris security policy applies when mounting. See the `mount(1M)` and `vfstab_adjunct(4)` man pages for more details.

FILES

`/etc/mnttab`

Table of mounted file systems.

`/etc/vfstab`

List of default parameters for each file system.

`/etc/security/tsol/vfstab_adjunct`

Mount-time attributes for file systems.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8
4/01 Reference
Manual
NOTES

`mount(1M)`, `mountall(1M)`, `mount(2)`, `mnttab(4)`, `vfstab(4)`, `vfstab_adjunct(4)`
`attributes(5)`

If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on the directory to which the symbolic link refers, rather than on top of the symbolic link itself.

NAME	mount_nfs – mount remote NFS resources
SYNOPSIS	<p>mount [-F nfs] [<i>generic_options</i>] [-o <i>specific_options</i>] [-O] [-S <i>attribute_list</i>] <i>resource</i></p> <p>mount [-F nfs] [<i>generic_options</i>] [-o <i>specific_options</i>] [-O] [-S <i>attribute_list</i>] <i>mount_point</i></p> <p>mount [-F nfs] [<i>generic_options</i>] [-o <i>specific_options</i>] [-O] [-S <i>attribute_list</i>] <i>resource mount_point</i></p>
DESCRIPTION	<p>The mount utility attaches a named <i>resource</i> to the file system hierarchy at the pathname location <i>mount_point</i>, which must already exist. If <i>mount_point</i> has any contents prior to the mount operation, the contents remain hidden until the <i>resource</i> is once again unmounted.</p> <p>If the resource is listed in the <i>/etc/vfstab</i> file, the command line can specify either <i>resource</i> or <i>mount_point</i>, and mount will consult <i>/etc/vfstab</i> for more information. If the -F option is omitted, mount takes the file system type from <i>/etc/vfstab</i>.</p> <p>If the resource is not listed in the <i>/etc/vfstab</i> file, then the command line must specify both the <i>resource</i> and the <i>mount_point</i>.</p> <p>A named <i>resource</i> can have one of the following formats:</p> <p><i>host:pathname</i></p> <p>Where <i>host</i> is the name of the NFS server host, and <i>pathname</i> is the path name of the directory on the server being mounted. The path name is interpreted according to the server's path name parsing rules and is not necessarily slash-separated, though on most servers, this will be the case.</p> <p><i>nfs:host[:port]/pathname</i></p> <p>This is an NFS URL and follows the standard convention for NFS URLs as described in <i>Internet RFC 2225 — NFS URL Scheme</i>. See the discussion of URL's and the public option under NFS FILE SYSTEMS below for a more detailed discussion.</p> <p>A comma-separated list of <i>host:pathname</i> and/or <i>nfs:host[:port]/pathname</i> resources</p> <p>See the discussion of Replicated file systems and failover under NFS FILE SYSTEMS below for a more detailed discussion.</p> <p>mount maintains a table of mounted file systems in <i>/etc/mnttab</i>, described in <i>mnttab(4)</i>. See <i>mount(1M)</i> for more details.</p> <p>Security attributes can be specified at mount time, with the -o or -S option on the mount command line or in the <i>vfstab_adjunct(4)</i> file. See the DESCRIPTION in the mount man page for more about specifying security attributes.</p> <p>Trusted Solaris security policy applies when mounting and unmounting file systems.</p>

mount_nfs(1M)

mount must run with the `sys_mount` and `net_privaddr` privileges. To succeed in all cases, mount also needs: `file_mac_read`, `file_dac_read`, `file_mac_write`, `file_dac_write`, `file_mac_search`, `file_dac_search`, `proc_setsid`, and `sys_trans_label`.

OPTIONS

See `mount(1M)` for the list of supported *generic_options*.

`-o specific_options`

Set file system specific options according to a comma-separated list with no intervening spaces. Most attributes for the `-S` option may also be specified for the `-o` option. See the `-S` option.

`acdirmax=n`

Hold cached attributes for no more than *n* seconds after directory update. The default value is 60.

`acdirmin=n`

Hold cached attributes for at least *n* seconds after directory update. The default value is 30.

`acregmax=n`

Hold cached attributes for no more than *n* seconds after file modification. The default value is 60.

`acregmin=n`

Hold cached attributes for at least *n* seconds after file modification. The default value is 3.

`actimeo=n`

Set *min* and *max* times for regular files and directories to *n* seconds.

`bg | fg`

If the first attempt fails, retry in the background, or, in the foreground. The default is `fg`.

`devices | nodevices`

Allow (disallow) opens on character and block devices. The default is `devices`.

Note: In the Trusted Solaris environment, device special files are typically located only in the `/dev` and `/devices` directories in the root file system. All other file systems should be mounted with the `nodevices` option to prevent recognition of devices that may reside in any other directories.

`grpuid`

By default, the GID associated with a newly created file will obey the System V semantics; that is, the GID is set to the effective GID of the calling process. This behavior may be overridden on a per-directory basis by setting the set-GID bit of the parent directory; in this case, the GID of a newly created file is set to the GID of the parent directory (see `open(2)` and `mkdir(2)`). Files created on file systems that are mounted with the `grpuid` option will obey BSD semantics independent of whether the set-GID bit of the parent directory is set; that is, the GID is unconditionally inherited from that of the parent directory.

mount_nfs(1M)

`hard | soft`

Return an error if the server does not respond, or continue the retry request until the server responds. The default value is `hard`.

`intr | nointr`

Allow (do not allow) keyboard interrupts to kill a process that is hung while waiting for a response on a hard-mounted file system. The default is `intr`, which makes it possible for clients to interrupt applications that may be waiting for a remote mount.

`kerberos`

This option has been deprecated in favor of the `sec=krb4` option.

`noac`

Suppress data and attribute caching.

`port=n`

The server IP port number. The default is `NFS_PORT`. If the `port` option is specified, and if the resource includes one or more NFS URLs, and if any of the URLs include a `port` number, then the `port` number in the option and in the URL must be the same.

`posix`

Request POSIX.1 semantics for the file system. Requires a mount Version 2 `mountd(1M)` on the server. See `standards(5)` for information regarding POSIX.

`priv | nopriv`

Forced privileges on executables are allowed or disallowed. The default is `priv`.

`proto=<netid>`

`<netid>` is a value of `network_id` field from entry in the `/etc/netconfig` file. By default, the transport protocol used for the NFS mount will be first available connection oriented transport supported on both the client and the server. If no connection oriented transport is found, then the first available connectionless transport is used. This default behavior can be overridden with the `proto=<netid>` option.

`public`

The `public` option forces the use of the public file handle when connecting to the NFS server. The resource specified may or may not have an NFS URL. See the discussion of URL's and the `public` option under `NFS FILE SYSTEMS` below for a more detailed discussion.

`quota | noquota`

Enable or prevent `quota(1M)` to check whether the user is over quota on this file system; if the file system has quotas enabled on the server, quotas will still be checked for operations on this file system. This option is not supported in the Trusted Solaris environment.

mount_nfs(1M)

remount

Remounts a read-only file system as read-write (using the `rw` option). This option cannot be used with other `-o` options, and this option works only on currently mounted read-only file systems.

retrans=*n*

Set the number of NFS retransmissions to *n*. The default value is 5. For connection-oriented transports, this option has no effect because it is assumed that the transport will perform retransmissions on behalf of NFS.

retry=*n*

The number of times to retry the mount operation. The default is 10000.

The default for the automounter is 0, in other words, do not retry. You might find it useful to increase this value on heavily loaded servers, where automounter traffic is dropped, causing unnecessary “server not responding” errors.

ro | rw

resource is mounted read-only or read-write. The default is `rw`.

rsize=*n*

Set the read buffer size to *n* bytes. The default value is 32768 when using Version 3 of the NFS protocol. The default can be negotiated down if the server prefers a smaller transfer size. When using Version 2, the default value is 8192.

sec=*mode*

Set the security *mode* for NFS transactions. If `sec=` is not specified, then the default action is to use `AUTH_SYS` over NFS Version 2 mounts, or to negotiate a *mode* over NFS Version 3 mounts. NFS Version 3 mounts negotiate a security mode when the server returns an array of security modes. The client will pick the first mode in the array that is supported on the client. Only one mode can be specified with the `sec=` option. See `nfssec(5)` for the available *mode* options.

secure

This option has been deprecated in favor of the `sec=dh` option.

suid | nosuid

Allow or disallow `setuid` execution. The default is `suid`.

timeo=*n*

Set the NFS timeout to *n* tenths of a second. The default value is 11 tenths of a second for connectionless transports, and 600 tenths of a second for connection-oriented transports.

vers=<NFS version number>

By default, the version of NFS protocol used between the client and the server is the highest one available on both systems. If the NFS server does not support NFS Version 3 protocol, then the NFS mount will use NFS Version 2 protocol.

Note: File systems being mounted from Trusted Solaris 1.2 servers should be specified with `vers=2`. Because the Trusted Solaris operating environment does

not recognize security attributes, such as labels, on file systems mounted from NFS Version 2 servers, all such file systems should be mounted as unlabeled file systems and should have mount-time security attributes supplied for them either with the `-S` option or in the `vfstab_adjunct` file.

`wsize=n`

Set the write buffer size to *n* bytes. The default value is 32768 when using Version 3 of the NFS protocol. The default can be negotiated down if the server prefers a smaller transfer size. When using Version 2, the default value is 8192.

`-S attribute_list`

See the definition of the `-S` option in the `OPTIONS` section of the `mount(1M)` man page.

`-O` Overlay mount. Allow the file system to be mounted over an existing mount point, making the underlying file system inaccessible. If a mount is attempted on a pre-existing mount point without setting this flag, the mount will fail, producing the error “device busy”.

NFS FILE SYSTEMS

Background versus Foreground

File systems mounted with the `-bg` option indicate that mount is to retry in the background if the server's mount daemon (`mountd(1M)`) does not respond. `mount` retries the request up to the count specified in the `retry=n` option. (Note that the default value for `retry` differs between `mount` and `automount`. See the description of `retry`, above.) Once the file system is mounted, each NFS request made in the kernel waits `timeo=n` tenths of a second for a response. If no response arrives, the time-out is multiplied by 2 and the request is retransmitted. When the number of retransmissions has reached the number specified in the `retrans=n` option, a file system mounted with the `soft` option returns an error on the request; one mounted with the `hard` option prints a warning message and continues to retry the request.

Hard versus Soft

File systems that are mounted read-write or that contain executable files should always be mounted with the `hard` option. Applications using `soft` mounted file systems may incur unexpected I/O errors, file corruption, and unexpected program core dumps. The `soft` option is not recommended.

Authenticated Requests

The server may require authenticated NFS requests from the client. Either `sec=dh` or `sec=krb4` authentication may be required. See `nfssec(5)`.

URLs and the public option

If the `public` option is specified, or if the *resource* includes an NFS URL, `mount` will attempt to connect to the server using the public file handle lookup protocol. See *Internet RFC 2054 — WebNFS Client Specification*. If the server supports the public file handle, the attempt is successful; `mount` will not need to contact the server's `rpcbind(1M)`, and the `mountd(1M)` daemons to get the port number of the mount server and the initial file handle of *pathname*, respectively. If the NFS client and server are separated by a firewall that allows all outbound connections through

mount_nfs(1M)

specific ports, such as `NFS_PORT`, then this enables NFS operations through the firewall. The `public` option and the NFS URL can be specified independently or together. They interact as specified in the following matrix:

	resource style	
	<i>host:pathname</i>	NFS URL
public option	+ force public file handle and fail mount if not supported. + use Native paths	+ force public file handle and fail mount if not supported. + use Canonical paths
default	+ use MOUNT protocol	+ try public file handle with Canonical paths. Fall back to MOUNT protocol if not supported.

A *Native path* is a path name that is interpreted according to conventions used on the native operating system of the NFS server. A *Canonical path* is a path name that is interpreted according to the URL rules. See *Internet RFC 1738 — Uniform Resource Locators (URL)*. Also, see **EXAMPLES for uses of *Native* and *Canonical* paths.**

Replicated file systems and failover

resource can list multiple read-only file systems to be used to provide data. These file systems should contain equivalent directory structures and identical files. It is also recommended that they be created by a utility such as `rdist(1)`. The file systems may be specified either with a comma-separated list of *host:pathname* entries and/or NFS URL entries, or with a comma-separated list of hosts, if all file system names are the same. If multiple file systems are named and the first server in the list is down, failover will use the next alternate server to access files. If the read-only option is not chosen, replication will be disabled. File access will block on the original if NFS locks are active for that file.

File Attributes

To improve NFS read performance, files and file attributes are cached. File modification times get updated whenever a write occurs. However, file access times may be temporarily out-of-date until the cache gets refreshed.

The attribute cache retains file attributes on the client. Attributes for a file are assigned a time to be flushed. If the file is modified before the flush time, then the flush time is extended by the time since the last modification (under the assumption that files that changed recently are likely to change soon). There is a minimum and maximum flush time extension for regular files and for directories. Setting `actimeo=n` sets flush time to *n* seconds for both regular files and directories.

Setting `actimeo=0` disables attribute caching on the client. This means that every reference to attributes will be satisfied directly from the server though file data will

mount_nfs(1M)

still be cached. While this guarantees that the client always has the latest file attributes from the server, it has an adverse effect on performance through additional latency, network load, and server load.

Setting the `noac` option also disables attribute caching, but has the further effect of disabling client write caching. While this guarantees that data written by an application will be written directly to a server, where it can be viewed immediately by other clients, it has a significant adverse effect on client write performance. Data written into memory-mapped file pages (`mmap(2)`) will not be written directly to this server.

EXAMPLES

EXAMPLE 1 Mounting An NFS File System

To mount an NFS file system:

```
example# mount serv:/usr/src /usr/src
```

EXAMPLE 2 Mounting An NFS File System Read-Only With No Suid Privileges

To mount an NFS file system read-only with no suid privileges:

```
example# mount -r -o nosuid serv:/usr/src /usr/src
```

EXAMPLE 3 Mounting An NFS File System Over Version 2, With The UDP Transport

To mount an NFS file system over Version 2, with the UDP transport:

```
example# mount -o vers=2,proto=udp serv:/usr/src /usr/src
```

EXAMPLE 4 Mounting An NFS File System Using An NFS URL

To mount an NFS file system using an NFS URL (a canonical path):

```
example# mount nfs://serv/usr/man /usr/man
```

EXAMPLE 5 Mounting An NFS File System Forcing Use Of The Public File Handle

To mount an NFS file system and force the use of the public file handle and an NFS URL (a canonical path) that has a non 7-bit ASCII escape sequence:

```
example# mount -o public nfs://serv/usr/%A0abc /mnt/test
```

EXAMPLE 6 Mounting An NFS File System Using A Native Path

To mount an NFS file system using a native path (where the server uses colons (":") as the component separator) and the public file handle:

```
example# mount -o public serv:C:doc:new /usr/doc
```

mount_nfs(1M)

EXAMPLE 6 Mounting An NFS File System Using A Native Path (Continued)

EXAMPLE 7 Mounting an NFS file system using AUTH_KERB authentication.

To mount an NFS file system using AUTH_KERB authentication:

```
example# mount -o sec=krb4 serv:/usr/src /usr/src
```

EXAMPLE 8 Mounting a replicated set of NFS file systems with the same pathnames.

To mount a replicated set of NFS file systems with the same pathnames:

```
example# mount serv-a,serv-b,serv-c:/usr/man /usr/man
```

EXAMPLE 9 Mounting a replicated set of NFS file systems with different pathnames.

To mount a replicated set of NFS file systems with different pathnames:

```
example# mount serv-x:/usr/man,serv-y:/var/man,nfs://serv-z/man /usr/man
```

SUMMARY OF TRUSTED SOLARIS CHANGES

The `-o quota` option has been removed; and the `nodevices` and `nopriv` options have been added.

Mount-time security attributes can be specified for file systems whose objects do not have any attributes (such as user and group IDs) and for file systems that do not have the Trusted Solaris extended security attributes (such as sensitivity labels). Trusted Solaris security policy applies when mounting.

`mount` must run with the `sys_mount` and `net_privaddr` privileges. To succeed in all cases, `mount` also needs: `file_mac_read`, `file_dac_read`, `file_mac_write`, `file_dac_write`, `file_mac_search`, `file_dac_search`, `proc_setsl`, and `sys_trans_label`.

FILES

<code>/etc/mnttab</code>	table of mounted file systems
<code>/etc/dfs/fstypes</code>	default distributed file system type
<code>/etc/vfstab</code>	table of automatically mounted resources

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8 4/01 Reference Manual

`mount(1M)`, `mountall(1M)`, `mountd(1M)`, `mkdir(2)`, `mount(2)`, `open(2)`, `umount(2)`, `mnttab(4)`, `vfstab_adjunct(4)`

SunOS 5.8 Reference Manual

`rdist(1)`, `quota(1M)`, `mmap(2)`, `attributes(5)`, `nfssec(5)`, `standards(5)`, `lofs(7FS)`

NOTES | The sensitivity label mount-time attributes are only useful for mounts from NFS servers that are not labels-cognizant. The mount-time sensitivity label must always be equal to the assigned `slabel`, if one is specified, in the NFS server's combination `tnrhdb(4)/ tnrhtp(4)` entry. An unlabeled file system is always mounted at the sensitivity label specified for the unlabeled server in the trusted networking databases; if a different sensitivity label is specified at mount time, the mount fails.

An NFS server should not attempt to mount its own file systems. See `lofs(7FS)`.

If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on *the directory to which the symbolic link refers*, rather than being mounted on top of the symbolic link itself.

SunOS 4.X used the `biod` maintenance procedure to perform parallel read-ahead and write-behind on NFS clients. SunOS 5.X made `biod` obsolete with multi-threaded processing, which transparently performs parallel read-ahead and write-behind.

Since the root (/) file system is mounted read-only by the kernel during the boot process, only the `remount` option (and options that can be used in conjunction with `remount`) affect the root (/) entry in the `/etc/vfstab` file.

mount_pcfs(1M)

NAME	mount_pcfs – mount pcfs file systems
SYNOPSIS	mount -F pcfs [<i>generic_options</i>] [-o <i>FSType-specific_options</i>] [-S <i>attribute_list</i>] <i>special</i> <i>mount_point</i> mount -F pcfs [<i>generic_options</i>] [-o <i>FSType-specific_options</i>] [-S <i>attribute_list</i>] <i>special mount_point</i>
DESCRIPTION	<p>mount attaches an MS-DOS file system (pcfs) to the file system hierarchy at the <i>mount_point</i>, which is the pathname of a directory. If <i>mount_point</i> has any contents prior to the mount operation, these are hidden until the file system is unmounted.</p> <p>If mount is invoked with <i>special</i> or <i>mount_point</i> as the only arguments, mount will search /etc/vfstab to fill in the missing arguments, including the <i>FSType-specific_options</i>; see mount(1M) for more details.</p> <p>The <i>special</i> argument can be one of two special device file types:</p> <ul style="list-style-type: none">■ A floppy disk, such as /dev/diskette0 or /dev/diskette1.■ A DOS logical drive on a hard disk expressed as <i>device-name:logical-drive</i>, where <i>device-name</i> specifies the special block device-file for the whole disk and <i>logical-drive</i> is either a drive letter (c through z) or a drive number (1 through 24). Examples are /dev/dsk/c0t0d0p0:c and /dev/dsk/c0t0d0p0:1. <p>The <i>special</i> device file type must have a formatted MS-DOS file system with either a 12-bit, 16-bit, or 32-bit File Allocation Table.</p> <p>Security attributes can be specified at mount time, with the -o or -S option on the mount command line or in the vfstab_adjunct(4) file. See the DESCRIPTION in the mount man page for more about specifying security attributes.</p> <p>To succeed, the mount command must have the sys_mount privilege. Mandatory and discretionary read access is required to both the mount point and the device being mounted; to override MAC and DAC restrictions requires privilege as described in Intro(2). To succeed in all cases, mount -F pcfs needs the file_mac_read, file_mac_write, and sys_trans_label privileges.</p>
OPTIONS	<p><i>generic_options</i> See mount(1M) for the list of supported options.</p> <p>-o Specify pcfs file system specific options. Most attributes for the -S option may also be specified for the -o option. See the -S option.</p> <p>The following options are available:</p> <p>rw ro Mount the file system read/write or read-only. The default is rw.</p>

foldcase|nofoldcase

Force uppercase characters in filenames to lowercase when reading them from the filesystem. This is for compatibility with the previous behavior of pcfs. The default is nofoldcase.

-S attribute_list

See the definition of the -S option in the OPTIONS section of the mount(1M) man page.

SUMMARY OF TRUSTED SOLARIS CHANGES

Trusted Solaris security policy applies when mounting and unmounting file systems.

Except when merely listing mounted file systems and resources, mount must run with the sys_mount privilege.

Mandatory and discretionary read access is required to both the mount point and the device being mounted; to override MAC and DAC restrictions requires privilege as described in Intro(2). To succeed in all cases, mount -F pcfs need the file_mac_read, file_mac_write, and sys_trans_label privileges.

Mount-time security attributes can be specified for file systems whose objects do not have any attributes (such as user and group IDs) and for file systems that do not have the Trusted Solaris extended security attributes (such as sensitivity labels). Trusted Solaris security policy applies when mounting.

FILES

/etc/mnttab table of mounted file systems

/etc/vfstab list of default parameters for each file system

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWesu

Trusted Solaris 8 4/01 Reference Manual NOTES

mount(1M), mountall(1M), mount(2), mnttab(4), vfstab(4), vfstab_adjunct(4)

attributes(5), pcfs(7FS)

If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on the directory to which the symbolic link refers, rather than on top of the symbolic link itself.

mount_tmpfs(1M)

NAME	mount_tmpfs – Mount tmpfs file systems
SYNOPSIS	mount [-F tmpfs] [-o size= sz] [-O] [-S <i>attribute_list</i>] <i>special mount_point</i>
DESCRIPTION	<p>tmpfs is a memory-based file system which uses kernel resources relating to the VM system and page cache as a file system.</p> <p>mount attaches a tmpfs file system to the file system hierarchy at the pathname location <i>mount_point</i>, which must already exist. If <i>mount_point</i> has any contents prior to the mount operation, these remain hidden until the file system is once again unmounted. The attributes (mode, owner, and group) of the root of the tmpfs filesystem are inherited from the underlying <i>mount_point</i>, along with some security attributes (e.g., sensitivity label), provided that those attributes are determinable. If not, the root's attributes are set to their default values.</p> <p>The <i>special</i> argument is usually specified as <i>swap</i> but is in fact disregarded and assumed to be the virtual memory resources within the system.</p> <p>Security attributes can be specified at mount time, with the -o or -S option on the mount command line or in the <i>vfstab_adjunct(4)</i> file. See the DESCRIPTION in the mount man page for more about specifying security attributes.</p> <p>To succeed, the mount command must have the <i>sys_mount</i> privilege. Mandatory and discretionary read access is required to both the mount point and the device being mounted; to override MAC and DAC restrictions requires privilege as described in <i>Intro(2)</i>. To succeed in all cases, mount -F <i>hsfs</i> also needs: <i>file_mac_read</i>, <i>file_dac_read</i>, <i>file_mac_write</i>, <i>file_dac_write</i>, <i>file_mac_search</i>, <i>file_dac_search</i>, <i>net_privaddr</i>, <i>proc_setsl</i>, and <i>sys_trans_label</i>.</p>
OPTIONS	<p>-o Specify <i>ufs</i> file system specific options in a comma-separated list with no intervening spaces. Most attributes for the -S option may also be specified for the -o option. See the -S option.</p> <p>If invalid options are specified, a warning message is printed and the invalid options are ignored. The following options are available:</p> <p>size=sz The <i>sz</i> argument controls the size of this particular tmpfs file system. If the argument is has a 'k' suffix, the number will be interpreted as a number of kilobytes. An 'm' suffix will be interpreted as a number of megabytes. No suffix is interpreted as bytes. In all cases, the actual size of the file system is the number of bytes specified, rounded up to the physical pagesize of the system.</p>

SUMMARY OF
TRUSTED
SOLARIS
CHANGES

	mount_tmpfs(1M)
suid nosuid	Setuid execution allowed or disallowed. The default is suid. nosuid without an explicit devices implies nodevices.
devices nodevices	Allow (disallow) access to character and block devices. The default is devices. Note: In the Trusted Solaris environment, device special files are typically located only in the /dev and /devices directories in the root file system. All other file systems should be mounted with the nodevices option to prevent recognition of devices that may reside in any other directories. The recognition of devices is also affected by the use of the devices or nodevices options to the share(1M) command, either on the command line or in the dfstab(4) file.
priv nopriv	Forced privileges on executables are allowed or disallowed. The default is priv. The recognition of forced privileges is also affected by the use of the priv or nopriv option to the share(1M) command, either on the command line or in the dfstab(4) file.
-O	Overlay mount. Allow the file system to be mounted over an existing mount point, making the underlying file system inaccessible. If a mount is attempted on a pre-existing mount point without setting this flag, the mount will fail, producing the error "device busy".
-S <i>attribute_list</i>	See the definition of the -S option in the OPTIONS section of the mount(1M) man page.

The nodevices and nopriv options have been added. Trusted Solaris security policy applies when mounting and unmounting file systems.

mount_tmpfs(1M)

mount must run with the `sys_mount` privilege. To succeed in all cases, mount also needs: `file_mac_read`, `file_dac_read`, `file_mac_write`, `file_dac_write`, `file_mac_search`, `file_dac_search`, `net_privaddr`, `proc_setsl`, and `sys_trans_label`.

Mount-time security attributes can be specified for file systems whose objects do not have any attributes (such as user and group IDs) and for file systems that do not have the Trusted Solaris extended security attributes (such as sensitivity labels). Trusted Solaris security policy applies when mounting.

FILES `/etc/mnttab` table of mounted file systems

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

**Trusted Solaris 8
4/01 Reference
Manual** `mount(1M)`, `mkdir(2)`, `mount(2)`, `open(2)`, `umount(2)`, `mnttab(4)`,
`vfstab_adjunct(4)`

**SunOS 5.8
Reference Manual** `attributes(5)`, `tmpfs(7FS)`

NOTES If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on the directory to which the symbolic link refers, rather than on top of the symbolic link itself.

NAME	mount_ufs – Mount ufs file systems
SYNOPSIS	<pre> mount -F ufs [<i>generic_options</i>] [-o <i>specific_options</i>] [-O] [-S <i>attribute_list</i>] <i>special</i> <i>mount_point</i> mount -F ufs [<i>generic_options</i>] [-o <i>specific_options</i>] [-O] [-S <i>attribute_list</i>] <i>special</i> <i>mount_point</i> </pre>
DESCRIPTION	<p>The <code>mount</code> utility attaches a <code>ufs</code> file system to the file system hierarchy at the <i>mount_point</i>, which is the pathname of a directory. If <i>mount_point</i> has any contents prior to the <code>mount</code> operation, these are hidden until the file system is unmounted.</p> <p>If <code>mount</code> is invoked with <i>special</i> or <i>mount_point</i> as the only arguments, <code>mount</code> will search <code>/etc/vfstab</code> to fill in the missing arguments, including the <i>specific_options</i>. See <code>mount(1M)</code> for more details.</p> <p>If <i>special</i> and <i>mount_point</i> are specified without any <i>specific_options</i>, the default is <code>rw</code>.</p> <p>If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on the directory to which the symbolic link refers, rather than on top of the symbolic link itself.</p> <p>Security attributes can be specified at mount time, with the <code>-o</code> or <code>-S</code> option on the <code>mount</code> command line or in the <code>vfstab_adjunct(4)</code> file. See the DESCRIPTION in the <code>mount(1M)</code> man page for more about specifying security attributes.</p> <p>Except when merely listing mounted file systems and resources, <code>mount</code> must run with the <code>sys_mount</code> privilege.</p> <p>To succeed in all cases, <code>mount</code> needs: <code>file_mac_read</code>, <code>file_dac_read</code>, <code>file_mac_write</code>, <code>file_dac_write</code>, <code>file_mac_search</code>, <code>file_dac_search</code>, <code>net_privaddr</code>, <code>proc_setsl</code>, <code>sys_mount</code>, and <code>sys_trans_label</code>.</p> <p>When mounting a UFS file system, <code>mount</code> should assert the <code>sys_fs_config</code> privilege. Otherwise, the <code>mount</code> succeeds, but logging is not enabled/disabled, <code>errno</code> is set to <code>EPERM</code>, and the user sees an error message.</p>
OPTIONS	<p>See <code>mount(1M)</code> for the list of supported <i>generic_options</i>.</p> <p>The following options are supported:</p> <p><code>-o <i>specific_options</i></code> Specify <code>ufs</code> file system specific options in a comma-separated list with no intervening spaces. Most attributes for the <code>-S</code> option may also be specified for the <code>-o</code> option. See the <code>-S</code> option.</p> <p>If invalid options are specified, a warning message is printed and the invalid options are ignored. The following options are available:</p>

mount_ufs(1M)

noatime

By default, the file system is mounted with normal access time (`atime`) recording. If `noatime` is specified, the file system will ignore access time updates on files, except when they coincide with updates to the `ctime` or `mtime`. See `stat(2)`. This option reduces disk activity on file systems where access times are unimportant (for example, a Usenet news spool).

`noatime` turns off access time recording regardless of `dfratime` or `nodfratime`.

devices | nodevices

Allow (disallow) opens on character and block devices. The default is `devices`.

Note: In the Trusted Solaris environment, device special files are typically located only in the `/dev` and `/devices` directories in the root file system. All other file systems should be mounted with the `nodevices` option to prevent recognition of devices that may reside in any other directories.

dfratime | nodfratime

By default, writing access time updates to the disk may be deferred (`dfratime`) for the file system until the disk is accessed for a reason other than updating access times. `nodfratime` disables this behavior.

f

Fake an `/etc/mnttab` entry, but do not actually mount any file systems. Parameters are not verified.

forcedirectio | noforcedirectio

If `forcedirectio` is specified and supported by the file system, then for the duration of the mount forced direct I/O will be used. If the filesystem is mounted using `forcedirectio`, then data is transferred directly between user address space and the disk. If the filesystem is mounted using `noforcedirectio`, then data is buffered in kernel address space when data is transferred between user address space and the disk. `forcedirectio` is a performance option that benefits only from large sequential data transfers. The default behavior is `noforcedirectio`.

global | noglobal

If `global` is specified and supported on the file system, and the system in question is part of a cluster, the file system will be globally visible on all nodes of the cluster. If `noglobal` is specified, the mount will not be globally visible. The default behavior is `noglobal`.

intr | nointr

Allow (do not allow) keyboard interrupts to kill a process that is waiting for an operation on a locked file system. The default is `intr`.

largefiles | nolargefiles

If `nolargefiles` is specified and supported by the file system, then for the duration of the mount it is guaranteed that all regular files in the file system have a size that will fit in the smallest object of type `off_t` supported by the

mount_ufs(1M)

system performing the mount. The mount will fail if there are any files in the file system not meeting this criterion. If `largefiles` is specified, there is no such guarantee. The default behavior is `largefiles`.

If `nolargefiles` is specified, `mount` will fail for `ufs` if the file system to be mounted has contained a large file (a file whose size is greater than or equal to 2 Gbyte) since the last invocation of `fsck` on the file system. The large file need not be present in the file system at the time of the mount for the mount to fail; it could have been created previously and destroyed. Invoking `fsck` (see `fsck_ufs(1M)`) on the file system will reset the file system state if no large files are present. After invoking `fsck`, a successful mount of the file system with `nolargefiles` specified indicates the absence of large files in the file system; an unsuccessful mount attempt indicates the presence of at least one large file.

logging | nologging

If `logging` is specified, then logging is enabled for the duration of the mounted file system. Logging is the process of storing transactions (changes that make up a complete UFS operation) in a log before the transactions are applied to the file system. Once a transaction is stored, the transaction can be applied to the file system later. This prevents file systems from becoming inconsistent, therefore eliminating the need to run `fsck`. And, because `fsck` can be bypassed, logging reduces the time required to reboot a system if it crashes, or after an unclean halt. The default behavior is `nologging`.

The log is allocated from free blocks on the file system, and is sized approximately 1 Mbyte per 1 Gbyte of file system, up to a maximum of 64 Mbytes. Logging can be enabled on any UFS, including root (/). The log created by UFS logging is continually flushed as it fills up. The log is totally flushed when the file system is unmounted or as a result of the `lockfs -f` command.

m

Mount the file system without making an entry in `/etc/mnttab`.

onerror=*action*

This option specifies the action that UFS should take to recover from an internal inconsistency on a file system. Specify *action* as `panic`, `lock`, or `umount`. These values cause a forced system shutdown, a file system lock to be applied to the file system, or the file system to be forcibly unmounted, respectively. The default is `panic`.

priv | nopriv

Forced privileges on executables are allowed or disallowed. The default is `priv`.

quota

This option is not supported in Trusted Solaris; any attempt to set this option is ignored. Quotas are turned on for the file system.

mount_ufs(1M)

remount

Remounts a read-only file system as read-write (using the `rw` option). This option can be used only in conjunction with the `f`, `logging|nologging`, `m`, and `noatime` options. This option works only on currently mounted read-only file systems.

rq

Read-write with quotas turned on. Equivalent to `rw, quota`.

ro | rw

Read-only or read-write. Default is `rw`.

suid | nosuid

Allow or disallow `setuid` execution. The default is `suid`. This option can also be used when mounting devices.

-O

Overlay mount. Allow the file system to be mounted over an existing mount point, making the underlying file system inaccessible. If a mount is attempted on a pre-existing mount point without setting this flag, the mount will fail, producing the error "device busy".

-S *attribute_list*

See the definition of the `-S` option in the **OPTIONS** section of the `mount(1M)` man page.

FILES

`/etc/mnttab` Table of mounted file systems
`/etc/vfstab` List of default parameters for each file system
`/etc/security/tsol/vfstab_adjunct`
Mount-time attributes for file systems

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

The `-o quota` option has been removed; the `nodevices` and `nopriv` options have been added.

Mount-time security attributes can be specified for file systems whose objects do not have any attributes (such as user and group IDs) and for file systems that do not have the Trusted Solaris extended security attributes (such as sensitivity labels). Trusted Solaris security policy applies when mounting.

`mount` must run with the `sys_mount` privilege.

mount_ufs(1M)

To succeed in all cases, the mount command needs the privileges: file_mac_read, file_dac_read, file_mac_write, file_dac_write, file_mac_search, file_dac_search, net_privaddr, proc_setsl, sys_mount, and sys_trans_label.

When mounting a UFS file system, mount should assert the sys_fs_config privilege. Otherwise, the mount succeeds, but logging is not enabled/disabled, errno is set to EPERM, and the user sees an error message.

mount(1M), mountall(1M), mountd(1M), mount(2), stat(2), mnttab(4), vfstab(4), vfstab_adjunct(4)

fsck(1M), fsck_ufs(1M), attributes(5), largefile(5)

Since the root (/) file system is mounted read-only by the kernel during the boot process, only the remount option (and options that can be used in conjunction with remount) affect the root (/) entry in the /etc/vfstab file.

named(1M)

NAME	in.named, named – Internet domain name server	
SYNOPSIS	in.named [-d <i>debuglevel</i>] [-q] [-r] [-f] [-p <i>remote/local-port</i>] [-w <i>dirname</i>] [[-b -c] <i>configfile</i>]	
DESCRIPTION	<p>in.named is the Internet domain name server. in.named spawns the named-xfer process whenever it needs to perform a zone transfer. See named-xfer(1M).</p> <p>The in.named name service is used by hosts on the Internet to provide access to the Internet distributed naming database. See <i>RFC 1034</i> and <i>RFC 1035</i> for more information on the Internet domain name system.</p> <p>With no arguments, in.named reads the default configuration file /etc/named.conf for any initial data, and listens for queries. Any additional arguments beyond those shown in the SYNOPSIS section are interpreted as the names of configuration files. If multiple configuration files are specified, only the last is used.</p> <p>The name server reads the configuration file to obtain instructions on where to find its initial data.</p> <p>In the Trusted Solaris environment, in.named listens for input requests on a multilevel port (MLP) and sends responses to the DNS client at the sensitivity label of the client's request. Thus, though in.named runs at the sensitivity label ADMIN_LOW, it can accept requests at any sensitivity label. in.named can also serve DNS clients and communicate with other DNS name servers on either Trusted Solaris hosts or non-trusted hosts.</p> <p>The DNS name server running on a Trusted Solaris machine is viewed as a supplier of public information, and the name database that it maintains is considered trusted. in.named requires the trusted path attribute, and it requires that the /etc/named.boot file, zone files, and other configuration files that it uses be at the sensitivity label ADMIN_LOW. As part of the name database, these files and their contents are also considered trusted; thus in.named can query any DNS name server specified in the files. The DNS name servers specified in these files may reside on either Trusted Solaris hosts or non-trusted hosts.</p>	
OPTIONS	<ul style="list-style-type: none"> -b <i>configfile</i> Use <i>configfile</i> rather than /etc/named.conf. This option allows filenames to begin with a leading dash. -c <i>configfile</i> Use <i>configfile</i> rather than /etc/named.conf. This option allows filenames to begin with a leading dash. -d <i>level</i> Print debugging information. <i>level</i> is a number indicating the level of messages printed. -f Run this process in the foreground. The process will not fork(2). By default, in.named runs as a daemon in the background. -p <i>remote/local-port</i> Use different port numbers. The default is the standard port number as returned 	

bygetservbyname(3SOCKET) for service domain.

The `-p` argument can specify up to two port numbers.

The specification of two port numbers requires a `/'`

(slash) separator. In this case, the first port is used

when contacting remote servers, and the second one is the service port bound by the local instance of

`in.named`. This option is used mostly for debugging purposes.

`-q` Trace all incoming queries. Note: this option is ignored in favor of the boot file directive, `options query-log`, when both options are used.

`-r` Turns recursion off in the server. Answers can come only from local (primary or secondary) zones. This option can be used on root servers. Note: This option will probably be eventually abandoned in favor of the boot file directive, `options no-recursion`.

`-w dirname` Change the current working directory of `in.named` to `dirname`.

`/etc/named.conf`
File
Directives

The following is a simple configuration file `/etc/named.conf` containing directives to guide the `in.named` process at startup time.

```
options {
    directory    "/usr/local/adm/named";
    pid-file     "/var/named/named.pid";
    named-xfer   "/usr/sbin/named-xfer";
    forwarders   {
        10.0.0.78;
        10.2.0.78;
    };
    transfers-in 10;
    forward only;
    fake-iquery yes;
    pollfd-chunk-size 20;
};

logging {
    category lame-servers { null; };
    category cname { null; };
};

zone "." in {
    type hint;
    file "root.cache";
};

zone "cc.berkeley.edu" in {
    type slave;
    file "128.32.137.3";
    masters { 128.32.137.8; };
};
```

named(1M)

```
zone "6.32.128.in-addr.arpa" in {
    type slave;
    file "128.32.137.3";
    masters { 128.32.137.8; };
};
zone "0.0.127.in-addr.arpa" in {
    type master;
    file "master/db.127";
};

zone "berkeley.edu" in {
    type master;
    file "berkeley.edu.zone";
};

zone "32.128.in-addr.arpa" in {
    type master;
    file "ucbhosts.rev";
};
```

The configuration file consists of sections and comments. Sections end with a `';`' and contain statements which are enclosed in `{ }` and may span multiple lines. The following sections are supported: options, zone, server, logging, acl, include, and key.

Comments Syntax

The following are examples of comments syntax in BIND 8.1:

```
/* This is a BIND comment as in C */
// This is a BIND comment as in C++
# This is a BIND comment as in common Unix shells and perl
```

WARNING: you cannot use the semicolon character `(;)` to start a comment.

Options Section

The syntax of the options section is as follows:

```
options {
    [ directory path_name; ]
    [ named-xfer path_name; ]
    [ pid-file path_name; ]
    [ auth-nxdomain yes_or_no; ]
    [ fake-iquery yes_or_no; ]
    [ fetch-glue yes_or_no; ]
    [ multiple-cnames yes_or_no; ]
    [ notify yes_or_no; ]
    [ recursion yes_or_no; ]
    [ forward ( only | first ); ]
    [ forwarders { [ in_addr ; [ in_addr ; ... ] ] }; ]
    [ check-names ( master | slave | response ) ( warn | fail | ignore ); ]
    [ allow-query { address_match_list }; ]
    [ allow-transfer { address_match_list }; ]
    [ listen-on [ port ip_port ] { address_match_list }; ]
    [ query-source [ address ( ip_addr | * ) ] [ port ( ip_port | * ) ] ; ]
    [ max-transfer-time-in number; ]
    [ transfer-format ( one-answer | many-answers ); ]
    [ transfers-in number; ]
}
```



```

[ transfers-out number; ]
[ transfers-per-ns number; ]
[ coresize size_spec ; ]
[ datasize size_spec ; ]
[ files size_spec ; ]
[ stacksize size_spec ; ]
[ clean-interval number; ]
[ interface-interval number; ]
[ scan-interval number; ]
[ topology { address_match_list }; ]
};

```

Definitions and Use of Options

The options section sets up global options to be used by BIND. This section may appear at only once in a configuration file; if more than one occurrence is found, the first occurrence determines the actual options used, and a warning will be generated. If there is no options section, an options block with each option set to its default will be used.

Pathnames

directory	The working directory of the server. Any non-absolute pathnames in the configuration file will be taken as relative to this directory. The default location for most server output files (for example, "named.run") is this directory. If a directory is not specified, the working directory defaults to ".", the directory from which the server was started. The directory specified should be an absolute path.
named-xfer	The pathname to the named-xfer program that the server uses for inbound zone transfers. If not specified, the default is operating system dependent, for example, "/usr/sbin/named-xfer").
pid-file	The pathname of the file the server writes its process ID in. If not specified, the default is operating system dependent, but is usually "/var/run/named.pid" or "/etc/named.pid". The pid-file is used by programs like "ndc" that want to send signals to the running nameserver.

Boolean Options

auth-nxdomain	If yes, then the AA bit is always set on NXDOMAIN responses, even if the server is not actually authoritative. The default is yes. Do not turn off auth-nxdomain unless you are sure you know what you are doing, as some older software will not like it.
fake-iquery	If yes, the server will simulate the obsolete DNS query type IQUERY. The default is no.
fetch-glue	If yes (the default), the server will fetch "glue" resource records it does not have when constructing the additional data section of a response. fetch-glue no can be used in conjunction with recursion no to prevent the server's cache from growing or becoming corrupted (at the cost of requiring more work from the client).

named(1M)

	multiple-cnames	If yes, then multiple CNAME resource records will be allowed for a domain name. The default is no. Allowing multiple CNAME records is against standards and is not recommended. Multiple CNAME support is available because previous versions of BIND allowed multiple CNAME records, and these records have been used for load balancing by a number of sites.
	notify	If yes (the default), DNS NOTIFY messages are sent when a zone the server is authoritative for changes. The use of NOTIFY speeds convergence between the master and its slaves. Slave servers that receive a NOTIFY message and understand it will contact the master server for the zone and see if they need to do a zone transfer, and if they do, they will initiate it immediately. The notify option may also be specified in the zone section, in which case it overrides the options notify statement.
	recursion	If yes, and a DNS query requests recursion, then the server will attempt to do all the work required to answer the query. If recursion is not on, the server will return a referral to the client if it doesn't know the answer. The default is yes. See also fetch-glue above.
Forwarding	The forwarding facility can be used to create a large sitewide cache on a few servers, reducing traffic over links to external name servers. It can also be used to allow queries by servers that do not have direct access to the Internet, but wish to look up exterior names anyway. Forwarding occurs only on those queries for which the server is not authoritative, and it does not have the answer in its cache.	
	forward	This option is only meaningful if the forwarders list is not empty. A value of first, the default, causes the server to query the forwarders first, and if that doesn't answer the question, the server will then look for the answer itself. If only is specified, the server will only query the forwarders.
	forwarders	Specifies the IP addresses to be used for forwarding. The default is the empty list (no forwarding).
	Future versions of BIND 8 will provide a more powerful forwarding system. The syntax described above will continue to be supported.	
Name Checking	The server can check domain names based upon their expected client contexts. For example, a domain name used as a hostname can be checked for compliance with the valid hostnames defined in the RFCs. Three checking methods are available:	
	ignore	No checking is done.
	warn	Names are checked against their expected client contexts. Invalid names are logged, but processing continues normally.
	fail	Names are checked against their expected client contexts. Invalid names are logged, and the offending data is rejected.

The server can check names in three areas: master zone files, slave zone files, and in responses to queries the server has initiated. If `check-names` response fail has been specified, and answering the client's question would require sending an invalid name to the client, the server will send a `REFUSED` response code to the client.

The defaults are:

```
check-names master fail;
check-names slave warn;
check-names response ignore;
```

`check-names` may also be specified in the zone section, in which case it overrides the options `check-names` statement. When used in a zone section, the area is not specified (because it can be deduced from the zone type).

Access Control

Access to the server can be restricted based on the IP address of the requesting system. See `address_match_list` for details on how to specify IP address lists.

<code>allow-query</code>	Specifies which hosts are allowed to ask ordinary questions. <code>allow-query</code> may also be specified in the zone section, in which case it overrides the options <code>allow-query</code> statement. If not specified, the default is to allow queries from all hosts.
<code>allow-transfer</code>	Specifies which hosts are allowed to receive zone transfers from the server. <code>allow-transfer</code> may also be specified in the zone section, in which case it overrides the options <code>allow-transfer</code> statement. If not specified, the default is to allow transfers from all hosts.

Interfaces

The interfaces and ports that the server will answer queries from may be specified using the `listen-on` option. `listen-on` takes an optional port, and an `address_match_list`. The server will listen on all interfaces allowed by the address match list. If a port is not specified, port 53 will be used.

Multiple `listen-on` statements are allowed. For example,

```
listen-on { 5.6.7.8; };
listen-on port 1234 { !1.2.3.4; 1.2/16; };
```

If no `listen-on` is specified, the server will listen on port 53 on all interfaces.

Query Address

If the server does not know the answer to a question, it will query other name servers. `query-source` specifies the address and port used for such queries. If address is `*` or is omitted, a wildcard IP address (`INADDR_ANY`) will be used. If port is `*` or is omitted, a random unprivileged port will be used. The default is:

```
query-source address * port *;
```

Note: `query-source` currently applies only to UDP queries; TCP queries always use a wildcard IP address and a random unprivileged port.

named(1M)

Zone Transfers	max-transfer-time-in	Inbound zone transfers (named-xfer processes) running longer than this many minutes will be terminated. The default is 120 minutes.
	transfer-format	The server supports two zone transfer methods. one-answer uses one DNS message per resource record transferred. many-answers packs as many resource records as possible into a message. many-answers is more efficient, but is only known to be understood by BIND 8.1 and patched versions of BIND 4.9.5. The default is one-answer. transfer-format may be overridden on a per-server basis by using the server section.
	transfers-in	The maximum number of inbound zone transfers that can be running concurrently. The default value is 10. Increasing transfers-in may speed up the convergence of slave zones, but it also may increase the load on the local system.
	transfers-out	This option will be used in the future to limit the number of concurrent outbound zone transfers. It is checked for syntax, but is otherwise ignored.
	transfers-per-ns	The maximum number of inbound zone transfers (named-xfer processes) that can be concurrently transferring from a given remote name server. The default value is 2. Increasing transfers-per-ns may speed up the convergence of slave zones, but it also may increase the load on the remote name server. transfers-per-ns may be overridden on a per-server basis by using the transfers statement in the server section.
Resource Limits	The server's usage of many system resources can be limited. Some operating systems do not support some of the limits, and a warning will be generated if an unsupported limit is set in the configuration file.	
	Scaled values are allowed when specifying resource limits. For example, 1G can be used instead of 1073741824 to specify a limit of one gigabyte, unlimited requests unlimited use, or the maximum available amount. Default uses the limit that was in force when the server was started. See ulimit(1) for a discussion of ulimit -a (ksh only) for defaults.	
	coresize	The maximum size of a core dump. The default is system dependent.
	datasize	The maximum amount of data memory the server may use. The default is system dependent.

named(1M)

files	The maximum number of files that the server may have open concurrently. The default is system dependent.
stacksize	The maximum amount of stack memory the server may use. The default is system dependent.

Topology

All other things being equal, when the server chooses a name server to query from a list of name servers, it prefers the one that is topologically closest to itself. The topology statement takes an `address_match_list` and interprets it in a special way. Each top-level list element is assigned a distance. Non-negated elements get a distance based on their position in the list, where the closer the match is to the start of the list, the shorter the distance is between it and the server. A negated match will be assigned the maximum distance from the server. If there is no match, the address will get a distance which is further than any non-negated list element, and closer than any negated element. For example,

```
topology {
    10/8;
    !1.2.3/24;
    { 1.2/16; 3/8; };
```

}; will prefer servers on network 10 the most, followed by hosts on network 1.2.0.0 (netmask 255.255.0.0) and network 3, with the exception of hosts on network 1.2.3 (netmask 255.255.255.0), which is preferred least of all. The default topology is

```
topology { localhost; localnets; };
```

The Server Section

The syntax of the server section is as follows:

```
server ip_addr {
    [ bogus yes_or_no; ]
    [ transfers number; ]
    [ transfer-format ( one-answer | many-answers ); ]
    [ keys { key_id [key_id ... ] }; ]
```

}; The server statement defines the characteristics to be associated with a remote name server.

If you discover that a server is giving out bad data, marking it as bogus will prevent further queries to it. The default value is no.

The server supports two zone transfer methods. The first, `one-answer`, uses one DNS message per resource record transferred. `many-answers` packs as many resource records as possible into a message. `many-answers` is more efficient, but is only known to be understood by BIND 8.1 and patched versions of BIND 4.9.5. You can specify which method to use for a server with the `transfer-format` option. If `transfer-format` is not specified, the `transfer-format` specified by the options statement will be used.

The transfers will be used in a future release of the server to limit the number of concurrent inbound zone transfers from the specified server. It is checked for syntax but is otherwise ignored.

named(1M)

The keys statement is intended for future use by the server. It is checked for syntax but is otherwise ignored.

The Zone Section

The syntax of the zone section is as follows:

```
zone domain_name [ ( in | hs | hesiod | chaos ) ] {
    type master;
    file path_name;
    [ check-names ( warn | fail | ignore ); ]
    [ allow-update { address_match_list }; ]
    [ allow-query { address_match_list }; ]
    [ allow-transfer { address_match_list }; ]
    [ notify yes_or_no; ]
    [ also-notify { ip_addr; [ ip_addr; ... ] }; ]
};

zone domain_name [ ( in | hs | hesiod | chaos ) ] {
    type ( slave | stub );
    [ file path_name; ]
    masters { ip_addr; [ ip_addr; ... ] };
    [ check-names ( warn | fail | ignore ); ]
    [ allow-update { address_match_list }; ]
    [ allow-query { address_match_list }; ]
    [ allow-transfer { address_match_list }; ]
    [ max-transfer-time-in number; ]
    [ notify yes_or_no; ]
    [ also-notify { ip_addr; [ ip_addr; ... ] }; ]
};

zone . [ ( in | hs | hesiod | chaos ) ] {
    type hint;
    file path_name;
    [ check-names ( warn | fail | ignore ); ]
};
```

Zone types are defined as follows:

- | | |
|--------|--|
| master | The master copy of the data in a zone. |
| slave | A slave zone is a replica of a master zone. The masters list specifies one or more IP addresses that the slave contacts to update its copy of the zone. If file is specified, then the replica will be written to the file. Use of file is recommended, since it often speeds server startup and eliminates a needless waste of bandwidth. |
| stub | A stub zone is like a slave zone, except that it replicates only the NS records of a master zone instead of the entire zone. |
| hint | The initial set of root name servers is specified using a hint zone. When the server starts up, it uses the root hints to find a root name server and get the most recent list of root name servers. |

Note: previous releases of BIND used the term `primary` for a master zone, `secondary` for a slave zone, and `cache` for a hint zone.

The zone's name may optionally be followed by a class. If a class is not specified, class in is used.

Zone options are described as follows:

check-names	See Name Checking.
allow-query	See the description of allow-query in the Access Control section.
allow-update	Specifies which hosts are allowed to submit dynamic DNS updates to the server. The default is to deny updates from all hosts.
allow-transfer	See the description of allow-transfer in the Access Control section.
max-transfer-time-in	See the description of max-transfer-time-in in the Zone Transfers section.
notify	See the description of notify in the Boolean Options section.
also-notify	also-notify is only meaningful if notify is active for this zone.

The set of machines that will receive a DNS NOTIFY message for this zone is made up of all the listed name servers for the zone (other than the primary master) plus any IP addresses specified with also-notify. also-notify is not meaningful for stub zones. The default is the empty list.

The Logging Section

The syntax of the logging section is as follows:

```
logging {
    [ channel channel_name {
        ( file path_name
          [ versions ( number | unlimited ) ]
          [ size size_spec ]
        | syslog ( kern | user | mail | daemon | auth | syslog | lpr |
                  news | uucp | cron | authpriv | ftp |
                  local0 | local1 | local2 | local3 |
                  local4 | local5 | local6 | local7 )
        | null );

        [ severity ( critical | error | warning | notice |
                    info | debug [ level ] | dynamic ); ]
        [ print-category yes_or_no; ]
        [ print-severity yes_or_no; ]
        [ print-time yes_or_no; ]
    }; ]

    [ category category_name {
        channel_name; [ channel_name; ... ]
    }; ]
    ...
}
```

named(1M)

```
};
```

The `logging` statement configures a wide variety of logging options for the name server. Its channel phrase associates output methods, format options and severity levels with a name that can then be used with the category phrase to select how various classes of messages are logged.

Only one logging statement is used to define as many channels and categories as are wanted. If there are multiple logging statements in a configuration, the first defined determines the logging, and warnings are issued for the others. If there is no logging statement, the default logging configuration will be:

```
logging {
    category default { default_syslog; default_debug; };
    category panic { default_syslog; default_stderr; };
    category packet { default_debug; };
    category eventlib { default_debug; };
};
```

The Channel Phrase

All log output goes to one or more "channels"; you can make as many of them as you want.

Every channel definition must include a clause that says whether messages selected for the channel go to a file, to a particular `syslog` facility, or are discarded. It can optionally also limit the message severity level that will be accepted by the channel (default is "info"), and whether to include a named-generated time stamp, the category name and/or severity level (default is not to include any).

The word `null` as the destination option for the channel will cause all messages sent to it to be discarded; other options for the channel are meaningless.

The file clause can include limitations both on how large the file is allowed to become, and how many versions of the file will be saved each time the file is opened.

The size option for files is simply a hard ceiling on log growth. If the file ever exceeds the size, then named will just not write anything more to it until the file is reopened; exceeding the size does not automatically trigger a reopen. The default behavior is to not limit the size of the file.

If you use the version logfile option, then named will retain that many backup versions of the file by renaming them when opening. For example, if you choose to keep 3 old versions of the file "lamers.log" then just before it is opened lamers.log.1 is renamed to lamers.log.2, lamers.log.0 is renamed to lamers.log.1, and lamers.log is renamed to lamers.log.0. No rolled versions are kept by default. The unlimited keyword is synonymous with 99 in current BIND releases.

The argument for the `syslog()` clause is a `syslog()` facility as described in the `syslog(3C)` manual page. How `syslogd(1M)` will handle messages sent to this facility is described in the `syslog.conf(4)` manual page. If you have a system which

uses a very old version of `syslog()` that only uses two arguments to the `openlog()` function, then this clause is silently ignored.

The severity clause works like the "priorities" to `syslog()`, except that they can also be used if you are writing straight to a file rather than using `syslog()`. Messages which are not at least of the severity level given will not be selected for the channel; messages of higher severity levels will be accepted.

If you are using `syslog()`, then the `syslog.conf` priorities will also determine what eventually passes through. For example, defining a channel facility and severity as `daemon` and `debug` but only logging `daemon.warning` by way of `syslog.conf` will cause messages of severity `info` and `notice` to be dropped. If the situation were reversed, with `named` writing messages of only `warning` or higher, then `syslogd` would print all messages it received from the channel.

The server can supply extensive debugging information when it is in debugging mode. If the server's global debug level is greater than zero, then debugging mode will be active. The global debug level is set either by starting the server with the `-d` option followed by a positive integer, or by sending the server the `SIGUSR1` signal (for example, by using `ndc trace`). The global debug level can be set to zero, and debugging mode turned off, by sending the server the `SIGUSR2` signal (`ndc notrace`). All debugging messages in the server have a debug level, and higher debug levels give more more detailed output. Channels that specify a specific debug severity, for example:

```
channel specific_debug_level {
    file "foo";
    severity debug 3;
};
```

will get debugging output of level 3 or less any time the server is in debugging mode, regardless of the global debugging level. Channels with dynamic severity use the server's global level to determine what messages to print.

If `print-time` has been turned on, then the date and time will be logged. `print-time` may be specified for a `syslog()` channel, but is usually pointless since `syslog()` also prints the date and time. If `print-category` is requested, then the category of the message will be logged as well. Finally, if `print-severity` is on, then the severity level of the message will be logged. The `print-options` may be used in any combination, and will always be printed in the following order: `time`, `category`, `severity`. Here is an example where all three `print-options` are on:

```
28-Apr-1997 15:05:32.863 default: notice: Ready to answer queries.
```

There are four predefined channels that are used for default logging for `in.named` as follows. How they are used is described in the next section.

```
channel default_syslog {
    syslog daemon;      # send to syslog's daemon facility
    severity info;      # only send priority info and higher
};
```

named(1M)

```
channel default_debug {
    file "named.run";      # write to named.run in the working directory
    severity dynamic;      # log at the server's current debug level
};

channel default_stderr { # writes to stderr
    file "<stderr>";      # this is illustrative only;
    # there's currently   # no way of specifying an internal file
                          # descriptor in the configuration language.
    severity info;        # only send priority info and higher
};

channel null {
    null;                  # toss anything sent to this channel
};
```

Once a channel is defined, it cannot be redefined. Thus you cannot alter the built-in channels directly, but you can modify the default logging by pointing categories at channels you have defined.

The Category Phase

There are many categories, so you can send the logs you want to see wherever you want, without seeing logs you do not want. If you do not specify a list of channels for a category, then log messages in that category will be sent to the default category instead. If do not specify a default category, the following "default default" is used:

```
category default { default_syslog; default_debug; };
```

For example, if you want to log security events to a file, but you also want keep the default logging behavior, specify the following:

```
channel my_security_channel {
    file "my_security_file";
    severity info;
};
category security { my_security_channel; default_syslog; default_debug; };
```

To discard all messages in a category, specify the null channel:

```
category lame-servers { null; };
category cname { null; };
```

The following categories are available:

default The catch-all. Many things still are not classified into categories, and they all end up here. Also, if you do not specify any channels for a category, the default category is used instead. If you do not define the default category, the following definition is used:

```
category default { default_syslog; default_debug; };
```

config	High-level configuration file processing.
parser	Low-level configuration file processing.
queries	A short log message is generated for every query the server receives.
lame-servers	Messages like "Lame server on ..."
statistics	Statistics.
panic	<p>If the server has to shut itself down due to an internal problem, it will log the problem in this category as well as in the problem's native category. If you do not define the panic category, the following definition is used:</p> <pre>category panic { default_syslog; default_stderr; };</pre>
update	Dynamic updates.
ncache	Negative caching.
xfer-in	Zone transfers the server is receiving.
xfer-out	Zone transfers the server is sending.
db	All database operations.
eventlib	<p>Debugging info from the event system. Only one channel may be specified for this category, and it must be a file channel. If you do not define the eventlib category, the following definition is used:</p> <pre>category eventlib { default_debug; };</pre>
packet	<p>Dumps of packets received and sent. Only one channel may be specified for this category, and it must be a file channel. If you do not define the packet category, the following definition is used:</p> <pre>category packet { default_debug; };</pre>
notify	The NOTIFY protocol.
cname	Messages like "... points to a CNAME".
security	Approved/unapproved requests.
os	Operating system problems.
insist	Internal consistency check failures.
maintenance	Periodic maintenance events.
load	Zone loading messages.

named(1M)

	<p>response-checks Messages arising from response checking, such as "Malformed response ...", "wrong ans. name ...", "unrelated additional info ...", "invalid RR type ...", and "bad referral ...".</p>
The Key Section	<p>The syntax of the key section is as follows:</p> <pre>key key_id { algorithm algorithm_id; secret secret_string; };</pre> <p>The key section defines a key ID which can be used in a server section to associate an authentication method with a particular name server.</p> <p>A key ID must be created with the key statement before it can be used in a server definition.</p> <p>The <code>algorithm_id</code> is a string that specifies a security/authentication algorithm. <code>secret_string</code> is the secret to be used by the algorithm.</p> <p>The key statement is intended for future use by the server. It is checked for syntax but is otherwise ignored.</p>
The Include Section	<p>The syntax of the include section is as follows:</p> <pre>include path_name;</pre> <p>The include statement inserts the specified file at the point that the include statement is encountered. It cannot be used within another statement, though, so a line such as <code>acl internal_hosts { "include internal_hosts.acl" }</code> is not allowed. Use <code>include</code> to break the configuration up into easily-managed chunks. For example:</p> <pre>include "/etc/security/keys.bind"; include "/etc/acls.bind";</pre> <p>could be used at the top of a BIND configuration file in order to include any ACL or key information.</p> <p>Be careful not to type <code>#include</code>, like you would in a C program, because <code>#</code> is used to start a comment.</p>
The ACL Format	<p>The syntax of the ACL section is as follows:</p> <pre>acl name { address_match_list };</pre> <p>The <code>acl</code> statement creates a named address match list. It gets its name from a primary use of address match lists: Access Control Lists (ACLs).</p>

Note that an address match list's name must be defined with `acl` before it can be used elsewhere; no forward references are allowed.

The following ACLs are built-in:

<code>any</code>	Allows all hosts.
<code>none</code>	Denies all hosts.
<code>localhost</code>	Allows the IP addresses of all interfaces on the system.
<code>localnets</code>	Allows any host on a network for which the system has an interface.

Zone File Format

The zone files are also known as the authoritative master files (data files) for a zone. In the boot file, references were made to these files as part of the specification of any primary directives.

Two classes of entries populate the zone files, directives and resource records. The start of the zone file is likely to contain one or two directives that establish a context that modifies the way subsequent records are interpreted.

Resource records for a zone determine how a zone is managed by establishing zone characteristics. For example, one type of zone record establishes the zone's mailbox information.

The very first record of each zone file should be a Start-of-Authority record (SOA) for a zone. A multiple-line SOA record is presented below. The meaning of the values in this sample will become clearer with the help of a list that describes the purpose of each field in the zone record (see the SOA list subitem under the `rr-type` list item in, Format of Resource Records in Zone Files).

```
@ IN SOA ucbvax.Berkeley.EDU. rwh.ucbvax.Berkeley.EDU. (
1989020501 ;serial
10800      ;refresh
3600       ;retry
3600000    ;expire
86400 )    ;minimum
```

Resource records normally end at the end of a line, but may be continued across lines between opening and closing parentheses (as demonstrated by the preceding sample).

Comments are introduced by semicolons. They continue to the end of the line.

Directives in Zone Files

There are two control directives that help determine how the zone file is processed, `$INCLUDE` and `$ORIGIN`.

The `$INCLUDE` directive refers to still another file within which zone characteristics are described. Such files typically contain groups of resource records, but they may also contain further directives.

named(1M)

The `$ORIGIN` directive establishes a current origin that is appended to any domain values that do not end with a `'.'` (dot). The placeholder domain represents the first resource record field as shown in Format of Resource Records in Zone Files. The format for these directives is:

```
$INCLUDE filename opt-current-domain
$ORIGIN current-domain
```

where:

<code>current-domain</code>	Specifies the value of the current origin that remains in effect for this configuration file unless a subsequent <code>\$ORIGIN</code> directive overrides it for the remaining portion of the file.
-----------------------------	--

<code>filename</code>	Specifies a file, the contents of which are, in effect, incorporated into the configuration file at the location of the corresponding <code>\$INCLUDE</code> directive.
-----------------------	---

<code>opt-current-domain</code>	Optionally defines a current origin that is applicable only to the records residing in the specified file in the corresponding <code>\$INCLUDE</code> directive. This directive overrides the origin given in a preceding <code>\$ORIGIN</code> directive, but only for the scope of the included text. See also <code>current-domain</code> . Neither the
---------------------------------	--

`opt-current-domain` argument of `$INCLUDE` nor the `$ORIGIN` directive in the included file can affect the current origin in effect for the remaining records in the main configuration file (as defined by those `$ORIGIN` directives that reside there).

Format of Resource Records in Zone Files

The format of the resource records is:

```
domain opt-ttl opt-class rr-type rr-data...where:
```

<code>domain</code>	Specifies the domain being described by the current line and any following lines that lack a value for this field. Beware of any domain values that you enter without full qualification, because the value of the current origin will be appended to them. The value of the current origin is appended when domain does not end with a dot.
---------------------	--

A domain value specified as the symbol `@` is replaced with the value of the current origin. The `current-domain` or any locally-overriding `opt-current-domain` value is used as its replacement. (For a discussion of these placeholders, see the earlier discussion of the `$ORIGIN` and `$INCLUDE` directives.)

A domain value specified as a `'.'` (dot) represents the root.

opt-ttl	Specifies the number of seconds corresponding to the time-to-live value applicable to the zone characteristic that is defined in the remaining fields. This field is optional. It defaults to zero. Zero is interpreted as the minimum value specified in the SOA record for the zone.
opt-class	Specifies the object address type; currently only one type is supported, IN, for objects connected to the Internet.
rr-type rr-data ...	Specifies values that describe a zone characteristic. Permissible rr-type and other field values are listed below. The field values are listed in the order that they must appear.
	A address Specifies the host address (in dotted-quad format). DCE or AFS server.
	CNAME canonical-name Specifies in a domain-name format the canonical name for the alias (domain).
	HINFO cpu-type OS-type Host information supplied in terms of a CPU type and an OS type.
	MX preference mail-exchanger Specifies in domain-name format a mail exchanger preceded by a preference value (between 0 and 32767), with lower numeric values representing higher logical preferences.
	NS authoritative-server Specifies in domain-name format an authoritative name server.
	NULL Specifies a null zone record.
	PTR domain-pointer Specifies in domain-name format a domain name pointer.
	RP mailbox txt-referral Offers details about how to reach a responsible person for the domain name.
	retry expire ttl
	SOA host-domain maintainer-addr serial- no refresh Establishes the start of a zone of authority in terms of the domain of the originating host (host-domain), the domain address of the maintainer (maintainer-addr), a serial number (serial-no), the refresh period in seconds (refresh), the retry

named(1M)

period in seconds (retry), the expiration period in seconds (expire), and the minimum time-to-live period in seconds (ttl). See RFC 1035.

The serial number should be changed each time the master file is changed. Secondary servers check the serial number at intervals specified by the refresh time in seconds; if the serial number changes, a zone transfer will be done to load the new data.

If a master server cannot be contacted when a refresh is due, the retry time specifies the interval at which refreshes should be attempted. If a master server cannot be contacted within the interval given by the expire time, all data from the zone is discarded by secondary servers. The minimum value is the time-to-live used by records in the file with no explicit time-to-live value.

The serial number can be given as a dotted number. However, this is a very unwise thing to do, since the translation to normal integers is via concatenation rather than multiplication and addition. You could spell out the year, month, day of month, and 0.99 version number and still fit it inside the unsigned 32-bit size of this field. This strategy should work for the foreseeable future (but is questionable after the year 4293).

For more detailed information, see *RFC 883*.

rr-data ... See the description of rr-type.

Consult *Name Server Operations Guide for BIND* for further information about the supported types of resource records.

EXIT STATUS The `in.named` process returns the following exit values:

0	Successful completion.
1	An error occurred.

FILES In the Trusted Solaris environment, these files have a sensitivity label of `ADMIN_LOW`:

<code>/etc/named.conf</code>	Name server configuration boot file.
<code>/etc/named.pid</code>	The process ID (on older systems).
<code>/var/tmp/named.run</code>	Debug output.
<code>/var/tmp/named.stats</code>	Nameserver statistics data.
<code>/var/tmp/nameddump.db</code>	Dump of the name servers database.
<code>/var/tmp/named.pid</code>	The process ID (on newer systems).

named(1M)

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

in.named accepts requests at any sensitivity label and replies at the sensitivity label of the client's request. in.named can serve DNS clients and can communicate with other DNS servers that are on Trusted Solaris hosts or non-trusted hosts.

Files used by in.named should be protected from unauthorized access by having the sensitivity label ADMIN_LOW.

Invoking in.named requires the trusted path attribute, an effective UID of 0, a process sensitivity label of ADMIN_LOW, and the following privileges: net_mac_read, net_privaddr, net_upgrade_sl, proc_setclr, sys_trans_label, sys_net_config, and sys_config.

Trusted Solaris 8 4/01 Reference Manual Solaris 9 Reference Manual

fork(2), resolver(3RESOLV), listen(3SOCKET), resolv.conf(4)

kill(1), named-xfer(1M), syslogd(1M), signal(3C), syslog(3C),
getservbyname(3SOCKET), syslog.conf(4), attributes(5)

Braden, R. (Editor), *Requirements for Internet Hosts - Applications and Support*, RFC 1123, Internet Engineering Task Force - Network Working Group, October 1989.

Mockapetris, Paul, *Domain Names - Concepts and Facilities*, RFC 1034, , Network Information Center, SRI International, Menlo Park, Calif., November 1987.

Mockapetris, Paul, *Domain Names - Implementation and Specification*, RFC 1035, Network Information Center, SRI International, Menlo Park, Calif., November 1987.

Mockapetris, Paul, *Domain System Changes and Observations*, RFC 973, Network Information Center, SRI International, Menlo Park, Calif., January 1986.

Partridge, Craig, *Mail Routing and the Domain System*, RFC 974, Network Information Center, SRI International, Menlo Park, Calif., January 1986.

Vixie, Paul, Dunlap, Keven J., Karels, Michael J., *Name Server Operations Guide for BIND* (public domain), Internet Software Consortium, 1995.

NOTES

The following signals have the specified effect when sent to the server process using the kill(1) command:

SIGHUP Causes the server to read /etc/named.conf and reload the database.

SIGHUP Also causes the server to check the serial number on all secondary zones. Normally the serial numbers are only checked at the

named(1M)

	intervals specified by the SOA record at the start of each zones-definition file.
SIGINT	Dumps the current database and cache to <code>/var/tmp/nameddump.db</code> .
SIGIOT	Dumps statistical data into <code>/var/tmp/named.stats</code> . Statistical data are appended to the file.
SIGUSR1	Turns on debugging at the lowest level when received the first time; receipt of each additional SIGUSR1 signal causes the server to increment the debug level.
SIGUSR2	Turns off debugging completely.
SIGWINCH	Toggles logging of all incoming queries through the syslog system daemon. See <code>syslogd(1M)</code> .

NAME	ndd – Get and set driver configuration parameters
SYNOPSIS	ndd [-set] <i>driver parameter</i> [<i>value</i>]
DESCRIPTION	<p>ndd gets and sets selected configuration parameters in some kernel drivers. Currently, ndd only supports the drivers that implement the TCP/IP Internet protocol family. Each driver chooses which parameters to make visible using ndd. Since these parameters are usually tightly coupled to the implementation, they are likely to change from release to release. Some parameters may be read-only.</p> <p>If the -set option is omitted, ndd queries the named <i>driver</i>, retrieves the value associated with the specified <i>parameter</i>, and prints it. If the -set option is given, ndd passes <i>value</i>, which must be specified, down to the named <i>driver</i> which assigns it to the named <i>parameter</i>.</p> <p>By convention, drivers that support ndd also support a special read-only <i>parameter</i> named “?” which can be used to list the parameters supported by the driver.</p>
EXAMPLES	<p>EXAMPLE 1 Getting Parameters Supported By The TCP Driver</p> <p>To see which parameters are supported by the TCP driver, use the following command:</p> <pre>example% ndd /dev/tcp \?</pre> <p>Note – The parameter name “?” may need to be escaped with a backslash to prevent its being interpreted as a shell meta character.</p> <p>EXAMPLE 2 Disabling packet forwarding</p> <p>The following command sets the value of the parameter <i>ip_forwarding</i> in the dual stack IP driver to zero. This disables IPv4 packet forwarding.</p> <pre>example% ndd -set /dev/ip ip_forwarding 0</pre> <p>Similarly, in order to disable IPv6 packet forwarding, the value of parameter <i>ip6_forwarding</i></p> <pre>example% ndd -set /dev/ip ip6_forwarding 0</pre> <p>EXAMPLE 3 Viewing current IP forwarding table</p> <p>To view the current IPv4 forwarding table, use the following command:</p> <pre>example% ndd /dev/ip ipv4_ire_status</pre> <p>To view the current IPv6 forwarding table, use the following command:</p> <pre>example% ndd /dev/ip ipv6_ire_status</pre>
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

ndd(1M)

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES
SunOS 5.8
Reference Manual
NOTES**

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

The -set option must inherit the `sys_net_config` privilege to set driver parameters.

`ioctl(2)`, `attributes(5)`, `arp(7P)`, `ip(7P)`, `ip6(7P)`, `tcp(7P)`, `udp(7P)`

The parameters supported by each driver may change from release to release. Like programs that read `/dev/kmem`, user programs or shell scripts that execute `ndd` should be prepared for parameter names to change.

The `ioctl()` command that `ndd` uses to communicate with drivers is likely to change in a future release. User programs should avoid making dependencies on it.

The meanings of many `ndd` parameters make sense only if you understand how the driver is implemented.

NAME	netstat – show network status														
SYNOPSIS	<pre> netstat [-anv] [-f <i>address_family</i>] netstat [-g -m -p -s -f <i>address_family</i>] [-n] [-P <i>protocol</i>] netstat -m netstat -i [-I <i>interface</i>] [-an] [-f <i>address_family</i>] [<i>interval</i>] netstat -r [-anv] [-f <i>address_family</i>] netstat -R [-anv] [-f <i>address_family</i>] netstat -M [-ns] [-f <i>address_family</i>] netstat -D [-I <i>interface</i>] [-f <i>address_family</i>] </pre>														
DESCRIPTION	<p>netstat displays the contents of various network-related data structures in various formats, depending on the options you select.</p> <p>The first form of the command displays a list of active sockets for each protocol. The second form selects one from among various other network data structures. The third form shows the state of the interfaces. The fourth form displays the routing table, the fifth form displays the routing table with extended metric information, the sixth form displays the multicast routing table, and the seventh form displays the state of DHCP on one or all interfaces.</p> <p>With no arguments, netstat prints connected sockets for PF_INET, PF_INET6, and PF_UNIX, unless modified otherwise by the -f option.</p>														
OPTIONS	<table> <tr> <td>-a</td><td>Show the state of all sockets, all routing table entries, or all interfaces, both physical and logical. Normally, sockets used by server processes are not shown. Only interface, host, network, and default routes are shown. Also, only the status of physical interfaces are shown.</td></tr> <tr> <td>-f <i>address_family</i></td><td>Limit all displays to those of the specified <i>address_family</i>. The value of <i>address_family</i> can be one of the following:</td></tr> <tr> <td>inet</td><td>For the AF_INET address family showing IPv4 information.</td></tr> <tr> <td>inet6</td><td>For the AF_INET6 address family showing IPv6 information.</td></tr> <tr> <td>unix</td><td>For the AF_UNIX address family.</td></tr> <tr> <td>-g</td><td>Show the multicast group memberships for all interfaces.</td></tr> <tr> <td>-i</td><td>Show the state of the interfaces that are used for IP traffic. Normally this shows status and statistics for the physical interfaces. When combined with the -a option, this will also report information for the logical interfaces. See ifconfig(1M).</td></tr> </table>	-a	Show the state of all sockets, all routing table entries, or all interfaces, both physical and logical. Normally, sockets used by server processes are not shown. Only interface, host, network, and default routes are shown. Also, only the status of physical interfaces are shown.	-f <i>address_family</i>	Limit all displays to those of the specified <i>address_family</i> . The value of <i>address_family</i> can be one of the following:	inet	For the AF_INET address family showing IPv4 information.	inet6	For the AF_INET6 address family showing IPv6 information.	unix	For the AF_UNIX address family.	-g	Show the multicast group memberships for all interfaces.	-i	Show the state of the interfaces that are used for IP traffic. Normally this shows status and statistics for the physical interfaces. When combined with the -a option, this will also report information for the logical interfaces. See ifconfig(1M).
-a	Show the state of all sockets, all routing table entries, or all interfaces, both physical and logical. Normally, sockets used by server processes are not shown. Only interface, host, network, and default routes are shown. Also, only the status of physical interfaces are shown.														
-f <i>address_family</i>	Limit all displays to those of the specified <i>address_family</i> . The value of <i>address_family</i> can be one of the following:														
inet	For the AF_INET address family showing IPv4 information.														
inet6	For the AF_INET6 address family showing IPv6 information.														
unix	For the AF_UNIX address family.														
-g	Show the multicast group memberships for all interfaces.														
-i	Show the state of the interfaces that are used for IP traffic. Normally this shows status and statistics for the physical interfaces. When combined with the -a option, this will also report information for the logical interfaces. See ifconfig(1M).														

netstat(1M)

	-m	Show the STREAMS statistics.
	-n	Show network addresses as numbers. <code>netstat</code> normally displays addresses as symbols. This option may be used with any of the display formats.
	-p	Show the net to media tables.
	-r	Show the routing tables. Normally, only interface, host, network, and default routes are shown, but when this option is combined with the -a option, all routes will be printed, including cache.
	-s	Show per-protocol statistics. When used with the -M option, show multicast routing statistics instead. When used with the -a option, per-interface statistics will be displayed, when available, in addition to statistics global to the system.
	-v	Verbose. Show additional information for the sockets and the routing table, including label information.
	-I <i>interface</i>	Show the state of a particular interface. <i>interface</i> can be any valid interface such as hme0 or le0. Normally, the status and statistics for physical interfaces are displayed. When this option is combined with the -a option, information for the logical interfaces is also reported.
	-M	Show the multicast routing tables. When used with the -s option, show multicast routing statistics instead.
	-P <i>protocol</i>	Limit display of statistics or state of all sockets to those applicable to <i>protocol</i> . The protocol can be one of ip, ipv6, icmp, icmpv6, igmp, udp, tcp, rawip. The command accepts protocol options only as all lowercase.
	-R	Show the routing tables with extended metric information, if any.
	-D	Show the status of DHCP configured interfaces.
OPERANDS	<i>interval</i>	If <i>interval</i> is specified, <code>netstat</code> displays interface information over the last <i>interval</i> seconds, repeating forever.
Active Sockets (First Form)	<p>The display for each active socket shows the local and remote address, the send and receive queue sizes (in bytes), the send and receive windows (in bytes), and the internal state of the protocol.</p> <p>The symbolic format normally used to display socket addresses is either <code>hostname.port</code> when the name of the host is specified, or <code>network.port</code> if a socket address specifies a network but no specific host.</p>	

TCP Sockets

The numeric host address or network number associated with the socket is used to look up the corresponding symbolic hostname or network name in the `hosts` or `networks` database.

If the network or hostname for an address is not known (or if the `-n` option is specified), the numerical network address is shown. Unspecified, or "wildcard", addresses and ports appear as `"*"`. For more information regarding the Internet naming conventions, refer to `inet(7P)` and `inet6(7P)`.

The possible state values for TCP sockets are as follows:

BOUND	Bound, ready to connect or listen.
CLOSED	Closed. The socket is not being used.
CLOSING	Closed, then remote shutdown; awaiting acknowledgment.
CLOSE_WAIT	Remote shutdown; waiting for the socket to close.
ESTABLISHED	Connection has been established.
FIN_WAIT_1	Socket closed; shutting down connection.
FIN_WAIT_2	Socket closed; waiting for shutdown from remote.
IDLE	Idle, opened but not bound.
LAST_ACK	Remote shutdown, then closed; awaiting acknowledgment.
LISTEN	Listening for incoming connections.
SYN_RECEIVED	Initial synchronization of the connection under way.
SYN_SENT	Actively trying to establish connection.
TIME_WAIT	Wait after close for remote shutdown retransmission.

Network Data Structures (Second Form)

The form of the display depends upon which of the `-g`, `-m`, `-p`, or `-s` options you select.

<code>-g</code>	Displays the list of multicast group membership.
<code>-m</code>	Displays the memory usage, for example, STREAMS mblks.
<code>-p</code>	Displays the net to media mapping table. For IPv4, the address resolution table is displayed. See <code>arp(1M)</code> . For IPv6, the neighbor cache is displayed.
<code>-s</code>	Displays the statistics for the various protocol layers.

The statistics use the MIB specified variables. The defined values for `ipForwarding` are:

<code>forwarding(1)</code>	Acting as a gateway.
<code>not-forwarding(2)</code>	Not acting as a gateway.

netstat(1M)

**Interface Status
(Third Form)**

The IPv6 and ICMPv6 protocol layers maintain per-interface statistics. If the `-a` option is specified with the `-s` option, then the per-interface statistics as well as the total sums are displayed. Otherwise, just the sum of the statistics are shown.

If you specify more than one of these options, `netstat` displays the information for each one of them.

The interface status display lists information for all current interfaces, one interface per line. If an interface is specified using the `-I` option, it displays information for only the specified interface.

The list consists of the interface name, `mtu` (maximum transmission unit, or maximum packet size) (see `ifconfig(1M)`), the network to which the interface is attached, addresses for each interface, and counter associated with the interface. The counters show the number of input packets, input errors, output packets, output errors, and collisions, respectively. For Point-to-Point interfaces, the Net/Dest field is the name or address on the other side of the link.

If the `-a` option is specified with either the `-i` option or the `-I` option, then the output includes additional information about the physical interface(s), input packets, input packets and output packets for each logical interface, for example the local IP address, associated with the physical interface(s).

If the `-n` option is specified, the list displays the IP address instead of the interface name.

If an optional *interval* is specified, the output will be continuously displayed in *interval* seconds until interrupted by the user.

The input interface is specified using the `-I` option. In this case, the list only displays traffic information in columns; the specified interface is first, the total count is second. This column list has the format of:

input		le0			output			(Total)			output	
packets	errs	packets	errs	colls	packets	errs	packets	errs	colls	packets	errs	colls
227681	0	659471	1	502	261331	0	99597	1	502			
10	0	0	0	0	10	0	0	0	0			
8	0	0	0	0	8	0	0	0	0			
10	0	2	0	0	10	0	2	0	0			

If the input interface is not specified, the first interface of address family `inet` or `inet6` will be displayed.

**Routing Table
(Fourth Form)**

The routing table display lists the available routes and the status of each. Each route consists of a destination host or network, and a gateway to use in forwarding packets. The *flags* column shows the status of the route (U if "up"), whether the route is to a gateway (G), and whether the route was created dynamically by a redirect (D). If the

**Routing Table
with Extended
Metric Information
(Fifth Form)**

-a option is specified, there will be routing entries with flags for combined routing and address resolution entries (A), broadcast addresses (B), and the local addresses for the host (L).

Interface routes are created for each interface attached to the local host; the gateway field for such entries shows the address of the outgoing interface.

The use column displays the number of packets sent using a combined routing and address resolution (A) or a broadcast (B) route. For a local (L) route, this count is the number of packets received, and for all other routes it is the number of times the routing entry has been used to create a new combined route and address resolution entry.

The *interface* entry indicates the network interface utilized for the route.

This form is the same as that of -r with the following additional information. If a route displayed has extended metric (emetric) information, it is displayed in the next line indented by a tab. Since a route may have multiple emetrics, each is displayed one line at a time. The format of display is the same as the format of specification in route(1M); that is, each field includes a keyword and, if there is value to follow, an equal sign (=) and the value. The fields are separated by commas (,). Depending on the nature of the route (that is, local or remote) and how it was entered into the routing table, there may be no emetric available for a particular route. In that case, no emetric is displayed.

**Multicast Routing
Tables (Sixth
Form)**

The multicast routing table consists of the virtual interface table and the actual routing table.

**DHCP Interface
Information
(Seventh Form)**

The DHCP interface information consists of the interface name, its current state, lease information, packet counts, and a list of flags.

The states correlate with the specifications set forth in *RFC 2131*.

Lease information includes:

- when the lease began;
- when lease renewal will begin; and
- when the lease will expire.

The flags currently defined include:

BOOTP The interface has a lease obtained through BOOTP.

BUSY The interface is busy with a DHCP transaction.

PRIMARY The interface is the primary interface. See dhcpinfo(1).

FAILED The interface is in failure state and must be manually restarted.

Packet counts are maintained for the number of packets sent, the number of packets received, and the number of lease offers declined by the DHCP client. All three

netstat(1M)

counters are initialized at zero and then incremented while obtaining a lease. The counters are reset when the period of lease renewal begins for the interface. Thus, the counters represent either the number of packets sent, received, and declined while obtaining the current lease, or the number of packets sent, received, and declined while attempting to obtain a future lease.

FILES /etc/default/inet_type DEFAULT_IP setting

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**

The -R option requires the net_rawaccess privilege.

TCP and UDP socket information has an additional sensitivity label column with the -v flag. If the local port used for a socket is a single-level port (SLP), the label of the port is printed. The string "Multi-level" is printed for a multi-level port (MLP) and the string "No Label" is printed if the socket is not bound to any port.

**Trusted Solaris 8
4/01 Reference
Manual**

arp(1M), ifconfig(1M)

crash(1M), iostat(1M), dhcpagent(1M), dhcpinfo(1), vmstat(1M), mibiisa(1M), savecore(1M), hosts(4), inet_type(4), networks(4), protocols(4), services(4), attributes(5), inet(7P), inet6(7P)

Droms, R., *RFC 2131, Dynamic Host Configuration Protocol*, Network Working Group, March 1997.

NOTES

When printing interface information, netstat honors the DEFAULT_IP setting in /etc/default/inet_type. If it is set to IP_VERSION4, then netstat will omit information relating to IPv6 interfaces, statistics, connections, routes and the like.

However you can override the DEFAULT_IP setting in /etc/default/inet_type on the command-line. For example, if you have used the command-line to explicitly request IPv6 information by using the inet6 address family or one of the IPv6 protocols, it will override the DEFAULT_IP setting.

If you need to examine network status information following a kernel crash, use the crash(1M) utility on the savecore(1M) output.

NAME	setfsattr, newsecfs – Set security attributes on an existing or newly created file system	
SYNOPSIS	<pre> /usr/sbin/setfsattr { [-l <i>sensitivity-level-range</i>] [-m <i>MLD-prefix</i>] [-p <i>allowed-privilege-set</i>] [-P <i>forced-privilege-set</i>] [-s <i>CMW-Label</i>] ... } { <i>special</i> <i>filesystem</i> } /usr/sbin/newsecfs { [-l <i>sensitivity-level-range</i>] [-M] [-m <i>MLD-prefix</i>] [-o <i>newfs options</i>] [-p <i>allowed-privilege-set</i>] [-P <i>forced-privilege-set</i>] [-s <i>CMW-Label</i>] ... } { <i>special</i> <i>filesystem</i> } </pre>	
DESCRIPTION	<p>setfsattr changes the security attributes of a file system. The file system may be specified either as a <i>filesystem</i> or as <i>special</i>, the device on which the file system resides. <i>filesystem</i> must be in <i>/etc/vfstab</i>, and it must be unmounted before setfsattr is invoked on it. setfsattr requires at least one option be specified; if not, an error is returned.</p> <p>newsecfs works similarly to setfsattr except that it runs newfs(1M) on the file system prior to setting the security attributes, then sets the label on the <i>lost+found</i> directory to [ADMIN_HIGH].</p>	
OPTIONS	<ul style="list-style-type: none"> -l <i>sensitivity-level-range</i> Set the filesystem sensitivity level range, a semicolon-separated pair of sensitivity labels. The labels must be valid sensitivity labels for the system. The first in the pair is the minimum sensitivity label, and it must be dominated by the second label, the maximum sensitivity label. The default is ADMIN_LOW;ADMIN_HIGH. -M Create the root directory of the file system as a multilevel directory (MLD). This option is available only with the newsecfs command. -m <i>MLD-prefix</i> Set the file system MLD prefix. The default is ".MLD.". The MLD prefix is the string that disables multilevel directory translation in pathname lookup. -o <i>newfs options</i> Set the file system newfs options. The options must be exactly the same as those expected by the newfs(1M) command. This option is available only with newsecfs. -p <i>allowed-privileges</i> Set the file system allowed-privilege set, specified as a text-string of comma-separated privilege names. The privileges in the allowed set must include all privileges in the forced set, or the operation fails. -P <i>forced-privileges</i> Set the filesystem forced-privilege set, specified as a text string of comma-separated privilege names. All privileges in the forced set must also be in the allowed set, or the operation fails. -s <i>CMW-Label</i> Set the filesystem CMW label. 	

newsecfs(1M)

USAGE To specify arguments that include semicolons or embedded spaces (such as for the `-l` and `-o` options), use quotes to enclose the arguments.

EXAMPLES **EXAMPLE 1** To create a new file system with a limited label range

To create a new file system with an allowable label range of Confidential to Secret, use this command:

```
$ newsecfs -l 'confidential;secret' raw_device
```

EXIT STATUS setfsattr exits with one of these values:

0 Success.

1 Failure.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWtsu

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

`fork(2)`

`mkfs(1M)`, `newfs(1M)`, `terminfo(4)`, `attributes(5)`

NAME	nfsd – NFS daemon
SYNOPSIS	/usr/lib/nfs/nfsd [-a] [-c #_conn] [-l <i>listen_backlog</i>] [-p <i>protocol</i>] [-t <i>device</i>] [<i>nservers</i>]
DESCRIPTION	<p>nfsd is the daemon that handles client filesystem requests. Users must have the <code>sys_nfs</code> privilege to run this daemon.</p> <p>The nfsd daemon is automatically invoked in run level 3 with the -a option.</p> <p>By default nfsd will start over the TCP and UDP transports.</p> <p>A previously invoked nfsd daemon started with or without options must be stopped before invoking another nfsd command.</p>
OPTIONS	<p>The following options are supported:</p> <ul style="list-style-type: none"> -a Start a NFS daemon over all available connectionless and connection-oriented transports, including TCP and UDP. -c #_conn This sets the maximum number of connections allowed to the NFS server over connection-oriented transports. By default, the number of connections is unlimited. -l Set connection queue length for the NFS TCP over a connection-oriented transport. The default value is 32 entries. -p <i>protocol</i> Start a NFS daemon over the specified protocol. -t <i>device</i> Start a NFS daemon for the transport specified by the given device.
OPERANDS	<p>The following operands are supported:</p> <ul style="list-style-type: none"> <i>nservers</i> Set the maximum number of concurrent NFS requests that the server can handle. This concurrency is achieved by up to <i>nservers</i> threads created as needed in the kernel. <i>nservers</i> should be based on the load expected on this server. 16 is the usual number of <i>nservers</i>. If <i>nservers</i> is not specified, the maximum number of concurrent NFS requests will default to 1.
USAGE	<p>If the <code>NFS_PORTMON</code> variable is set, then clients are required to use privileged ports (ports < <code>IPPORT_RESERVED</code>) in order to get NFS services. In the Trusted Solaris environment, this variable is set to 1 by default. This variable has been moved from the "nfs" module to the "nfssrv" module. To set the variable, edit the <code>/etc/system</code> file and add this entry:</p> <pre>set nfssrv:nfs_portmon = 1</pre>
EXIT STATUS	<ul style="list-style-type: none"> 0 Daemon started successfully. 1 Daemon failed to start.

nfsd(1M)

FILES

.nfsXXX	Client machine pointer to an open-but-unlinked file
/etc/init.d/nfs.server	Shell script for starting nfsd
/etc/system	System configuration information file

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

Trusted Solaris 8
4/01 Reference
Manual

The sys_nfs and net_mac_read privileges are required to run this daemon. NFS_PORTMON has been set to 1 by default.

mountd(1M), sharetab(4)
ps(1), system(4), attributes(5)

NFS Administration Guide

NOTES

1. The NFS service uses kernel threads to process all of the NFS requests. Currently, system utilization associated with these threads is not charged to the nfsd process. Therefore, ps(1) can report 0 cpu time associated with the NFS daemon, even though NFS processing is taking place on the server.
2. Manually starting and restarting nfsd is not recommended. If it is necessary to do so, use the NFS server start/stop script (/etc/init.d/nfs.server). See *NFS Administration Guide* for more information.

NAME	nfsstat – NFS statistics
SYNOPSIS	nfsstat [-cnrsmza]
DESCRIPTION	<p>nfsstat displays statistical information about the NFS and RPC (Remote Procedure Call), interfaces to the kernel. It can also be used to reinitialize this information. If no options are given the default is</p> <pre>nfsstat -csnra</pre> <pre>nfsstat -cnrs</pre> <p>That is, display everything, but reinitialize nothing.</p> <p>To succeed with no option or with any option other than <code>-z</code>, <code>nfsstat</code> requires MAC and DAC read access to <code>/dev/mem</code>. To succeed with the <code>-z</code> option, <code>nfsstat</code> requires MAC and DAC write access to <code>/dev/mem</code> and the <code>sys_config</code> privilege.</p>
OPTIONS	<p><code>-a</code> Display NFS_ACL information.</p> <p><code>-c</code> Display client information. Only the client side NFS, RPC, and NFS_ACL information is printed. Can be combined with the <code>-n</code>, <code>-r</code>, and <code>-a</code> options to print client side NFS, RPC, and NFS_ACL information only.</p> <p><code>-m</code> Display statistics for each NFS mounted file system. This includes the server name and address, mount flags, current read and write sizes, the retransmission count, the attribute cache timeout values, failover information, and the timers used for dynamic retransmission. Note that the dynamic retransmission timers are displayed only where dynamic retransmission is in use. By default, NFS mounts over the TCP protocols and NFS Version 3 mounts over either TCP or UDP do not use dynamic retransmission. If you specify the <code>-m</code> option, this is the only option <code>nfsstat</code> uses. Any options specified in addition to <code>-m</code> are checked for validity, then ignored.</p> <p><code>-n</code> Display NFS information. NFS information for both the client and server side will be printed. Can be combined with the <code>-c</code> and <code>-s</code> options to print client or server NFS information only.</p> <p><code>-r</code> Display RPC information.</p> <p><code>-s</code> Display server information.</p> <p><code>-z</code> Zero (reinitialize) statistics. This option requires the <code>sys_config</code> privilege and can be combined with any of the above options to zero particular sets of statistics after printing them.</p>
DISPLAYS	<p>The server RPC display includes the following fields:</p> <pre>calls The total number of RPC calls received.</pre>

nfsstat(1M)

badcalls	The total number of calls rejected by the RPC layer (the sum of badlen and xdr call as defined below).
nullrecv	The number of times an RPC call was not available when it was thought to be received.
badlen	The number of RPC calls with a length shorter than a minimum-sized RPC call.
xdr call	The number of RPC calls whose header could not be XDR decoded.
dupchecks	The number of RPC calls that looked up in the duplicate request cache.
dupreqs	The number of RPC calls that were found to be duplicates.

The server NFS display shows the number of NFS calls received (`calls`) and rejected (`badcalls`), and the counts and percentages for the various calls that were made.

The server NFS_ACL display shows the counts and percentages for the various calls that were made.

The client RPC display includes the following fields:

calls	The total number of RPC calls made.
badcalls	The total number of calls rejected by the RPC layer.
badxids	The number of times a reply from a server was received which did not correspond to any outstanding call.
timeouts	The number of times a call timed out while waiting for a reply from the server.
newcreds	The number of times authentication information had to be refreshed.
badverfs	The number of times the call failed due to a bad verifier in the response.
timers	The number of times the calculated time-out value was greater than or equal to the minimum specified time-out value for a call.
cantconn	The number of times the call failed due to a failure to make a connection to the server.
nomem	The number of times the call failed due to a failure to allocate memory.
interrupts	The number of times the call was interrupted by a signal before completing.

retrans	The number of times a call had to be retransmitted due to a timeout while waiting for a reply from the server. Applicable only to RPC over connection-less transports.
cantsend	The number of times a client was unable to send an RPC request over a connectionless transport when it tried to do so.

The client NFS display shows the number of calls sent and rejected, as well as the number of times a CLIENT handle was received (`clgets`), the number of times the CLIENT handle cache had no unused entries (`cltoomany`), as well as a count of the various calls and their respective percentages.

The client NFS_ACL display shows the counts and percentages for the various calls that were made.

The `-m` option includes information about mount flags set by mount options, mount flags internal to the system, and other mount information. See `mount_nfs(1M)`.

The following mount flags are set by mount options:

sec	sec has one of the following values:
none	No authentication.
sys	UNIX-style authentication (UID, GID).
short	Short hand UNIX style authentication.
dh	des—style authentication (encrypted timestamps).
krb4	kerberos v4—style authentication.
krb5	kerberos v5—style authentication.
krb5i	kerberos v5—style authentication with integrity.
hard	Hard mount.
soft	Soft mount.
intr	Interrupts allowed on hard mount.
nointr	No interrupts allowed on hard mount.
noac	Client is not caching attributes.
rsize	Read buffer size in bytes.
wsiz	Write buffer size in bytes.
retrans	NFS retransmissions.
timeo	Initial NFS timeout, in tenths of a second.
nocto	No close-to-open consistency.
llock	Local locking being used (no lock manager).

nfsstat(1M)

grp_{id} System V group id inheritance.

rpctimesync RPC time sync.

The following mount flags are internal to the system:

printed "Not responding" message printed.

down Server is down.

dynamic Dynamic transfer size adjustment.

link Server supports links.

symlink Server supports symbolic links.

readdir Use readdir instead of readdirplus.

acl Server supports NFS_ACL.

The following flags relate to additional mount information:

vers NFS version.

proto Protocol.

The -m option also provides attribute cache timeout values. The following fields in -m output provide timeout values for attribute cache:

acregmin Minimum seconds to hold cached file attributes.

acregmax Maximum seconds to hold cached file attributes.

acdirmin Minimum seconds to hold cached directory attributes.

acdirmax Maximum seconds to hold cached directory attributes.

The following fields in -m output provide failover information:

noresponse How many times servers have failed to respond.

failover How many times a new server has been selected.

remap How many times files have been re-evaluated to the new server.

currserver Which server is currently providing NFS service. See the *NFS Administration Guide* for additional details.

The fields in -m output shown below provide information on dynamic retransmissions. Note that these items are displayed only where dynamic retransmission is in use.

srtt The value for the smoothed round-trip time, in milliseconds.

dev Estimated deviation, in milliseconds.

cur Current backed-off retransmission value, in milliseconds.

EXIT STATUS The following exit values are returned:

0 Successful completion.
 >0 An error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES To succeed with no option or with any option other than `-z`, `nfsstat` requires MAC and DAC read access to `/dev/mem`. To succeed with the `-z` option, `nfsstat` requires MAC and DAC write access to `/dev/mem` and the `sys_config` privilege.

**Trusted Solaris 8
4/01 Reference
Manual**

`mount_nfs(1M)`

`attributes(5)`

Solaris 8 Advanced Installation Guide

NFS Administration Guide

nis_cachemgr(1M)

NAME	nis_cachemgr – NIS+ utility to cache location information about NIS+ servers
SYNOPSIS	<code>/usr/sbin/nis_cachemgr [-i] [-v]</code>
DESCRIPTION	<p>The <code>nis_cachemgr</code> daemon maintains a cache of NIS+ directory objects and active servers for domains. It is responsible for locating servers for a domain on behalf of client processes. This improves performance because only one process has to search for servers. The cache contains location information necessary to contact the NIS+ servers. This includes transport addresses, information needed to authenticate the server, and a time to live field which gives a hint on how long the directory object can be cached. The cache helps to improve the performance of the clients that are traversing the NIS+ name space. <code>nis_cachemgr</code> should be running on all the machines that are using NIS+. However, it is not required that the <code>nis_cachemgr</code> program be running in order for NIS+ requests to be serviced.</p> <p>The cache maintained by this program is shared by all the processes that access NIS+ on a machine. The cache is maintained in a file that is memory mapped (see <code>mmap(2)</code>) by all the processes. On start up, <code>nis_cachemgr</code> initializes the cache from the cold start file (see <code>nisinit(1M)</code>) and preserves unexpired entries that already exist in the cache file. Thus, the cache survives machine reboots.</p> <p>The <code>nis_cachemgr</code> program is normally started from a system startup script. <code>nisshowcache(1M)</code> can be used to look at the cached objects and active servers. It must be started by a user with a UID of 0 and at a sensitivity label of <code>ADMIN_LOW</code>. Upon startup <code>nis_cachemgr</code> must inherit the <code>net_mac_read</code>, <code>net_upgrade_sl</code>, <code>file_dac_read</code>, <code>file_dac_write</code>, and <code>sys_system_door</code> privileges.</p> <p>The <code>nisprefadm(1M)</code> command (see the SunOS 5.8 Reference Manual) can be used to control which NIS+ servers the <code>nis_cachemgr</code> program will try to select.</p> <p>The <code>nis_cachemgr</code> program makes NIS+ requests under the NIS+ principal name of the host on which it runs. Before running <code>nis_cachemgr</code>, security credentials for the host should be added to the <code>cred.org_dir</code> table in the host's domain using <code>nisaddcred(1M)</code>. Credentials of type DES will be needed if the NIS+ service is operating at security level 2 (see <code>rpc.nisd(1M)</code>). See the WARNINGS section, below. Additionally, a "keylogin -r" should be done on the machine.</p>
OPTIONS	<p>-i Force <code>nis_cachemgr</code> to ignore the previous cache file and reinitialize the cache from just the cold start file. By default, the cache manager initializes itself from both the cold start file and the old cache file, thereby maintaining the entries in the cache across machine reboots.</p> <p>-v This flag sets verbose mode. In this mode, the <code>nis_cachemgr</code> program logs not only errors and warnings, but also additional status messages. The additional messages are logged using <code>syslog(3C)</code> with a priority of <code>LOG_INFO</code>.</p>

nis_cachemgr(1M)

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**

The `nis_cachemgr` must be started by a user with a UID of 0 and at a sensitivity level of `ADMIN_LOW`. At startup it must inherit the `net_mac_read`, `net_upgrade_sl`, `file_dac_read`, `file_dac_write`, and `sys_system_door` privileges.

FILES

<code>/var/nis/NIS_SHARED_DIRCACHE</code>	The shared cache file
<code>/var/nis/NIS_COLD_START</code>	The coldstart file
<code>/etc/init.d/rpc</code>	Initialization scripts for NIS+

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

**Trusted Solaris 8
4/01 Reference
Manual
Solaris 8
Reference Manual**

`rpc.nisd(1M)`, `rpc(3NSL)`

`keylogin(1)`, `nisaddcred(1M)`, `nisinit(1M)`, `nisprefadm(1M)`,
`nisshowcache(1M)`, `mmap(2)`, `syslog(3C)`, `nisfiles(4)`, `attributes(5)`

DIAGNOSTICS

The `nis_cachemgr` daemon logs error messages and warnings using `syslog(3C)`. Error messages are logged to the `DAEMON` facility with a priority of `LOG_ERR`. Warning messages are logged with a priority of `LOG_WARNING`. Additional status messages can be obtained using the `-v` option.

nisclient(1M)

NAME	nisclient – initialize NIS+ credentials for NIS+ principals
SYNOPSIS	<pre> /usr/lib/nis/nisclient -c [-x] [-o] [-v] [-l <network_password>] [-d <NIS+_domain>] client_name... /usr/lib/nis/nisclient -i [-x] [-v] -h <NIS+_server_host> [-a <NIS+_server_addr>] [-k <key_domain>] [-d <NIS+_domain>] [-s 0 2] /usr/lib/nis/nisclient -u [-x] [-v] /usr/lib/nis/nisclient -r [-x] </pre>
DESCRIPTION	<p>The nisclient shell script can be used to:</p> <ul style="list-style-type: none"> ■ create NIS+ credentials for hosts and users ■ initialize NIS+ hosts and users ■ restore the network service environment <p>NIS+ credentials are used to provide authentication information of NIS+ clients to NIS+ service. Upon startup, nisclient must inherit the file_dac_read, file_dac_write, file_mac_read, sys_system_door, net_mac_read, net_reply_equal, net_setclr, net_setid, net_setpriv, net_upgrade_sl, proc_owner, and sys_netconfig privileges.</p> <p>Use the first synopsis (-c option) to create individual NIS+ credentials for hosts or users. You must be logged in as a NIS+ principal in the domain for which you are creating the new credentials. You must also have write permission to the local "cred" table. The <i>client_name</i> argument accepts any valid host or user name in the NIS+ domain (for example, the <i>client_name</i> must exist in the hosts or passwd table). nisclient verifies each <i>client_name</i> against both the host and passwd tables, then adds the proper NIS+ credentials for hosts or users. Note that if you are creating NIS+ credentials outside of your local domain, the host or user must exist in the host or passwd tables in both the local and remote domains.</p> <p>By default, nisclient will not overwrite existing entries in the credential table for the hosts and users specified. To overwrite, use the -o option. After the credentials have been created, nisclient will print the command that must be executed on the client machine to initialize the host or the user. The -c option requires a network password for the client which is used to encrypt the secret key for the client. You can either specify it on the command line with the -l option or the script will prompt you for it. You can change this network password later with passwd(1) or chkey(1).</p> <p>nisclient -c is not intended to be used to create NIS+ credentials for all users and hosts which are defined in the passwd and hosts tables. To define credentials for all users and hosts, use nispopulate(1M).</p> <p>Use the second synopsis (-i option) to initialize a NIS+ client machine. The -i option can be used to convert machines to use NIS+ or to change the machine's domainname. You must be logged in as root on the machine that is to become a NIS+ client. Your administrator must have already created the NIS+ credential for this host by using</p>

`nisclient -c` or `nispopulate -C`. You will need the network password your administrator created. `nisclient` will prompt you for the network password to decrypt your secret key and then for this machine's root login password to generate a new set of secret/public keys. If the NIS+ credential was created by your administrator using `nisclient -c`, then you can simply use the initialization command that was printed by the `nisclient` script to initialize this host instead of typing it manually.

To initialize an unauthenticated NIS+ client machine, use the `-i` option with `-s 0`. With these options, the `nisclient -i` option will not ask for any passwords.

During the client initialization process, files that are being modified are backed up as `files.no_nisplus`. The files that are usually modified during a client initialization are: `/etc/defaultdomain`, `/etc/nsswitch.conf`, `/etc/inet/hosts`, and, if it exists, `/var/nis/NIS_COLD_START`. Notice that a file will not be saved if a backup file already exists.

The `-i` option does not set up a NIS+ client to resolve hostnames using DNS. Please refer to the DNS documentation for information on setting up DNS. (See `resolv.conf(4)`).

It is not necessary to initialize either NIS+ root master servers or machines that were installed as NIS+ clients using `suninstall(1M)`.

Use the third synopsis (`-u` option) to initialize a NIS+ user. You must be logged in as the user on a NIS+ client machine in the domain where your NIS+ credentials have been created. Your administrator should have already created the NIS+ credential for your username using `nisclient -c` or `nispopulate(1M)`. You will need the network password your administrator used to create the NIS+ credential for your username. `nisclient` will prompt you for this network password to decrypt your secret key and then for your login password to generate a new set of secret/public keys.

Use the fourth synopsis (`-r` option) to restore the network service environment to whatever you were using before `nisclient -i` was executed. You must be logged in as root on the machine that is to be restored. The restore will only work if the machine was initialized with `nisclient -i` because it uses the backup files created by the `-i` option.

Reboot the machine after initializing a machine or restoring the network service.

OPTIONS

The following options are supported:

- | | |
|--|--|
| <code>-a <NIS+_server_addr></code> | Specifies the IP address for the NIS+ server. This option is used <i>only</i> with the <code>-i</code> option. |
| <code>-c</code> | Adds DES credentials for NIS+ principles. |
| <code>-d <NIS+_domain></code> | Specifies the NIS+ domain where the credential should be created when used in conjunction with the <code>-c</code> |

nisclient(1M)

	option. It specifies the name for the new NIS+ domain when used in conjunction with the <code>-i</code> option. The default is your current domainname.
<code>-h <NIS+_server_host></code>	Specifies the NIS+ server's hostname. This option is used <i>only</i> with the <code>-i</code> option.
<code>-i</code>	Initializes a NIS+ client machine.
<code>-l <network_password></code>	Specifies the network password for the clients. This option is used <i>only</i> with the <code>-c</code> option. If this option is not specified, the script will prompt you for the network password.
<code>-k <key_domain></code>	This option specifies the domain where root's credentials are stored. If a domain is not specified, then the system default domain is assumed.
<code>-o</code>	Overwrites existing credential entries. The default is not to overwrite. This is used <i>only</i> with the <code>-c</code> option.
<code>-r</code>	Restores the network service environment.
<code>-s 0 2</code>	Specifies the authentication level for the NIS+ client. Level 0 is for unauthenticated clients and level 2 is for authenticated (DES) clients. The default is to set up with level 2 authentication. This is used <i>only</i> with the <code>-i</code> option. <code>nisclient</code> always uses level 2 authentication (DES) for both <code>-c</code> and <code>-u</code> options. There is no need to run <code>nisclient</code> with <code>-u</code> and <code>-c</code> for level 0 authentication. To configure authentication mechanisms other than DES at security level 2, use <code>nisauthconf(1M)</code> before running <code>nisclient</code> .
<code>-u</code>	Initializes a NIS+ user.
<code>-v</code>	Runs the script in verbose mode.
<code>-x</code>	Turns the "echo" mode on. The script just prints the commands that it would have executed. Notice that the commands are not actually executed. The default is off.

EXAMPLES

EXAMPLE 1 Adding the DES credential in the local domain

To add the DES credential for host `sunws` and user `fred` in the local domain:

```
example% /usr/lib/nis/nisclient -c sunws fred
```

EXAMPLE 2 Adding the DES credential in a specified domain

To add the DES credential for host `sunws` and user `fred` in domain `xyz.sun.com.`:

```
example% /usr/lib/nis/nisclient -c -d xyz.sun.com. sunws fred
```


EXAMPLE 3 Initializing the host in a specific domain

To initialize host sunws as a NIS+ client in domain xyz.sun.com. where nisplus_server is a server for the domain xyz.sun.com.:

```
example# /usr/lib/nis/nisclient -i -h nisplus_server -d xyz.sun.com
```

The script will prompt you for the IP address of nisplus_server if the server is not found in the /etc/hosts file. The -d option is needed only if your current domain name is different from the new domain name.

EXAMPLE 4 Initializing the host as an unauthenticated client in a specific domain

To initialize host sunws as an unauthenticated NIS+ client in domain xyz.sun.com. where nisplus_server is a server for the domain xyz.sun.com:

```
example# /usr/lib/nis/nisclient -i -S 0 \
-h nisplus_server -d xyz.sun.com. -a 129.140.44.1
```

EXAMPLE 5 Initializing the user as a NIS+ principal

To initialize user fred as a NIS+ principal, log in as user fred on a NIS+ client machine.

```
example% /usr/lib/nis/nisclient -u
```

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**

At startup, nisclient must inherit the file_dac_read, file_dac_write, file_mac_read, sys_system_door, net_mac_read, net_reply_equal, net_setclr, net_setid, net_setpriv, net_upgrade_sl, proc_owner, and sys_netconfig privileges.

FILES

/var/nis/NIS_COLD_START	This file contains a list of servers, their transport addresses, and their Secure RPC public keys that serve the machines default domain.
/etc/defaultdomain	The system default domainname.
/etc/nsswitch.conf	Configuration file for the name-service switch.
/etc/inet/hosts	Local host name database.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWnisu

nisclient(1M)

Trusted Solaris 8
4/01 Reference
SunOS 5.8
Reference Manual

passwd(1), nispopulate(1M), nsswitch.conf(4), resolv.conf(4)
chkey(1), keylogin(1), nis+(1), keyserver(1M), nisaddcred(1M),
nisauthconf(1M), nisinit(1M), suninstall(1M), attributes(5)

NAME	rpc.nisd, nisd – NIS+ service daemon
SYNOPSIS	/usr/sbin/rpc.nisd [-ACDFhlu ^v] [-Y [-B [-t <i>netid</i>]]] [-d <i>dictionary</i>] [-L <i>load</i>] [-S <i>level</i>]
DESCRIPTION	<p>The <code>rpc.nisd</code> daemon is an RPC service that implements the NIS+ service. This daemon must be running on all machines which serve a portion of the NIS+ namespace. A Trusted Solaris system must be the root master in the NIS+ configuration.</p> <p><code>rpc.nisd</code> is usually started from a system startup script. It must be started through a role that has a UID of 0 and run with a sensitivity label of <code>ADMIN_LOW</code>. (For example, the role might be assigned the predefined NIS+ security administration and NIS+ administration profiles.) <code>rpc.nisd</code> must be run from the Trusted Path and inherit the <code>net_mac_read</code>, <code>net_upgrade_sl</code>, and <code>proc_setsl</code> privileges.</p> <p>The <code>-B</code> option causes <code>rpc.nisd</code> to start an auxiliary process, <code>rpc.nisd_resolv</code>, which provides <code>ypserv</code> compatible DNS forwarding for NIS host requests. <code>rpc.nisd_resolv</code> can also be started independently. See <code>rpc.nisd_resolv(1M)</code> for more information on using <code>rpc.nisd_resolv</code> independently.</p>
OPTIONS	<p><code>-A</code> Authentication verbose mode. The daemon logs all the authentication related activities to <code>syslogd(1M)</code> with <code>LOG_INFO</code> priority.</p> <p><code>-C</code> Open diagnostic channel on <code>/dev/console</code>.</p> <p><code>-D</code> Debug mode (don't fork).</p> <p><code>-F</code> Force the server to do a checkpoint of the database when it starts up. Forced checkpoints may be required when the server is low on disk space. This option removes updates from the transaction log that have propagated to all of the replicas.</p> <p><code>-h</code> Print list of options.</p> <p><code>-u</code> Allow updates from non-Trusted Solaris TCB clients.</p> <p><code>-v</code> Verbose. With this option, the daemon sends a running narration of what it is doing to the <code>syslog</code> daemon (see <code>syslogd(1M)</code>) at <code>LOG_INFO</code> priority. This option is most useful for debugging problems with the service (see also <code>-A</code> option).</p> <p><code>-Y</code> Put the server into NIS (YP) compatibility mode. When operating in this mode, the NIS+ server will respond to NIS Version 2 requests using the version 2 protocol. Because the YP protocol is not authenticated, only those items that have read access to nobody (the unauthenticated request) will be visible through the V2 protocol. It supports only the standard Version 2 maps in this mode (see <code>-B</code> option and <code>NOTES</code> in <code>ypfiles(4)</code>).</p> <p><code>-B</code> Provide <code>ypserv</code> compatible DNS forwarding for NIS host requests. The DNS resolving process, <code>rpc.nisd_resolv</code>, is started and controlled by <code>rpc.nisd</code>. This option requires that the <code>/etc/resolv.conf</code> file be setup</p>

nisd(1M)

for communication with a DNS nameserver. The `nslookup` utility can be used to verify communication with a DNS nameserver. See `resolv.conf(4)` and `nslookup(1M)`.

<code>-t netid</code>	Use <i>netid</i> as the transport for communication between <code>rpc.nisd</code> and <code>rpc.nisd_resolv</code> . The default transport is <code>ticots(7D)</code> (<code>tcp</code> on SunOS 4.x systems).
<code>-d dictionary</code>	Specify an alternate dictionary for the NIS+ database. The primary use of this option is for testing. Note that the string is not interpreted, rather it is simply passed to the <code>db_initialize()</code> function.
<code>-L number</code>	Specify the “load” the NIS+ service is allowed to place on the server. The load is specified in terms of the <i>number</i> of child processes that the server may spawn. This <i>number must</i> be at least 1 for the callback functions to work correctly. The default is 128.
<code>-S level</code>	Set the authorization security level of the service. The argument is a number between 0 and 2. By default, the daemon runs at security level 2.
0	Security level 0 is designed to be used for testing and initial setup of the NIS+ namespace. When running at level 0, the daemon does not enforce any access controls. Any client is allowed to perform any operation, including updates and deletions.
1	At security level 1, the daemon accepts both <code>AUTH_SYS</code> and <code>AUTH_DES</code> credentials for authenticating clients and authorizing them to perform NIS+ operations. This is not a secure mode of operation since <code>AUTH_SYS</code> credentials are easily forged. It should not be used on networks in which any untrusted users may potentially have access.
2	At security level 2, the daemon only accepts authentication using the security mechanisms configured by <code>nisauthconf(1M)</code> . The default security mechanism is <code>AUTH_DES</code> . Security level 2 is the default if the <code>-S</code> option is not used.

EXAMPLES

EXAMPLE 1 Setting up the NIS+ service.

The following example sets up the NIS+ service.

```
example% rpc.nisd
```

ENVIRONMENT
VARIABLES

EXAMPLE 1 Setting up the NIS+ service. *(Continued)*

EXAMPLE 2 Setting Up NIS+ Service Emulating YP With DNS Forwarding

The following example sets up the NIS+ service, emulating YP with DNS forwarding.

example% **rpc.nisd -YB**

FILES

- NETPATH** The transports that the NIS+ service will use can be limited by setting this environment variable (see `netconfig(4)`).
- /var/nis/data/parent.object**
This file describes the namespace that is logically above the NIS+ namespace. The most common type of parent object is a DNS object. This object contains contact information for a server of that domain.
- /var/nis/data/root.object**
This file describes the root object of the NIS+ namespace. It is a standard XDR-encoded NIS+ directory object that can be modified by authorized clients using the `nis_modify(3NSL)` interface.
- /etc/init.d/rpc**
Initialization script for NIS+.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWnisu

SUMMARY OF
TRUSTED
SOLARIS
CHANGES

A Trusted Solaris system must be the root master of the NIS+ configuration. The Trusted Solaris environment adds the `-u`. `rpc.nisd` must be run from the Trusted Path and inherit the `net_mac_read`, `net_upgrade_sl`, and `proc_setsl`. The daemon must be started by a role with a UID of 0 and run with a sensitivity label of `ADMIN_LOW`.

Trusted Solaris 8
4/01 Reference
Manual

SunOS 5.8
Reference Manual

- `nis_cachemgr(1M)`, `nissetup(1M)`, `nslookup(1M)`, `rpc.nisd_resolv(1M)`, `rpc.nispasswd(1M)`, `nis_modify(3NSL)`, `resolv.conf(4)`
- `nisauthconf(1M)`, `nisinit(1M)`, `syslogd(1M)`, `netconfig(4)`, `nisfiles(4)`, `attributes(5)`, `ticots(7D)`

nisd_resolv(1M)

NAME	rpc.nisd_resolv, nisd_resolv – NIS+ service daemon				
SYNOPSIS	rpc.nisd_resolv [-v -V] [-F [-C <i>fd</i>]] [-t <i>xx</i>] [-p <i>yy</i>]				
DESCRIPTION	<p>rpc.nisd_resolv is an auxiliary process which provides DNS forwarding service for NIS hosts requests to both yperv and rpc.nisd that are running in the NIS compatibility mode. It is generally started by invoking rpc.nisd(1M) with the -B option or yperv(1M) with the -d option. Although it is not recommended, rpc.nisd_resolv can also be started independently with the following options.</p> <p>This command is not supported in the Trusted Solaris environment because yperv and other NIS(YP) compatibility is unsupported.</p>				
OPTIONS	<p>-F Run in foreground.</p> <p>-C <i>fd</i> Use <i>fd</i> for service xprt (from nisd).</p> <p>-v Verbose. Send output to the syslog daemon.</p> <p>-V Verbose. Send output to stdout.</p> <p>-t <i>xx</i> Use transport <i>xx</i>.</p> <p>-p <i>yy</i> Use transient program# <i>yy</i>.</p>				
SUMMARY OF TRUSTED SOLARIS ATTRIBUTES CHANGES	<p>This command is not supported in the Trusted Solaris environment.</p> <p>See attributes(5) for descriptions of the following attributes:</p> <table><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWnisu</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWnisu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWnisu				
Trusted Solaris 8 4/01 Reference Manual Notes	<p>rpc.nisd(1M)</p> <p>nslookup(1M), resolv.conf(4), attributes(5)</p> <p>This command requires that the /etc/resolv.conf file be setup for communication with a DNS nameserver. The nslookup utility can be used to verify communication with a DNS nameserver. See resolv.conf(4) and nslookup(1M).</p>				

NAME	rpc.nispasswdd, nispasswdd – NIS+ password update daemon	
SYNOPSIS	/usr/sbin/rpc.nispasswdd [-a <i>attempts</i>] [-c <i>minutes</i>] [-D] [-g] [-v]	
DESCRIPTION	<p>rpc.nispasswdd daemon is an ONC+ RPC service that services password update requests from nispasswd(1). It updates password entries in the NIS+ passwd table.</p> <p>rpc.nispasswdd is normally started from a system startup script after the NIS+ server (rpc.nisd(1M)) has been started. rpc.nispasswdd will determine whether it is running on a machine that is a master server for one or more NIS+ directories. If it discovers that the host is not a master server, then it will promptly exit. It will also determine if rpc.nisd(1M) is running in NIS(YF) compatibility mode (the -Y option) and will register as yppasswdd for NIS(YF) clients as well.</p> <p>ypserv and other NIS (YP) compatibility is not supported.</p> <p>rpc.nispasswdd will syslog all failed password update attempts, which will allow an administrator to determine whether someone was trying to "crack" the passwords.</p> <p>rpc.nispasswdd has to be run by a superuser.</p>	
OPTIONS	<ul style="list-style-type: none"> -a <i>attempts</i> Set the maximum number of attempts allowed to authenticate the caller within a password update request session. Failed attempts are syslogd(1M) and the request is cached by the daemon. After the maximum number of allowed attempts the daemon severs the connection to the client. The default value is set to 3. -c <i>minutes</i> Set the number of minutes a failed password update request should be cached by the daemon. This is the time during which if the daemon receives further password update requests for the same user and authentication of the caller fails, then the daemon will simply not respond. The default value is set to 30 minutes. -D Debug. Run in debugging mode. -g Generate DES credential. By default the DES credential is not generated for a user if who does not have one. By specifying this option, if a user does not have a credential, then one will be generated and stored in the NIS+ cred table. -v Verbose. With this option, the daemon sends a running narration of what it is doing to the syslog daemon. This option is useful for debugging problems. 	
EXIT STATUS	<ul style="list-style-type: none"> 0 success 1 an error has occurred. 	
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>rpc.nispasswdd must be run with a UID of 0 and with a sensitivity label of ADMIN_LOW. On startup, rpc.nispasswdd must inherit the net_mac_read and net_upgrade_sl privileges. ypserv and other NIS (YP) compatibility is not supported.</p>	

nispaswdd(1M)

FILES /etc/init.d/rpc initialization script for NIS+

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWnisu

Trusted Solaris 8 rpc.nisd(1M), nsswitch.conf(4)
4/01 Reference
SunOS 5.6 nispaswd(1), passwd(1), syslogd(1M), attributes(5)
Reference Manual

NAME	nispopulate – Populate the NIS+ tables in a NIS+ domain
SYNOPSIS	<pre> /usr/lib/nis/nispopulate -Y [-x] [-f] [-n] [-u] [-v] [-S 0 2] [-l <network_passwd>] [-d <NIS+_domain>] -h <NIS_server_host> [-a <NIS_server_addr>] -y <NIS_domain> [table] ... /usr/lib/nis/nispopulate -F [-x] [-f] [-u] [-v] [-S 0 2] [-d <NIS+_domain>] [-l <network_passwd>] [-p <directory_path>] [table] ... /usr/lib/nis/nispopulate -C [-x] [-f] [-v] [-d <NIS+_domain>] [-l <network_passwd>] [hosts passwd] </pre>
DESCRIPTION	<p>The nispopulate shell script can be used to populate NIS+ tables in a specified domain from their corresponding files or NIS maps. nispopulate assumes that the tables have been created either through nisserver(1M) or nissetup(1M).</p> <p>The table argument accepts standard names that are used in the administration of Solaris systems and non-standard <i>key-value</i> type tables. See nisaddent(1M) for more information on <i>key-value</i> type tables. If the table argument is not specified, nispopulate will automatically populate each of the standard tables. These standard (default) tables are: auto_master, auto_home, ethers, group, hosts, ipnodes, networks, passwd, protocols, services, rpc, netmasks, bootparams, netgroup, aliases and shadow. Note that the shadow table is only used when populating from files. The non-standard tables that nispopulate accepts are those of <i>key-value</i> type. These tables must first be created manually with the nistbladm(1) command.</p> <p>Use the first synopsis (-Y) to populate NIS+ tables from NIS maps. nispopulate uses ypxfr(1M) to transfer the NIS maps from the NIS servers to the /var/yp/<NIS_domain> directory on the local machine. Then, it uses these files as the input source. Note that <NIS_domain> is case sensitive. Make sure there is enough disk space for that directory.</p> <p>Use the second synopsis (-F) to populate NIS+ tables from local files. nispopulate will use those files that match the table name as input sources in the current working directory or in the specified directory.</p> <p>Note that when populating the hosts, ipnodes, and passwd tables, nispopulate will automatically create the NIS+ credentials for all users and hosts (ipnodes) which are defined in the hosts, ipnodes, and passwd tables, respectively. A network passwd is required to create these credentials. This network password is used to encrypt the secret key for the new users and hosts. This password can be specified using the -l option or it will use the default password, "nisplus". nispopulate will not overwrite any existing credential entries in the credential table. Use nisclient(1M) to overwrite the entries in the cred table. It creates both LOCAL and DES credentials for users, and only DES credentials for hosts. To disable automatic credential creation, specify the "-S 0" option.</p> <p>The third synopsis (-C) is used to populate NIS+ credential table with level 2 authentication (DES) from the hosts, ipnodes and passwd tables of the specified</p>

nispopulate(1M)

domain. The valid table arguments for this operation are `hosts`, `ipnodes` and `passwd`. If this argument is not specified then it will use `hosts`, `ipnodes` and `passwd` as the input source. If other authentication mechanisms are configured using `nisauthconf(1M)`, the NIS+ credential table will be loaded with credentials for those mechanisms.

If `nispopulate` was earlier used with "`-S 0`" option, then no credentials were added for the hosts or the users. If later the site decides to add credentials for all users and hosts, then this (`-C`) option can be used to add credentials.

OPTIONS

<code>-a <NIS_server_addr></code>	Specifies the IP address for the NIS server. This option is <i>only</i> used with the <code>-Y</code> option.
<code>-C</code>	Populate the NIS+ credential table from <code>hosts</code> , <code>ipnodes</code> , and <code>passwd</code> tables using DES authentication (security level 2). If other authentication mechanisms are configured using <code>nisauthconf(1M)</code> , the NIS+ credential table will be populated with credentials for those mechanisms.
<code>-d <NIS+_domain.></code>	Specifies the NIS+ domain. The default is the local domain.
<code>-F</code>	Populates NIS+ tables from files.
<code>-f</code>	Forces the script to populate the NIS+ tables without prompting for confirmation.
<code>-h <NIS_server_host></code>	Specifies the NIS server hostname from where the NIS maps are copied from. This is only used with the <code>-Y</code> option. This hostname must be present in the NIS+ <code>hosts</code> or <code>ipnodes</code> table, or in the <code>/etc/hosts</code> or <code>/etc/inet/ipnodes</code> file. If the hostname is not defined, the script will prompt you for its IP address, or you can use the <code>-a</code> option to specify the address manually.
<code>-l <network_passwd></code>	Specifies the network password for populating the NIS+ credential table. This is only used when you are populating the <code>hosts</code> , <code>ipnodes</code> , and <code>passwd</code> tables. The default <code>passwd</code> is "nisplus".
<code>-n</code>	Does not overwrite local NIS maps in <code>/var/yp/<NISdomain></code> directory if they already exist. The default is to overwrite the existing NIS maps in the local <code>/var/yp/<NISdomain></code> directory. This is <i>only</i> used with the <code>-Y</code> option.
<code>-p <directory_path></code>	Specifies the directory where the files are stored. This is <i>only</i> used with the <code>-F</code> option. The default is the current working directory.

nispopulate(1M)

-s 0 2	Specifies the authentication level for the NIS+ clients. Level 0 is for unauthenticated clients and no credentials will be created for users and hosts in the specified domain. Level 2 is for authenticated (DES) clients and DES credentials will be created for users and hosts in the specified domain. The default is to set up with level 2 authentication (DES). There is no need to run nispopulate with -C for level 0 authentication. Also, if other authentication mechanisms are configured with nisauthconf(1M), credentials for those mechanisms will also be populated for the NIS+ clients.
-u	Updates the NIS+ tables (ie., adds, deletes, modifies) from either files or NIS maps. This option should be used to bring an NIS+ table up to date when there are only a small number of changes. The default is to add to the NIS+ tables without deleting any existing entries. Also, see the -n option for updating NIS+ tables from existing maps in the /var/yp directory.
-v	Runs the script in verbose mode.
-x	Turns the "echo" mode on. The script just prints the commands that it would have executed. Note that the commands are not actually executed. The default is off.
-Y	Populate the NIS+ tables from NIS maps.
-y <NIS_domain>	Specifies the NIS domain to copy the NIS maps from. This is <i>only</i> used with the -Y option. The default domainname is the same as the local domainname.

EXAMPLES

EXAMPLE 1 Using nispopulate

To populate all the NIS+ standard tables in the domain *xyz.sun.com.* from NIS+ maps of the *yp.sun.COM* domain as input source where host *yp_host* is a YP server of *yp.sun.COM*:

```
nis_server# /usr/lib/nis/nispopulate -Y -y yp.sun.COM \
-h yp_host -d xyz.sun.com.
```

EXAMPLE 2 Updating all NIS+ standard tables

To update all of the NIS+ standard tables from the same NIS domain and hosts shown above:

```
nis_server# /usr/lib/nis/nispopulate -Y -u -y yp.sun.COM -h yp_host \
-d xyz.sun.com.
```

nispopulate(1M)

EXAMPLE 2 Updating all NIS+ standard tables *(Continued)*

EXAMPLE 3 Populating the hosts table

To populate the hosts table in domain *xyz.sun.com.* from the hosts file in the */var/nis/files* directory and using "somepasswd" as the network password for key encryption:

```
nis_server# /usr/lib/nis/nispopulate -F -p \
/var/nis/files -l somepasswd hosts
```

EXAMPLE 4 Populating the passwd table

To populate the passwd table in domain *xyz.sun.com.* from the passwd file in the */var/nis/files* directory without automatically creating the NIS+ credentials:

```
nis_server# /usr/lib/nis/nispopulate -F -p /var/nis/files \
-d xys.sun.com. -S 0 passwd
```

EXAMPLE 5 Populating the credential table

To populate the credential table in domain *xyz.sun.com.* for all users defined in the passwd table.

```
nis_server# /usr/lib/nis/nispopulate -C -d xys.sun.com. passwd
```

EXAMPLE 6 Creating and populating a non-standard key-value type NIS+ table

To create and populate a non-standard key-value type NIS+ table, "private", from the file */var/nis/files/private*: (nispopulate assumes that the private.org_dirkey-value type table has already been created).

```
nis_server# /usr/bin/nistbladm -D access=og=rmcd,nw=r \
-c private key=S,nogw= value=,nogw= private.org.dir
nis_server# /usr/lib/nis/nispopulate -F -p /var/nis/files private
```

ENVIRONMENT VARIABLES

nispopulate normally creates temporary files in the directory */tmp*. You may specify another directory by setting the environment variable *TMPDIR* to your chosen directory. If *TMPDIR* is not a valid directory, then nispopulate will use */tmp*).

FILES

<i>/etc/inet/hosts</i>	Local host name database
<i>/etc/inet/ipnodes</i>	Local database associating names of nodes with IP addresses
<i>/var/yp</i>	NIS(YP) domain directory
<i>/var/nis</i>	NIS+ domain directory
<i>/tmp</i>	Temporary directory.

ATTRIBUTES

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

nispopulate(1M)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWnisu

nissetup creates the following additional tables: tnrhdb, tnrhdp, and tntime.
nissetup(1M)
nis+(1), nistbladm(1), nisaddcred(1M), nisaddent(1M), nisauthconf(1M),
nisclient(1M), nisserver(1M), rpc.nisd(1M), ypxfr(1M), attributes(5)

nissetup(1M)

NAME	nissetup – Initialize a NIS+ domain				
SYNOPSIS	/usr/lib/nis/nissetup [-Y] [<i>domain</i>]				
DESCRIPTION	<p>nissetup is a shell script that sets up a NIS+ domain to service clients that wish to store system administration information in a domain named <i>domain</i>. This domain should already exist prior to executing this command (see nismkdir(1) and nisinit(1M)).</p> <p>A NIS+ domain consists of a NIS+ directory and its subdirectories: org_dir and groups_dir. org_dir stores system administration information and groups_dir stores information for group access control.</p> <p>nissetup creates the subdirectories org_dir and groups_dir in <i>domain</i>. Both subdirectories will be replicated on the same servers as the parent domain. After the subdirectories are created, nissetup creates the default tables that NIS+ serves. These are auto_master, auto_home, bootparams, cred, ethers, group, hosts, mail_aliases, netmasks, networks, passwd, protocols, rpc, services, tnrhdb, tnrhttp, tntime, and timezone. The nissetup script uses the nistbladm(1) command to create these tables. The script can be easily customized to add site specific tables that should be created at setup time.</p> <p>This command is normally executed just once per domain.</p>				
OPTIONS	<p>-Y Specify that the domain will be served as both a NIS+ domain as well as an NIS domain using the backward compatibility flag. This will set up the domain to be less secure by making all the system tables readable by unauthenticated clients as well.</p>				
SUMMARY OF TRUSTED SOLARIS CHANGES ATTRIBUTES	<p>In Trusted Solaris 2.5, 2.5.1, 7 and 8, nissetup creates the following additional tables: tnrhdb, tnrhttp, and tntime.</p> <p>See attributes(5) for descriptions of the following attributes:</p> <table><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWnisu</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWnisu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWnisu				
SunOS 5.8 Reference Manual	nissetup(1), nismkdir(1), nistbladm(1), nisaddent(1M), nisinit(1M), nisserver(1M), attributes(5)				
NOTES	<p>While this command creates the default tables, it does not initialize them with data. This is accomplished with the nisaddent(1M) command.</p> <p>It is easier to use the nisserver(1M) script to create subdirectories and the default tables.</p>				

NAME	nscd – Name service cache daemon								
SYNOPSIS	/usr/sbin/nscd [-f <i>configuration-file</i>] [-g] [-e <i>cachename</i> , yes no] [-i <i>cachename</i>]								
DESCRIPTION	<p>nscd is a process that provides a cache for the most common name service requests. It starts up during multi-user boot. The default <i>configuration-file</i> <code>/etc/nscd.conf</code> determines the behavior of the cache daemon. See <code>nscd.conf(4)</code>.</p> <p>nscd provides caching for the <code>passwd(4)</code>, <code>group(4)</code>, <code>hosts(4)</code>, <code>ipnodes(4)</code>, <code>exec_attr(4)</code>, <code>prof_attr(4)</code>, and <code>user_attr(4)</code> databases through standard <code>libc</code> interfaces, such as <code>gethostbyname(3NSL)</code>, <code>getipnodebyname(3SOCKET)</code>, <code>gethostbyaddr(3NSL)</code>, and others. Each cache has a separate time-to-live for its data; modifying the local database (<code>/etc/hosts</code>, and so forth) causes that cache to become invalidated within ten seconds. The shadow file is specifically not cached. <code>getspnam(3C)</code> calls remain uncached as a result.</p> <p>nscd also acts as its own administration tool. If an instance of <code>nscd</code> is already running, commands are passed to the running version transparently.</p> <p>In order to preserve NIS+ security, the startup script for <code>nscd</code> (<code>/etc/init.d/nscd</code>) checks the permissions on the <code>passwd</code>, <code>group</code> and <code>host</code> tables if NIS+ is being used. If those tables are not readable by unauthenticated users, then caching is disabled so that each process continues to authenticate itself as before.</p> <p>nscd runs at the sensitivity label <code>ADMIN_LOW</code>. However, it can communicate with DNS name servers at any sensitivity label. It requires the Trusted Path attribute.</p>								
OPTIONS	<p>Several of the options described below require a <i>cachename</i> specification. Supported values are <code>passwd</code>, <code>group</code>, and <code>hosts</code>.</p> <table> <tr> <td>-f <i>configuration-file</i></td><td>Causes <code>nscd</code> to read its configuration data from the specified file.</td></tr> <tr> <td>-g</td><td>Prints current configuration and statistics to standard output. This is the only option executable by non-root users.</td></tr> <tr> <td>-e <i>cachename</i>, yes no</td><td>Enables or disables the specified cache.</td></tr> <tr> <td>-i <i>cachename</i></td><td>Invalidate the specified cache.</td></tr> </table>	-f <i>configuration-file</i>	Causes <code>nscd</code> to read its configuration data from the specified file.	-g	Prints current configuration and statistics to standard output. This is the only option executable by non-root users.	-e <i>cachename</i> , yes no	Enables or disables the specified cache.	-i <i>cachename</i>	Invalidate the specified cache.
-f <i>configuration-file</i>	Causes <code>nscd</code> to read its configuration data from the specified file.								
-g	Prints current configuration and statistics to standard output. This is the only option executable by non-root users.								
-e <i>cachename</i> , yes no	Enables or disables the specified cache.								
-i <i>cachename</i>	Invalidate the specified cache.								
EXAMPLES	<p>EXAMPLE 1 Stopping and restarting the <code>nscd</code> daemon.</p> <pre>example# /etc/init.d/nscd stop example# /etc/init.d/nscd start</pre>								
SUMMARY OF TRUSTED SOLARIS CHANGES	To invoke <code>nscd</code> requires the Trusted Path attribute, a process sensitivity label of <code>ADMIN_LOW</code> , and the following privileges: <code>net_upgrade_sl</code> , <code>net_mac_read</code> , <code>proc_setclr</code> , <code>sys_trans_label</code> , <code>sys_net_config</code> , <code>file_dac_write</code> , and <code>file_setid</code> . If <code>nscd</code> 's clearance is not <code>ADMIN_HIGH</code> , it will be set to <code>ADMIN_HIGH</code> .								

nscd(1M)

FILES /etc/nscd.conf determines behavior of cache daemon.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

**Trusted Solaris 8
4/01 Reference
Manual** exec_attr(4), nsswitch.conf(4), prof_attr(4), user_attr(4)
**Solaris 8
4/01 Reference
Manual** getsppnam(3C), gethostbyname(3NSL), getipnodebyname(3SOCKET), group(4),
hosts(4), ipnodes(4), nscd.conf(4), passwd(4), attributes(5)

WARNINGS The nscd interface is included in this release on an uncommitted basis only, and is subject to change or removal in a future minor release.

NAME	nslookup – query name servers interactively	
SYNOPSIS	nslookup [- option]... <i>host</i> [<i>server</i>] nslookup [- option]... - [<i>server</i>] nslookup	
DESCRIPTION	<p>nslookup sends queries to Internet domain name servers. It has two modes: interactive and non-interactive. Interactive mode allows the user to contact servers for information about various hosts and domains or to display a list of hosts in a domain. Non-interactive mode is used to display just the name and requested information for a host or domain.</p> <p>If the name server with which nslookup must communicate is on a non-trusted host, nslookup can communicate with that host if the host’s default sensitivity label matches the nslookup process’ sensitivity label. To communicate with a name server on a non-trusted host whose default sensitivity label does not match, nslookup must be run with the net_upgrade_sl, net_downgrade_sl, and net_mac_read privileges.</p>	
OPTIONS	-option	Set the permissible options, as shown in the following list. These are the same options that the set command supports in interactive mode (see set in the Commands section for more complete descriptions).
	all	List the current settings
	class=classname	Restrict search according to the specified class
	d2	Set exhaustive debug mode on
	nod2	Set exhaustive debug mode off
	debug	Set debug mode on
	nodebug	Set debug mode off
	defname	Set domain-appending mode on
	nodefname	Set domain-appending mode off
	domain=string	Establish the appendable domain
	ignoretc	Set it to ignore packet truncation errors
	noignoretc	Set it to acknowledge packet truncation errors
OPERANDS	host	Inquires about the specified <i>host</i> . In this non-interactive command format, nslookup does not prompt for additional commands.
	–	Causes nslookup to prompt for more information, such as host names, before sending one or more queries.

nslookup(1M)

	<p><i>server</i> Directs inquiries to the name server specified here in the command line rather than the one read from the <code>/etc/resolv.conf</code> file (see <code>resolv.conf(4)</code>). <i>server</i> can be either a name or an Internet address. If the specified host cannot be reached, <code>nslookup</code> resorts to using the name server specified in <code>/etc/resolv.conf</code>.</p>
Non-interactive Mode	<p>Non-interactive mode is selected when the name or Internet address of the host to be looked up is given as the first argument.</p> <p>Within non-interactive mode, space-separated options can be specified. They must be entered before the host name, to be queried. Each option must be prefixed with a hyphen.</p> <p>For example, to request extensive host information and to set the timeout to 10 seconds when inquiring about <code>gypsy</code>, enter:</p> <pre>example% nslookup-query=hinfo-timeout=10gypsy</pre> <p>To avoid repeated entry of an option that you almost always use, place a corresponding <code>set</code> command in a <code>.nslookuprc</code> file located inside your home directory. (See <code>Commands</code> for more information about <code>set</code>.) The <code>.nslookuprc</code> file can contain several <code>set</code> commands if each is followed by a RETURN.</p>
Entering and Leaving Interactive Mode	<p>Interactive mode is selected when</p> <ul style="list-style-type: none">■ No arguments are supplied.■ A '-' (hyphen) character is supplied as the <i>host</i> argument. <p>To exit from an interactive <code>nslookup</code> session, type Control-d or type the command <code>exit</code> followed by RETURN.</p>
Supported Command Interactions	<p>The commands associated with interactive mode are subject to various limitations and run-time conventions.</p> <p>The maximum length of a command line is 255 characters. When the RETURN key is pressed, command-line execution begins. While a command is running, its execution can be interrupted by typing Control-c.</p> <p>The first word entered on the command line must be the name of a <code>nslookup</code> command unless you wish to enter the name of a host to inquire about. Any unrecognized command is handled as a host name to inquire about. To force a command to be treated as a host name to be inquired about, precede it with a backslash character.</p>
Commands	<p><code>exit</code> Exit the <code>nslookup</code> program.</p> <p><code>help</code> Display a brief summary of commands.</p>

?

Display a brief summary of commands.

`host [server]`

Look up information for *host* using the current default server, or using *server* if it is specified.

If the *host* supplied is an Internet address and the query type is A or 1PTR, the name of the host is returned. If the *host* supplied is a name and it does not have a trailing period, the default domain name is appended to the name. (This behavior depends on the state of the `set` options `-domain`, `-srchlist`, `-defname`, and `-search`).

To look up a host that is not in the current domain, append a period to the name.

`finger [name] [>> filename]`

Connect with the finger server on the current host, which is defined by the most recent successful host lookup.

If no *name* value is specified, a list of login account names on the current host is generated.

Similar to a shell command interpreter, output can be redirected to a file using the usual redirection symbols: `>` and `>>`.

`ls [-options] domain [>> filename]`

List the information available for *domain*, optionally creating or appending to *filename*. The default output contains host names and their Internet addresses.

Output can be redirected to *filename* using the `>` and `>>` redirection symbols. When output is directed to a file, hash marks are shown for every 50 records received from the server. The permissible values for *options* are:

<code>a</code>	Lists aliases of hosts in the domain. This is a synonym for the command <code>ls -tCNAME</code> .
<code>d</code>	Lists all records for the domain. This is a synonym for the command <code>ls -tANY</code> .
<code>h</code>	Lists CPU and operating system information for the domain. This is a synonym for the command <code>ls -tHINFO</code> .
<code>s</code>	Lists well-known services of hosts in the domain. This is a synonym for the command <code>ls -tWKS</code> .
<code>t <i>querytype-value</i></code>	lists all records of the specified type (see <i>querytype</i> within the discussion of the <code>set</code> command).

`set token=value`

`set keyword`

nslookup(1M)

Establish a preferred mode of search operation. Permissible *token* and *keyword* values are:

<code>all</code>	Display the current values of frequently-used options. Information about the current default server and host is also displayed.
<code>cl[ass]=classname</code>	Limit the search according to the protocol group (<i>classname</i>) for which lookup information is desired. Permissible <i>classname</i> values are: ANY A wildcard selecting all classes IN The Internet class (the default) CHAOS The Chaos class. HESIOD The MIT Athena Hesiod class.
<code>d2</code> <code>nod2</code>	Enable or disable exhaustive debugging mode. Essentially all fields of every packet are displayed. By default, this option is disabled.
<code>deb[ug]</code> <code>nodeb[ug]</code>	Enable or disable debugging mode. When debugging mode is enabled, much more information is produced about the packet sent to the server and the resulting answer. By default, this option is disabled.
<code>def[name]</code> <code>nodef[name]</code>	Enable or disable appending the default domain name to a single-component lookup request (one that lacks a dot). By default, this option is enabled for <code>nslookup</code> . The default value for the domain name is the value given in <code>/etc/resolv.conf</code> , unless: there is an environmental value for <code>LOCALDOMAIN</code> when <code>nslookup</code> is run; a recent value has been specified through the <code>srchlist</code> command or the <code>set domain</code> command.
<code>do[main]=string</code>	Change the default domain name to be appended to all lookup requests to <i>string</i> . For this option to have any effect, the <code>-defname</code> option must also be enabled and the <code>-search</code> option must be set in a compatible way. The domain search list contains the parents of the default domain if it has at least two components in its name. For example, if the default domain is <code>CC.Berkeley.EDU</code> , the search list is <code>CC.Berkeley.EDU</code> and <code>Berkeley.EDU</code> . Use the

nslookup(1M)

set srchlist command to specify a different list.
Use the set all command to display the list.

ignoretc
noignoretc

Ignore packet truncation errors. By default, this option is disabled.

srch[list]=*name1/name2/...*

Change the default domain name to *name1* and the domain search list to *name1*, *name2*, etc. A maximum of 6 names can be specified, along with slash characters to separate them. For example,

```
example% set srchlist=lcs.MIT.EDU/ai.MIT.EDU/MIT.EDU
```

sets the domain to *lcs.MIT.EDU* and the search list to all three names. This command overrides the default domain name and search list of the set domain command. Use the set all command to display the list.

search
nosearch

Enable or disable having the domain names in the domain search list appended to the request, generating a series of lookup queries if necessary until an answer is received. To take effect, the lookup request must contain at least one dot (period); yet it must not contain a trailing period. By default, this option is enabled.

po[rt]=*value*

Specify the default TCP/UDP name server port. By default, this value is 53.

q[querytype]=*value*

ty[pe]=*value*

Change the type of information returned from a query to one of:

A	The Internet address of the host
CNAME	The canonical name for an alias
HINFO	The host CPU and operating system type
MD	The mail destination
MX	The mail exchanger
MB	The mailbox domain name
MG	The mail group member
MINFO	The mailbox or mail list information
NS	The name server
PTR	The host name if the query is in the form of an Internet address; otherwise the pointer to other information
SOA	The domain's start-of-authority information

nslookup(1M)

TXT The text information

UINFO The user information

WKS The supported well-known services (Other types specified in the *RFC 1035* document are valid, but they are not as useful.)

recurse
norecurse
 Enable or disable having to query other name servers before abandoning a search. By default, this feature is enabled.

ret[ry]=count
 Set the maximum number of times to retry a request before abandoning a search. When a reply to a request is not received within a certain amount of time (changed with *set timeout*), the timeout period is doubled and the request is resent. The retry value controls how many times a request is resent before the request is aborted. The default for *count* is 4.

ro[ot]=host
 Change the name of the root server to *host*. This affects the *root* command. The default root server is *ns.internic.net*.

t[timeout]=interval
 Change the amount of time to wait for a reply to *interval* seconds. Each retry doubles the timeout period. The default *interval* is 5 seconds.

vc
novc
 Enable or disable the use of a virtual circuit when sending requests to the server. By default, this feature is disabled.

root
 Change the default server to the server for the root of the domain name space. Currently, the host *ns.internic.net* is used; this command is a synonym for *server ns.internic.net*. The name of the root server can be changed with the *set root* command.

server domain
lserver domain
 Change the default server to *domain*. *lserver* uses the initial server to look up information about *domain* while *server* uses the current default server. If an authoritative answer can not be found, the names of servers that might have the answer are returned.

view filename
 Sort the output of previous *ls* command(s) and display it one text screenful at a time, similar to *more(1)*.

EXAMPLES

EXAMPLE 1 Searching the Internet domain namespace.

To effectively search the Internet domain namespace, it helps to know its structure. At present, the Internet domain name-space is tree-structured, with one top level domain for each country except the U.S.A. There are also some traditional top level domains, not explicitly tied to any particular country. These include:

COM	Commercial establishments
EDU	Educational institutions
ORG	Not-for-profit organizations
GOV	Government agencies
MIL	MILNET hosts

If you are looking for a specific host, you need to know something about the host's organization in order to determine the top-level domain that it belongs to. For instance, if you want to find the Internet address of a machine at UCLA, do the following:

- Connect with the root server using the `root` command. The root server of the name space has knowledge of the top-level domains.
- Since UCLA is a university, its domain name is `ucla.edu`. Connect with a server for the `ucla.edu` domain with the command `server ucla.edu`. The response produces the names of hosts that act as servers for that domain. Note: the root server does not have information about `ucla.edu`, but knows the names and addresses of hosts that do. Once located by the root server, all future queries will be sent to the UCLA name server.
- To request information about a particular host in the domain (for instance, `locus`), just type the host name. To request a listing of hosts in the UCLA domain, use the `ls` command. The `ls` command requires a domain name (in this case, `ucla.edu`) as an argument.

If you are connected with a name server that handles more than one domain, all lookups for host names must be fully specified with its domain. For instance, the domain `harvard.edu` is served by `seismo.css.gov`, which also services the `css.gov` and `cornell.edu` domains. A lookup request for the host `aiken` in the `harvard.edu` domain must be specified as `aiken.harvard.edu`. However, the `set domain=name` and `set defname` commands can be used to automatically append a domain name to each request.

After a successful lookup of a host, use the `finger(1)` command to see who is on the system, or to finger a specific person. (`finger` requires the type to be A.)

To get other information about the host, use the `set querytype=value` command to change the type of information desired and request another lookup.

ENVIRONMENT
VARIABLES

HOSTALIASES References the file containing host aliases

nslookup(1M)

LOCALDOMAIN Overrides default domain

EXIT STATUS The process returns the following values:

0 On success.

1 On failure.

FILES /etc/resolv.conf
initial domain name and name server addresses

\$HOME/.nslookuprc
initial option commands

/usr/lib/nslookup.help
summary of commands

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES If the name server with which nslookup must communicate is on a non-trusted host, nslookup can communicate with that host if the host's default sensitivity label matches the nslookup process' sensitivity label. To communicate with a name server on a non-trusted host whose default sensitivity label does not match, nslookup must be run with the net_upgrade_sl, net_downgrade_sl, and net_mac_read privileges.

Trusted Solaris 8 in.named(1M), resolver(3RESOLV), resolv.conf(4)

4/01 Reference Manual finger(1), more(1), nstest(1M), attributes(5)

SUNWcsu

Reference Manual

Mockapetris, Paul, *Domain Names - Concepts and Facilities*, RFC 1034, Network Information Center, SRI International, Menlo Park, Calif., November 1987.

Mockapetris, Paul, *Domain Names - Implementation and Specification*, RFC 1035, Network Information Center, SRI International, Menlo Park, Calif., November 1987.

DIAGNOSTICS If the lookup request is successful, an error message is produced. Possible errors are:

Timed out

The server did not respond to a request after a certain amount of time (changed with set timeout=*value*) and a certain number of retries (changed with set retry=*value*).

No response from server

No name server is running on the server machine.

No records

The server does not have resource records of the current query type for the host, although the host name is valid. The query type is specified with the `set querytype` command.

Non-existent domain

The host or domain name does not exist.

Connection refused**Network is unreachable**

The connection to the name or finger server can not be made at the current time. This error commonly occurs with `ls` and `finger` requests.

Server failure

The name server found an internal inconsistency in its database and could not return a valid answer.

Refused

The name server refused to service the request.

Format error

The name server found that the request packet was not in the proper format. This may indicate an error in `nslookup`.

pbind(1M)

NAME	pbind – Control and query bindings of processes to processors						
SYNOPSIS	pbind -b <i>processor_id</i> <i>pid</i> ... pbind -u <i>pid</i> ... pbind [-q] [<i>pid</i> ...]						
DESCRIPTION	<p>The pbind utility controls and queries bindings of processes to processors. pbind binds all the LWPs (lightweight processes) of a process to a processor, or removes or displays the bindings.</p> <p>When an LWP is bound to a processor, it will be executed only by that processor except when the LWP requires a resource that is provided only by another processor. The binding is not exclusive, that is, the processor is free execute other LWPs as well.</p> <p>Bindings are inherited, so new LWPs and processes created by a bound LWP will have the same binding. Binding an interactive shell to a processor, for example, binds all commands executed by the shell.</p> <p>This command may be used to bind or unbind any process for which the user has permission to signal — any process that has the same effective UID as the user. The <code>proc_owner</code> privilege is needed to bind or unbind any process with an effective user ID different from that of the user.</p>						
OPTIONS	<p>The following options are supported:</p> <table><tr><td>-b <i>processor_id</i></td><td>Binds all the LWPs of the specified processes to the processor <i>processor_id</i>. Specify <i>processor_id</i> as the processor ID of the processor to be controlled or queried. <i>processor_id</i> must be present and on-line. Use the <code>psrinfo</code> command to determine whether or not <i>processor_id</i> is present and online. See <code>psrinfo(1M)</code>.</td></tr><tr><td>-q</td><td>Displays the bindings of the specified processes, or of all processes. If a process is composed of multiple LWPs, which have different bindings, the bindings of only one of the bound LWPs will be displayed.</td></tr><tr><td>-u</td><td>Removes the bindings of all LWPs of the specified processes, allowing them to be executed on any on-line processor.</td></tr></table>	-b <i>processor_id</i>	Binds all the LWPs of the specified processes to the processor <i>processor_id</i> . Specify <i>processor_id</i> as the processor ID of the processor to be controlled or queried. <i>processor_id</i> must be present and on-line. Use the <code>psrinfo</code> command to determine whether or not <i>processor_id</i> is present and online. See <code>psrinfo(1M)</code> .	-q	Displays the bindings of the specified processes, or of all processes. If a process is composed of multiple LWPs, which have different bindings, the bindings of only one of the bound LWPs will be displayed.	-u	Removes the bindings of all LWPs of the specified processes, allowing them to be executed on any on-line processor.
-b <i>processor_id</i>	Binds all the LWPs of the specified processes to the processor <i>processor_id</i> . Specify <i>processor_id</i> as the processor ID of the processor to be controlled or queried. <i>processor_id</i> must be present and on-line. Use the <code>psrinfo</code> command to determine whether or not <i>processor_id</i> is present and online. See <code>psrinfo(1M)</code> .						
-q	Displays the bindings of the specified processes, or of all processes. If a process is composed of multiple LWPs, which have different bindings, the bindings of only one of the bound LWPs will be displayed.						
-u	Removes the bindings of all LWPs of the specified processes, allowing them to be executed on any on-line processor.						
OPERANDS	<p>The following operands are supported:</p> <table><tr><td><i>pid</i></td><td>The process ID of the process to be controlled or queried.</td></tr></table>	<i>pid</i>	The process ID of the process to be controlled or queried.				
<i>pid</i>	The process ID of the process to be controlled or queried.						
EXAMPLES	<p>EXAMPLE 1 Binding processes</p> <p>The following example binds processes 204 and 223 to processor 2.</p> <pre>example% pbind -b 2 204 223</pre> <p>This command displays the following output:</p>						

EXAMPLE 1 Binding processes (Continued)

```
process id 204: was 2, now 2
process id 223: was 3, now 2
```

EXAMPLE 2 Unbinding a process

The following example unbinds process 204.

```
example% pbind -u 204
```

EXAMPLE 3 Querying Bindings

The following example demonstrates that process 1 is bound to processor 0, process 149 has at least one LWP bound to CPU3, and process 101 has no bound LWPs.

```
example% pbind -q 1 149 101
```

This command displays the following output:

```
process id 1: 0
process id 149: 3
process id 101: not bound
```

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

EXIT STATUS

The following exit values are returned:

```
0          Successful completion.
>0         An error occurred.
```

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**
Trusted Solaris 8
4/01 Reference
Manual

The `proc_owner` privilege is needed to bind or unbind any process with an effective user ID different from that of the user.

`psradm(1M)`

`psrinfo(1M)`, `psrset(1M)`, `processor_bind(2)`, `processor_info(2)`,
`sysconf(3C)`, `attributes(5)`

DIAGNOSTICS

```
pbind: cannot query pid 31: No such process
       The process specified did not exist or has exited.
```

```
pbind: cannot bind pid 31: Not owner
       The user does not have permission to bind the process.
```

pbind(1M)

pbind: cannot bind pid 31: Invalid argument
The specified processor is not online.

NAME	pfsh – Profile shell
SYNOPSIS	pfsh [-acefhiknprstuvx] [<i>argument...</i>] (obsolete)
DESCRIPTION	The pfsh command is documented along with two other shell commands, pfcsh and pfksh, in Trusted Solaris 8 and later releases in the pfexec(1) man page.

pkgchk(1M)

NAME	pkgchk – Check package installation accuracy
SYNOPSIS	<p>pkgchk [-l -acfnqv] [-i <i>file</i>] [-p <i>path...</i>] [-R <i>root_path</i>] [[-m <i>pkgmap</i> [-e <i>envfile</i>]] [<i>pkginst</i>]...]</p> <p>pkgchk -d <i>device</i> [-l -fv] [-i <i>file</i>] [-M] [-p <i>path...</i>] [-v <i>fs_file</i>] [<i>pkginst</i>...]</p>
DESCRIPTION	<p>pkgchk checks the accuracy of installed files or, by using the -l option, displays information about package files. pkgchk checks the integrity of directory structures and files. Discrepancies are written to standard error along with a detailed explanation of the problem.</p> <p>The pkgchk utility handles the Trusted Solaris security attributes specified in the <code>tsolinfo(4)</code> file. pkgchk checks for the accuracy of the security attributes of installed files in addition to other attributes. The -l option displays security attributes of the package files in addition to other attributes. The -f option corrects security attributes of installed files in addition to other attributes.</p> <p>The first synopsis defined above is used to list or check the contents and/or attributes of objects that are currently installed on the system, or in the indicated pkgmap. Package names may be listed on the command line, or by default, the entire contents of a machine will be checked.</p> <p>The second synopsis is used to list or check the contents of a package which has been spooled on the specified device, but not installed. Note that attributes cannot be checked for spooled packages.</p>
OPTIONS	<p>The following options are supported:</p> <ul style="list-style-type: none"> -a Audit the file attributes only and do not check file contents. Default is to check both. -c Audit the file contents only and do not check file attributes. Default is to check both. -d <i>device</i> Specify the device on which a spooled package resides. <i>device</i> can be a directory path name or the identifiers for tape, floppy disk, or removable disk (for example, <code>/var/tmp</code> or <code>/dev/diskette</code>). -e <i>envfile</i> Request that the package information file named as <i>envfile</i> be used to resolve parameters noted in the specified pkgmap file. -f Correct file attributes if possible. If used with the -x option, this option removes hidden files. When pkgchk is invoked with this option, it creates directories, named pipes, links, and special devices if they do not already exist. If the -d option calls out an uninstalled package, the -f option will only take effect if the package is in directory (not stream) format. All file attributes will be set to agree with the entries in the pkgmap file except that <code>setuid</code>, <code>setgid</code>, and sticky bits will not be set in the mode.

	-i <i>file</i>	Read a list of path names from <i>file</i> and compare this list against the installation software database or the indicated pkgmap file. Path names which are not contained in <i>file</i> are not checked.
	-l	List information on the selected files that make up a package. This option is not compatible with the -a, -c, -f, -g, and -v options.
	-m pkgmap	Check the package against the package map file, pkgmap.
	-M	Instruct pkgchk not to use the <i>\$root_path/etc/vfstab</i> file for determining the client's mount points. This option assumes the mount points are correct on the server and it behaves consistently with Solaris 2.5 and earlier releases.
	-n	Do not check volatile or editable files' contents. This should be used for most post-installation checking.
	-p <i>path</i>	Only check the accuracy of the path name or path names listed. <i>path</i> can be one or more path names separated by commas (or by white space, if the list is quoted).
	-q	Quiet mode. Do not give messages about missing files.
	-R <i>root_path</i>	Define the full name of a directory to use as the <i>root_path</i> . All files, including package system information files, are relocated to a directory tree starting in the specified <i>root_path</i> . The <i>root_path</i> may be specified when installing to a client from a server (for example, /export/root/client1).
	-v	Verbose mode. Files are listed as processed.
	-V <i>fs_file</i>	Specify an alternative <i>fs_file</i> to map the client's file systems. For example, used in situations where the <i>\$root_path/etc/vfstab</i> file is non-existent or unreliable.
	-x	Search exclusive directories, looking for files which exist that are not in the installation software database or the indicated pkgmap file.
OPERANDS	<i>pkginst</i>	<p>The package instance or instances to be checked. The format <i>pkginst.*</i> can be used to check all instances of a package. The default is to display all information about all installed packages.</p> <p>The asterisk character (*) is a special character to some shells and may need to be escaped. In the C-Shell, "*" must be surrounded by single quotes (') or preceded by a backslash (\);</p>
EXAMPLES	<p>EXAMPLE 1 Using pkgchk for Displaying Package Installation Information</p> <p>The following example displays package installation information for /usr/bin/ls:</p> <pre>example% pkgchk -l -p /usr/bin/ls</pre>	

pkgchk(1M)

EXAMPLE 1 Using pkgchk for Displaying Package Installation Information (Continued)

EXIT STATUS

0 Successful completion.

>0 An error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

The pkgchk utility is extended to handle the Trusted Solaris security attributes specified in the tsolinfo(4) file. pkgchk checks for the accuracy of the security attributes of installed files in addition to other attributes. The -l option displays security attributes of the package files in addition to other attributes. The -f option corrects security attributes of installed files in addition to other attributes.

Trusted Solaris 8 4/01 Reference Manual

tsolinfo(4)

pkginfo(1), pkgtrans(1), pkgadd(1M), pkgask(1M), pkgrm(1M), attributes(5)

Application Packaging Developer's Guide

NAME	halt, poweroff – Stop the processor				
SYNOPSIS	<pre>/usr/sbin/halt [-dlnqy]</pre> <pre>/usr/sbin/poweroff [-dlnqy]</pre>				
DESCRIPTION	<p>halt and poweroff write out any pending information to the disks and then stop the processor. poweroff will have the machine remove power, if possible.</p> <p>halt and poweroff normally log the system shutdown to the system log daemon, syslogd(1M), and place a shutdown record in the login accounting file /var/adm/wtmp. These actions are inhibited if the -n or -q options are present.</p>				
OPTIONS	<p>The following options are supported:</p> <ul style="list-style-type: none"> -d Force a system crash dump before rebooting. See dumpadm(1M) for information on configuring system crash dumps. -l Suppress sending a message about who executed halt to the system log daemon, syslogd(1M), about who executed halt -n Prevent the sync(4) before stopping. -q Quick halt. No graceful shutdown is attempted. -y Halt the system, even from a dialup terminal. 				
FILES	/var/adm/wtmp History of user access and administration information				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
SUMMARY OF TRUSTED SOLARIS CHANGES	This command requires the sys_boot privilege and an effective uid of 0 in order to run.				
Trusted Solaris 8 4/01 Reference Manual	init(1M), reboot(1M)				
Solaris 9 1/01 Reference Manual	dumpadm(1M), shutdown(1M), sync(1M), syslogd(1M), attributes(5)				
NOTES	<p>Unlike shutdown(1M) and init(1M), halt does not execute the rc0 scripts.</p> <p>The poweroff utility is equivalent to init 5.</p>				

praudit(1M)

NAME	praudit – Print contents of an audit trail file				
SYNOPSIS	praudit [-lrs] [-ddel] [filename...]				
DESCRIPTION	<p>praudit reads the listed <i>filenames</i> (or standard input, if no <i>filename</i> is specified) and interprets the data as audit trail records as defined in <code>audit.log(4)</code>. By default, times, user and group IDs (UIDs and GIDs, respectively) are converted to their ASCII representation. Record type and event fields are converted to their ASCII representation. A maximum of 100 audit files can be specified on the command line.</p> <p>In the Trusted Solaris environment, the <code>praudit</code> command is run at the label <code>ADMIN_HIGH</code>, and requires either the <code>file_dac_read</code> privilege or an effective UID of 0 to succeed. The <code>PAF_LABEL_VIEW</code> process attribute flag for the current process affects how <code>ADMIN_HIGH</code> or <code>ADMIN_LOW</code> binary labels are translated to their text equivalents. See <code>pattr(1)</code> and <code>getpattr(2)</code> for more information.</p>				
OPTIONS	<p>-l Prints one line per record. The record type and event fields are always converted to their short text representation as is done for the <code>-s</code> option.</p> <p>-r Print records in their raw form. Times, UIDs, GIDs, record types, and events are displayed as integers. This option and the <code>-s</code> option are exclusive. If both are used, a format usage error message is output.</p> <p>-s Print records in their short form. All numeric fields are converted to text and displayed. The short text representations for the record type and event fields are used. This option and the <code>-r</code> option are exclusive. If both are used, a format usage error message is output.</p> <p>-ddel Use <i>del</i> as the field delimiter instead of the default delimiter, which is the comma. If <i>del</i> has special meaning for the shell, it must be quoted. The maximum size of a delimiter is four characters.</p>				
FILES	<p>/etc/security/audit_event Audit event definition and class mappings.</p> <p>/etc/security/audit_class Audit class definitions.</p>				
ATTRIBUTES	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>The <code>praudit</code> command uses the <code>PAF_LABEL_VIEW</code> process attribute flag and converts security labels as well as times and IDs. It is run at <code>ADMIN_HIGH</code> and requires either the <code>file_dac_read</code> privilege or an effective UID of 0 to succeed.</p> <p>The functionality described in this man page is available only if auditing is enabled. By default, auditing is enabled in the Trusted Solaris environment.</p>				

		praudit(1M)
Trusted Solaris 8 4/01 Reference Manual	audit(2), getauditflags(3BSM), audit.log(4), audit_class(4), audit_event(4)	
	<i>Trusted Solaris Audit Administration Manual</i>	
SunOS 5.8 Reference Manual	group(4), passwd(4), attributes(5)	

prtconf(1M)

NAME	prtconf – print system configuration
SPARC	<code>/usr/sbin/prtconf [-V] [-F] [-x] [-vpPD]</code>
IA	<code>/usr/sbin/prtconf [-V] [-x] [-vpPD]</code>
DESCRIPTION	The <code>prtconf</code> command prints the system configuration information. The output includes the total amount of memory, and the configuration of system peripherals formatted as a device tree.
OPTIONS	<p>The following options are supported:</p> <ul style="list-style-type: none">-D For each system peripheral in the device tree, displays the name of the device driver used to manage the peripheral.-F (SPARC only). Returns the device pathname of the console frame buffer, if one exists. If there is no frame buffer, <code>prtconf</code> returns a non-zero exit code. This flag must be used by itself. It returns only the name of the console, frame buffer device or a non-zero exit code. For example, if the console frame buffer on a SPARCstation 1 is <code>cgthree</code> in SBus slot #3, the command returns: <code>/sbus@1,f80000000/cgthree@3,0</code>. This option could be used to create a symlink for <code>/dev/fb</code> to the actual console device.-p Displays information derived from the device tree provided by the firmware (PROM) on SPARC platforms or the booting system on IA platforms.-P Includes information about pseudo devices. By default, information regarding pseudo devices is omitted.-v Specifies verbose mode.-V Displays platform-dependent PROM (on SPARC platforms) or booting system (on IA platforms) version information. This flag must be used by itself. The output is a string. The format of the string is arbitrary and platform-dependent.-x Reports if the firmware on this system is 64-bit ready. Some existing platforms may need a firmware upgrade in order to run the 64-bit kernel. If the operation is not applicable to this platform or the firmware is already 64-bit ready, it exits silently with a return code of zero. If the operation is applicable to this platform and the firmware is not 64-bit ready, it displays a descriptive message on <code>stdout</code> and exits with a non-zero return code. The hardware platform documentation contains more information about the platforms that may need a firmware upgrade in order to run the 64-bit kernel. <p>This flag overrides all other flags and must be used by itself.</p>

EXAMPLES**EXAMPLE 1** Running prtconf on a SPARC Sun4/65 Series Machine

Running prtconf on a Sun4/65 series machine produces the following sample output:

```
example% prtconf
System Configuration: Sun Microsystems sun4c
Memory size: 16 Megabytes
System Peripherals (Software Nodes):
Sun 4_65
  options, instance #0
  zs, instance #0
  zs, instance #1
  fd (driver not attached)
  audio (driver not attached)
  sbus, instance #0
    dma, instance #0
    esp, instance #0
      sd (driver not attached)
      st (driver not attached)
      sd, instance #0
      sd, instance #1 (driver not attached)
      sd, instance #2 (driver not attached)
      sd, instance #3
      sd, instance #4 (driver not attached)
      sd, instance #5 (driver not attached)
      sd, instance #6 (driver not attached)
    le, instance #0
    cgsix (driver not attached)
  auxiliary-io (driver not attached)
  interrupt-enable (driver not attached)
  memory-error (driver not attached)
  counter-timer (driver not attached)
  eeprom (driver not attached)
  pseudo, instance #0
```

EXAMPLE 2 Running prtconf on an IA Machine

Running prtconf on an IA machine produces the following sample output:

```
example% prtconf
System Configuration: Sun Microsystems i86pc
Memory size: 32 Megabytes
System Peripherals (Software Nodes):
i86pc
  eisa, instance #0
    kd, instance #0
    ata, instance #0
      cmdk, instance #0
    aha, instance #0
      cmdk, instance #1 (driver not attached)
      cmdk, instance #2 (driver not attached)
      cmdk, instance #3 (driver not attached)
      cmdk, instance #4 (driver not attached)
      cmdk, instance #5 (driver not attached)
```

prtconf(1M)

EXAMPLE 2 Running prtconf on an IA Machine (Continued)

```
cmdk, instance #6 (driver not attached)
cmdk, instance #7
chanmux, instance #0
asy, instance #0
asy, instance #1
elx, instance #0
elx, instance #1 (driver not attached)
elx, instance #2 (driver not attached)
elx, instance #3 (driver not attached)
fdc, instance #0
fd, instance #0
fd, instance #1
options, instance #0
objmgr, instance #0
pseudo, instance #0
example%
```

EXIT STATUS

The following exit values are returned:

0	No error occurred.
non-zero	With the -F option (SPARC only), a non-zero return value means that the output device is not a framebuffer. With the -x option, a non-zero return value means that the firmware is not 64-bit ready. In all other cases, a non-zero return value means that an error occurred.

SUMMARY OF TRUSTED SOLARIS CHANGES

The -v option of this command can be run from an administrative role.

The file_mac_read privilege is required in order to run the prtconf command.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWesu (32-bit)
	SUNWesxu (64-bit)

SEE ALSO

modinfo(1M), sysdef(1M), attributes(5)

Sun Hardware Platform Guide

SPARC Only

openprom(7D)

NOTES

The output of the prtconf command is highly dependent on the version of the PROM installed in the system. The output will be affected in potentially all circumstances.

prtconf(1M)

The `driver not attached` message means that no driver is currently attached to that instance of the device. In general, drivers are loaded and installed (and attached to hardware instances) on demand, and when needed, and may be uninstalled and unloaded when the device is not in use.

psradm(1M)

NAME	psradm – Change processor operational status
SYNOPSIS	psradm -f -i -n [-v] <i>processor_id</i> ... psradm -a-f -i -n [-v]
DESCRIPTION	<p>The psradm utility changes the operational status of processors. The legal states for the processor are <i>on-line</i>, <i>off-line</i>, and <i>no-intr</i>.</p> <p>An <i>on-line</i> processor processes LWPs (lightweight processes) and may be interrupted by I/O devices in the system.</p> <p>An <i>off-line</i> processor does not process any LWPs. Usually, an <i>off-line</i> processor is not interruptible by I/O devices in the system. On some processors or under certain conditions, it may not be possible to disable interrupts for an <i>off-line</i> processor. Thus, the actual effect of being <i>off-line</i> may vary from machine to machine.</p> <p>A <i>no-intr</i> processor processes LWPs but is not interruptible by I/O devices.</p> <p>A processor may not be taken <i>off-line</i> if there are LWPs that are bound to the processor. On some architectures, it might not be possible to take certain processors <i>off-line</i> if, for example, the system depends on some resource provided by the processor.</p> <p>At least one processor in the system must be able to process LWPs. At least one processor must also be able to be interrupted. Since an <i>off-line</i> processor may be interruptible, it is possible to have an operational system with one processor <i>no-intr</i> and all other processors <i>off-line</i> but with one or more accepting interrupts.</p> <p>If any of the specified processors are powered off, psradm may power on one or more processors.</p> <p>To succeed, this command needs the <code>sys_config</code> privilege.</p>
OPTIONS	<p>The following options are supported:</p> <ul style="list-style-type: none">-a Perform the action on all processors, or as many as possible.-f Take the specified processors <i>off-line</i>.-i Set the specified processors <i>no-intr</i>.-n Bring the specified processors <i>on-line</i>.-v Output a message giving the results of each attempted operation.
OPERANDS	<p>The following operands are supported:</p> <p><i>processor_id</i> The processor ID of the processor to be set <i>on-line</i> or <i>off-line</i>.</p>

EXAMPLES**EXAMPLE 1** Set processors 2 and 3 off-line.

The following example sets processors 2 and 3 off-line.

```
psradm -f 2 3
```

EXAMPLE 2 Set processors 1 and 2 no-intr.

The following example sets processors 1 and 2 no-intr.

```
psradm -i 1 2
```

EXAMPLE 3 Set all processors on-line.

The following example sets all processors on-line.

```
psradm -a -n
```

EXIT STATUS

The following exit values are returned:

0 Successful completion.

>0 An error occurred.

FILES`/etc/wtmpx` records logging processor status changes`/etc/wtmp` records logging processor status changes**ATTRIBUTES**See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SQUARES

Trusted Squares
4/01 Reference Manual
DIAGNOSTICS

To succeed, this command needs the `sys_config` privilege.`p_online(2)``psrinfo(1M)`, `psrset(1M)`, `attributes(5)``psradm: processor 4: Invalid argument`

The specified processor does not exist in the configuration.

`psradm: processor 3: Device busy`

The specified processor could not be taken off-line because it either has LWPs bound to it, is the last on-line processor in the system, or is needed by the system because it provides some essential service.

psradm(1M)

psradm: processor 3: Device busy

The specified processor could not be set no-intr because it is the last interruptible processor in the system, or or it is the only processor in the system that can service interrupts needed by the system.

psradm: processor 3: Device busy

The specified processor is powered off, and it cannot be powered on because some platform-specific resource is unavailable.

psradm: processor 0: Not owner

The user does not have permission to change processor status.

psradm: processor 2: Operation not supported

The specified processor is powered off, and the platform does not support power on of individual processors.

NAME	in.rarpd, rarpd – DARPA Reverse Address Resolution Protocol server
SYNOPSIS	<pre>/usr/sbin/in.rarpd [-d] -a</pre> <pre>/usr/sbin/in.rarpd [-d] device unit</pre>
DESCRIPTION	<p><code>in.rarpd</code> starts a daemon that responds to Reverse Address Resolution Protocol (RARP) requests. The daemon forks a copy of itself that runs in background. It must be started from the trusted path, with a UID of 0 and the label <code>ADMIN_LOW</code>. To succeed, it must inherit the <code>sys_net_conf</code> and <code>net_broadcast</code> privileges.</p> <p>RARP is used by machines at boot time to discover their Internet Protocol (IP) address. The booting machine provides its Ethernet address in a RARP request message. Using the <code>ethers</code> and <code>hosts</code> databases, <code>in.rarpd</code> maps this Ethernet address into the corresponding IP address which it returns to the booting machine in a RARP reply message. The booting machine must be listed in both databases for <code>in.rarpd</code> to locate its IP address. <code>in.rarpd</code> issues no reply when it fails to locate an IP address.</p> <p><code>in.rarpd</code> uses the STREAMS-based Data Link Provider Interface (DLPI) message set to communicate directly with the datalink device driver.</p>
OPTIONS	<p>The following options are supported:</p> <ul style="list-style-type: none"> -a Get the list of available network interfaces from IP using the <code>SIOCGIFADDR</code> ioctl and start a RARP daemon process on each interface returned. -d Print assorted debugging messages while executing.
EXAMPLES	<p>EXAMPLE 1 Starting an <code>in.rarpd</code> daemon for each network interface name returned from <code>/dev/ip</code>:</p> <p>The following command starts an <code>in.rarpd</code> for each network interface name returned from <code>/dev/ip</code>:</p> <pre>example# /usr/sbin/in.rarpd -a</pre> <p>EXAMPLE 2 Starting an <code>in.rarpd</code> daemon on the device <code>/dev/le</code> with the device instance number 0</p> <p>The following command starts one <code>in.rarpd</code> on the device <code>/dev/le</code> with the device instance number 0.</p> <pre>example# /usr/sbin/in.rarpd le 0</pre>
SUMMARY OF TRUSTED SOLARIS CHANGES	<code>in.rarpd</code> should be started from the trusted path with a UID0 and sensitivity label of <code>ADMIN_LOW</code> . It must inherit the <code>sys_net_config</code> and <code>net_broadcast</code> privileges.
FILES	<pre>/etc/ethers File or other source, as specified by nsswitch.conf(4).</pre> <pre>/etc/hosts File or other source, as specified by nsswitch.conf(4).</pre>

rarpd(1M)

/tftpboot Directory for remote boot scripts.
/dev/ip List of available network interfaces.
/dev/arp Address resolution protocol list.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

**Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual**

ifconfig(1M), nsswitch.conf(4)
boot(1M), ethers(4), hosts(4), netconfig(4), attributes(5), dlpi(7P)
RFC-903, *A Reverse Address Resolution Protocol*, Network Information Center, SRI
International.
Unix International, *Data Link Provider Interface*, Version 2, May 7, 1991, Sun
Microsystems, 800-6915-01.

NAME	rdate – Set system date from a remote host				
SYNOPSIS	rdate <i>hostname</i>				
DESCRIPTION	rdate sets the local date and time from the <i>hostname</i> given as an argument. This program needs to inherit the <code>sys_config</code> privilege to run properly. Typically <code>rdate</code> can be inserted as part of a startup script.				
USAGE	The <code>rdate</code> command is IPv6-enabled. See <code>ip6(7P)</code> .				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>This program needs to inherit the <code>sys_config</code> privilege to run properly.</p> <p><code>attributes(5)</code>, <code>ip6(7P)</code></p>				

Reference Manual

h2>rdisc(1M)

NAME	in.rdisc, rdisc – Network router discovery daemon
SYNOPSIS	<pre>/usr/sbin/in.rdisc [-a] [-f] [-s] [send-address] [receive-address] /usr/sbin/in.rdisc -r [-p preference] [-T interval] [send-address] [receive-address]</pre>
DESCRIPTION	<p>in.rdisc implements the ICMP router discovery protocol. The first form of the command is used on hosts and the second form is used on routers. On a host, in.rdisc is invoked at boot time to populate the network routing tables with default routes. On a router, it is also invoked at boot time in order to start advertising the router to all the hosts.</p>
Host (First Form)	<p>On a host, in.rdisc listens on the ALL_HOSTS (224.0.0.1) multicast address for ROUTER_ADVERTISE messages from routers. The received messages are handled by first ignoring those listed router addresses with which the host does not share a network. Among the remaining addresses, the ones with the highest preference are selected as default routers and a default route is entered in the kernel routing table for each one of them.</p> <p>Optionally, in.rdisc can avoid waiting for routers to announce themselves by sending out a few ROUTER_SOLICITATION messages to the ALL_ROUTERS (224.0.0.2) multicast address when it is started.</p> <p>A timer is associated with each router address. The address will no longer be considered for inclusion in the routing tables if the timer expires before a new <i>advertise</i> message is received from the router. The address will also be excluded from consideration if the host receives an <i>advertise</i> message with the preference being maximally negative.</p>
Router (Second Form)	<p>When in.rdisc is started on a router, it uses the SIOCGIFCONF ioctl(2) to find the interfaces configured into the system and it starts listening on the ALL_ROUTERS multicast address on all the interfaces that support multicast. It sends out <i>advertise</i> messages to the ALL_HOSTS multicast address advertising all its IP addresses. A few initial <i>advertise</i> messages are sent out during the first 30 seconds and after that it will transmit <i>advertise</i> messages approximately every 600 seconds.</p> <p>When in.rdisc receives a <i>solicitation</i> message, it sends an <i>advertise</i> message to the host that sent the <i>solicitation</i> message.</p> <p>When in.rdisc is terminated by a signal, it sends out an <i>advertise</i> message with the preference being maximally negative.</p>
OPTIONS	<p>-a Accept all routers independent of the preference they have in their <i>advertise</i> messages. Normally, in.rdisc only accepts (and enters in the kernel routing tables) the router or routers with the highest preference.</p> <p>-f Run in.rdisc forever even if no routers are found. Normally, in.rdisc gives up if it has not received any <i>advertise</i> message</p>

-r	Act as a router, rather than a host.
-s	Send three <i>solicitation</i> messages initially to quickly discover the routers when the system is booted. When -s is specified, <code>in.rdisc</code> exits with a non-zero exit code if it can not find any routers. This can be overridden with the -f option.
-p <i>preference</i>	Set the preference transmitted in the <i>solicitation</i> messages. The default is zero.
-T <i>interval</i>	Set the interval between transmitting the <i>advertise</i> messages. The default time is 600 seconds.

`in.rdisc` must be started from the trusted path. To modify kernel routing tables, it must inherit the `sys_net_config` privilege. To open a raw socket, it needs the `net_rawaccess` privilege. To send multicast or broadcast packets, it needs the `net_broadcast` privilege.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

ioctl(2), attributes(5), icmp(7P), inet(7P)

Deering, S.E., editor, *ICMP Router Discovery Messages*, RFC 1256, Network Information Center, SRI International, Menlo Park, California, September 1991.

reboot(1M)

NAME	reboot – Restart the operating system				
SYNOPSIS	/usr/sbin/reboot [-dlmq] [<i>boot arguments</i>]				
DESCRIPTION	<p>The reboot utility restarts the kernel. The kernel is loaded into memory by the PROM monitor, which transfers control to the loaded kernel.</p> <p>Although reboot can be run by an administrative role with appropriate privilege and authorization at any time, shutdown(1M) is normally used first to warn all users logged in of the impending loss of service.</p> <p>The reboot utility performs a sync(1M) operation on the disks, and then a multi-user reboot is initiated. See init(1M) for details.</p> <p>The reboot utility normally logs the reboot to the system log daemon, syslogd(1M), and places a shutdown record in the login accounting file /var/adm/wtmpx. These actions are inhibited if the -n or -q options are present.</p> <p>Normally, the system will reboot itself at power-up or after crashes.</p>				
OPTIONS	<p>-d Force a system crash dump before rebooting. See dumpadm(1M) for information on configuring system crash dumps.</p> <p>-l Suppress sending a message to the system log daemon, syslogd(1M) about who executed reboot.</p> <p>-n Avoid the sync(1M) operation. Use of this option can cause file system damage.</p> <p>-q Quick. Reboot quickly and ungracefully, without shutting down running processes first.</p> <p><i>bootarguments</i> These arguments are accepted for compatibility, and are passed unchanged to the uadmin(2) system call.</p>				
EXAMPLES	<p>EXAMPLE 1 Example of the reboot utility.</p> <p>In the example below, the delimiter ‘—’ (two hyphens) must be used to separate the options of reboot from the arguments of boot(1M).</p> <pre>example# reboot -dl — -rv</pre>				
FILES	/var/adm/wtmpx Login accounting file				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				

reboot(1M)

reboot requires the sys_boot privilege in order to run. Shutting down a computer requires authorization.

halt(1M), init(1M), uadmin(2)

boot(1M), crash(1M), dumpadm(1M), fsck(1M), shutdown(1M), sync(1M),
syslogd(1M), reboot(3C), attributes(5)

reject(1M)

NAME	accept, reject – Accept or reject print requests
SYNOPSIS	accept <i>destination</i> ... reject [-r <i>reason</i>] <i>destination</i> ...
DESCRIPTION	<p>accept allows the queueing of print requests for the named destinations.</p> <p>reject prevents queueing of print requests for the named destinations.</p> <p>Use lpstat -a to check if destinations are accepting or rejecting print requests.</p> <p>accept and request must be run on the print server; they have no meaning on a client system.</p>
OPTIONS	<p>The following options are supported for reject.</p> <p>-r <i>reason</i> Assigns a reason for rejection of print requests for <i>destination</i>. Enclose <i>reason</i> in quotes if it contains blanks. <i>reason</i> is reported by lpstat -a. By default, <i>reason</i> is unknown reason for existing destinations, and new printer for destinations added to the system but not yet accepting requests.</p>
OPERANDS	<p>The following operands are supported.</p> <p><i>destination</i> The name of the destination accepting or rejecting print requests. Destination specifies the name of a printer or class of printers [see lpadmin(1M)]. Specify <i>destination</i> using atomic name. See printers.conf(4) for information regarding the naming conventions for atomic names.</p>
EXIT STATUS	<p>The following exit values are returned:</p> <p>0 Successful completion.</p> <p>non-zero An error occurred.</p>
FILES	/var/spool/lp/* LP print queue.
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWpcu
CSI	Enabled (see NOTES)

SUMMARY OF TRUSTED SOLARIS CHANGES	See the lpstat(1) man page for Trusted Solaris security that affects the accept command.
Trusted Solaris 8 4/01 Reference Manual	enable(1) , lp(1) , lpstat(1) , lpadmin(1M) , lpsched(1M)

reject(1M)

printers.conf (4), attributes(5)

accept and reject only affect queuing on the print server's spooling system. Requests made from a client system remain queued in the client system's queuing mechanism until they are cancelled or accepted by the print server's spooling system.

accept is CSI-enabled except for the *destination* name.

rem_drv(1M)

NAME	rem_drv – Remove a device driver from the system				
SYNOPSIS	rem_drv [-b <i>basedir</i>] <i>device_driver</i>				
DESCRIPTION	<p>The rem_drv command informs the system that the device driver <i>device_driver</i> is no longer valid. If possible, rem_drv unloads <i>device_driver</i> from memory. Entries for the device in the /devices namespace are removed. rem_drv also updates the system driver configuration files.</p> <p>If rem_drv has been executed, the next time the system is rebooted it will automatically perform a reconfiguration boot (see kernel(1M)).</p>				
OPTIONS	<p>-b <i>basedir</i> Sets the path to the root directory of the diskless client. Used on the server to execute rem_drv for a client. The client machine must be rebooted to unload the driver.</p>				
EXAMPLES	<p>EXAMPLE 1 Examples of rem_drv.</p> <p>The following example removes the sd driver from use:</p> <pre>example% rem_drv sd</pre> <p>The next example removes the driver from the sun1 diskless client. The driver will not be uninstalled nor unloaded until the client machine is rebooted.</p> <pre>example% rem_drv -b /export/root/sun1 sd</pre>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
SUMMARY OF TRUSTED SOLARIS CHANGES	To succeed, this command needs the sys_devices privilege. This command is intended to be invoked at ADMIN_LOW with effective user ID 0; if invoked by other users, this command needs the file_dac_write privilege.				
Trusted Solaris 8 4/01 Reference Manual	<p>add_drv(1M), drvconfig(1M)</p> <p>kernel(1M), attributes(5)</p>				

NAME	remove_allocatable – remove entries from allocation databases								
SYNOPSIS	/usr/sbin/remove_allocatable [-f] -n <i>name</i>								
DESCRIPTION	<p>remove_allocatable removes database entries for allocatable devices and certain non-allocatable devices to control access to an allocatable device.</p> <p>remove_allocatable removes the device's entries from the device_allocate(4) and device_maps(4) databases.</p>								
OPTIONS	<p>-f Force the removal of an entry. remove_allocatable exits with an error if this option is not specified when an entry with the specified device name no longer exists.</p> <p>-n <i>name</i> Removes the entry for a device named <i>name</i> from the device_allocate and device_maps.</p>								
ERRORS	<p>When successful, remove_allocatable returns an exit status of 0 (true). remove_allocatable returns a nonzero exit status in the event of an error. The exit codes are as follows:</p> <table> <tr> <td>1</td><td>Invocation syntax error</td></tr> <tr> <td>2</td><td>Unknown system error</td></tr> <tr> <td>3</td><td>Device <i>name</i> not found. This error occurs only when the -f option is not specified.</td></tr> <tr> <td>4</td><td>Permission denied. User does not have DAC or MAC access to database.</td></tr> </table>	1	Invocation syntax error	2	Unknown system error	3	Device <i>name</i> not found. This error occurs only when the -f option is not specified.	4	Permission denied. User does not have DAC or MAC access to database.
1	Invocation syntax error								
2	Unknown system error								
3	Device <i>name</i> not found. This error occurs only when the -f option is not specified.								
4	Permission denied. User does not have DAC or MAC access to database.								
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWtsu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu				
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWtsu								
FILES	<p>/etc/security/device_allocate Mandatory access control file for devices</p> <p>/etc/security/device_maps List of physical devices associated with a device name and type</p>								
Trusted Solaris 8 4/01 Reference Manual	allocate(1), device_allocate(4), device_clean(1M), device_maps(4), add_allocatable(1M)								
SunOS 5.8 Reference Manual	attributes(5)								

rexecd(1M)

NAME	in.rexecd, rexecd – Remote execution server
SYNOPSIS	in.rexecd
DESCRIPTION	<p>in.rexecd is the server for the <code>rexec(3SOCKET)</code> routine. The server provides remote execution facilities with authentication based on user names and passwords. It is invoked automatically as needed by <code>inetd(1M)</code>, and then executes the following protocol:</p> <ol style="list-style-type: none">1) The server reads characters from the socket up to a null (<code>\\0</code>) byte. The resultant string is interpreted as an ASCII number, base 10.2) If the number received in step 1 is non-zero, it is interpreted as the port number of a secondary stream to be used for the <code>stderr</code>. A second connection is then created to the specified port on the client's machine.3) A null terminated user name of at most 16 characters is retrieved on the initial socket.4) A null terminated password of at most 16 characters is retrieved on the initial socket.5) A null terminated command to be passed to a shell is retrieved on the initial socket. The length of the command is limited by the upper bound on the size of the system's argument list.6) <code>rexecd</code> then validates the user as is done at login time and, if the authentication was successful, changes to the user's home directory, and establishes the user and group protections of the user. Access is denied unless the user has the remote login authorization. If the <code>/etc/nologin</code> file exists, access is denied. If any of these steps fail the connection is aborted and a diagnostic message is returned.7) A null byte is returned on the connection associated with the <code>stderr</code> and the command line is passed to the normal login shell of the user. The shell inherits the network connections established by <code>rexecd</code>.
USAGE	in.rexecd and rexecd are IPv6-enabled. See <code>ip6(7P)</code> .
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES Trusted Solaris 8 4/01 Reference Manual	<p>Login is not allowed unless the user has the <code>remote login</code> authorization. If the <code>/etc/nologin</code> file exists, the user is not allowed to log in.</p> <p><code>inetd(1M)</code>, <code>inetd.conf(4)</code></p> <p><code>rexec(3SOCKET)</code>, <code>attributes(5)</code>, <code>ip6(7P)</code></p>
--	--

DIAGNOSTICS All diagnostic messages are returned on the connection associated with the `stderr`, after which any network connections are closed. An error is indicated by a leading byte with a value of 1 (0 is returned in step 7 above upon successful completion of all the steps prior to the command execution).

<code>username too long</code>	The name is longer than 16 characters.
<code>password too long</code>	The password is longer than 16 characters.
<code>command too long</code>	The command line passed exceeds the size of the argument list (as configured into the system).
<code>Login incorrect</code>	No password file entry for the user name existed.
<code>Password incorrect</code>	The wrong password was supplied.
<code>No remote directory</code>	The <code>chdir</code> command to the home directory failed.
<code>Try again.</code>	A fork by the server failed.
<code>/usr/bin/sh: ...</code>	The user's login shell could not be started.

rlogind(1M)

NAME	in.rlogind, rlogind – Remote login server
SYNOPSIS	/usr/sbin/in.rlogind -U -T
DESCRIPTION	<p>in.rlogind is the server for the rlogin(1) program. The server provides a remote login facility with authentication based on privileged port numbers.</p> <p>in.rlogind is invoked by inetd(1M) when a remote login connection is established, and executes the following protocol:</p> <ul style="list-style-type: none">■ The server checks the client's source port. If the port is not in the range 0-1023, the server aborts the connection.■ The server checks the client's source address. If an entry for the client exists in both /etc/hosts and /etc/hosts.equiv, a user logging in from the client is not prompted for a password. If the address is associated with a host for which no corresponding entry exists in /etc/hosts, the user is prompted for a password, regardless of whether an entry for the client is present in /etc/hosts.equiv. See hosts(4) and hosts.equiv(4). <p>Once the source port and address have been checked, in.rlogind allocates a pseudo-terminal and manipulates file descriptors so that the slave half of the pseudo-terminal becomes the stdin, stdout, and stderr for a login process. The login process is an instance of the login(1) program, invoked with the -r.</p> <p>The login process then proceeds with the in.rshd(1M) authentication process.</p> <p>The parent of the login process manipulates the master side of the pseudo-terminal, operating as an intermediary between the login process and the client instance of the rlogin program. In normal operation, a packet protocol is invoked to provide Ctrl-S and Ctrl-Q type facilities and propagate interrupt signals to the remote programs. The login process propagates the client terminal's baud rate and terminal type, as found in the environment variable, TERM; see environ(4).</p> <p>The -U option is used to pass the UID of the client to login(1). The -T option is used if the client has the trusted path attribute.</p>
USAGE	rlogind and in.rlogind are IPv6-enabled. See ip6(7P).
SUMMARY OF TRUSTED ATTRIBUTES CHANGES	<p>Two new options (-U and -T) are used in the call to login(1).</p> <p>See attributes(5) for descriptions of the following attributes:</p>

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8
4/01 Reference
Manual
New and
Reference Manual

login(1), in.rshd(1M), inetd(1M), inetd.conf(4)
rlogin(1), environ(4), hosts(4), hosts.equiv(4), attributes(5), ip6(7P)

DIAGNOSTICS	<p>All diagnostic messages are returned on the connection associated with the <code>stderr</code>, after which any network connections are closed. An error is indicated by a leading byte with a value of 1.</p> <p>Hostname for your address unknown. No entry in the host name database existed for the client's machine.</p> <p>Try again. A <i>fork</i> by the server failed.</p> <p>/usr/bin/sh: ... The user's login shell could not be started.</p>
NOTES	<p>The authentication procedure used here assumes the integrity of each client machine and the connecting medium. This is insecure, but it is useful in an "open" environment.</p> <p>A facility to allow all data exchanges to be encrypted should be present.</p>

rmmount(1M)

NAME	rmmount – removable media mounter for CD-ROM and floppy
SYNOPSIS	<code>/usr/sbin/rmmount [-D]</code>
DESCRIPTION	<p>The <code>rmmount</code> utility is a removable media mounter that is executed by Volume Management whenever a CD-ROM or floppy is inserted. The Volume Management daemon, <code>vold(1M)</code>, manages CD-ROM and floppy devices. <code>rmmount</code> is also called by <code>volrmmount(1)</code>, and by device allocation clean scripts invoked by the <code>allocate(1)</code> command.</p> <p>If the media is read-only (either CD-ROM or floppy with write-protect tab set), the file system is mounted read-only.</p> <p>If a file system is not identified, <code>rmmount</code> does not mount a file system. See the <i>System Administration Guide, Volume 1</i> for more information on the location of CD-ROM and floppy media without file systems. Also see <code>volfs(7FS)</code>.</p> <p>If a file system type has been determined, it is then checked to see that it is “clean.” If the file system is “dirty,” <code>fsck -p</code> (see <code>fsck(1M)</code>) is run in an attempt to clean it. If <code>fsck</code> fails, the file system is mounted read-only.</p> <p>After the mount is complete, “actions” associated with the media type are executed. These actions allow for the notification to other programs that new media are available. These actions are shared objects and are described in the configuration file, <code>/etc/rmmount.conf</code>.</p> <p>Actions are executed in the order in which they appear in the configuration file. The action function can return either 1 or 0. If it returns 0, no further actions will be executed. This allows the function to control which applications are executed.</p> <p>In order to execute an action, <code>rmmount</code> performs a <code>dlopen(3DL)</code> on the shared object and calls the action function defined within it. The definition of the interface to actions can be found in <code>/usr/include/rmmount.h</code>. The actions are run in the <code>rmmount</code> process, that is, at label <code>ADMIN_LOW</code>, with a UID of 0, and with the privileges of <code>rmmount</code> (see SUMMARY OF TRUSTED SOLARIS CHANGES). If an action does not require privileges, it should fork a child process to do the work of the action, and the child process whould change its effective UID. For example, if the action should run as the user logged in on the console, the process UID should be set to the UID of <code>/dev/console</code>.</p> <p>File systems mounted by <code>rmmount</code> are always mounted with the <code>nosuid</code> flag set, thereby disabling set-uid programs and access to block or character devices in that file system. Upon ejection, <code>rmmount</code> unmounts mounted file systems and executes actions associated with the media type. If a file system is “busy” (that is, it contains the current working directory of a live process), the ejection will fail.</p>
OPTIONS	<code>-D</code> Turn on the debugging output from the <code>rmmount</code> <code>dprintf</code> calls.
ENVIRONMENT VARIABLES	Several environment variables must be passed to <code>rmmount</code> .

VOLUME_ACTION	The event type to be handled. The event types are insert, remount, clear_mounts, eject, and unmount.
VOLUME_MEDIATYPE	The type of media be mounted or unmounted, for example, cdrom or floppy
VOLUME_NAME	A name for the mounted media. This is typically the volume name taken from the media's label.
VOLUME_PATH	Pathname of the device special file for the device on which the filesystem resides.
VOLUME_SYMDEV	Symbolic name of the device containing the volume, for example, cdrom0 or floppy0.

Two other environment variables may optionally be passed to rmmount.

VOLUME_PCFS_ID	Partition letter for a PCFS file system.
VOLUME_MOUNT_MODE	Mount mode. If not supplied, the default mode is "r" for read-only file systems, and "rw" for others.

The rmmount command mounts the media on device \$VOLUME_PATH at the mount point /\$VOLUME_MEDIATYPE/\$VOLUME_NAME. It also creates a symbolic link \$VOLUME_MEDIATYPE/\$VOLUME_SYMDEV pointing to the /\$VOLUME_MEDIATYPE/\$VOLUME_NAME mount point.

With the values of the environment variables passed by vold(1M), the file system is mounted in one of the following locations:

Mount Location	State of Media
/floppy/floppy0	symbolic link to mounted floppy in local floppy drive
/floppy/floppy_name	mounted named floppy
/floppy/unnamed_floppy	mounted unnamed floppy
/cdrom/cdrom0	symbolic link to mounted CD-ROM in local CD-ROM drive
/cdrom/CD-ROM_name	mounted named CD-ROM
/cdrom/CD-ROM_name/partition	mounted named CD-ROM with partitioned file system
/cdrom/unnamed_cdrom	mounted unnamed CD-ROM
/etc/rmmount.conf	removable media mounter configuration file.
/usr/lib/rmmount/*.so.1	shared objects used by rmmount.

FILES

rmmount(1M)

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWvolu

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**

The rmmount command must inherit the file_chown, file_dac_read, file_dac_write, file_mac_read, file_mac_write, file_owner, file_setdac, and sys_mount privileges.

**Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual**

allocate(1)

volcancel(1), volcheck(1), volmissing(1), volrmmount(1), fsck(1M), vold(1M), dlopen(3DL), rmmount.conf(4), vold.conf(4), attributes(5), volfs(7FS)

System Administration Guide, Volume 1

NAME	route – Manually manipulate the routing tables
SYNOPSIS	<pre> route [-fnvq] <i>command</i> [[<i>modifiers</i>] <i>args</i>] route [-fnvq] -t <i>file1</i> route [-fnvq] add delete [<i>host</i> <i>net</i>]<i>destination</i> [<i>gateway</i> [<i>args</i>]] [<i>extended_metric</i>] route [-fnvq] change get [<i>host</i> <i>net</i>]<i>destination</i> [<i>gateway</i> [<i>args</i>]] [<i>extended_metric</i>] route [-n] monitor route [-n] flush </pre>
DESCRIPTION	<p>route manually manipulates the network routing tables. These tables are normally maintained by the system routing daemon, such as <i>in.routed</i>(1M) and <i>in.ripngd</i>(1M).</p> <p>This utility supports a limited number of general options, but a rich command language. It enables the user to specify any arbitrary request that could be delivered via the programmatic interface discussed in <i>route</i>(7P).</p> <p>route uses a routing socket and the new message types RTM_ADD, RTM_DELETE, RTM_GET, and RTM_CHANGE.</p> <p>The route utility must inherit the <i>sys_net_config</i> privilege to operate directly on the routing table for the specific host or network indicated by <i>destination</i>. It also must run with a uid of 0, or have the <i>file_dac_read</i> and <i>file_dac_write</i> privileges.</p>
Trusted Solaris Routing	<p>In the Trusted Solaris environment, as in the Solaris operating environment, routing can be configured by the administrator to be <i>static</i>, that is, determined by a router or route description in a static file, or <i>dynamic</i>, that is, determined at the time of routing. By default, dynamic routing is in effect. The static routing file to establish default gateways (routers) in both environments is <i>/etc/defaultrouter</i>.</p> <p>In the Trusted Solaris environment, an administrator can establish gateways for specific networks and default gateways in the file <i>/etc/tsolgateways</i>. Routing decisions are made in the order:</p> <ol style="list-style-type: none"> 1. Get routing information from the file <i>/etc/tsolgateways</i>. If it does not exist, 2. Get routing information from the file <i>/etc/defaultrouter</i>. If it does not exist, 3. Start dynamic routing. <p>For trusted routing, security attributes must also be associated with a route. The additional security routing information (SRI) includes sensitivity label range, CIPSO DOI, RIPS0 label, and RIPS0 error. The SRI and the simple metric together compose the extended metric (<i>extended_metric</i>), which is necessary for trusted routing.</p>

route(1M)

If `-e` is specified in *extended_metric*, the metric and SRI are obtained from a file composed of a series of lines, each specifying an extended metric value. Both the `-e` and `-m` options use the same format. For readability only, the one-line format is shown here as two lines:

```
metric= val,min_sl=val,max_sl=val,doi= val
ripso_label= val,ripso_error=val,ripso_only,cipso_only
```

The *val* for *metric* is an integer from 1 to 15. The *val* for *min_sl* and *max_sl* is a sensitivity label in either hex or string form. The *val* for *doi* is a nonzero integer. The *val* for *ripso_label* is the classification, followed by a space, followed by a list of protections separated by semicolons (;). Both the classification and the protections are specified either in hex or string form. The *val* for *ripso_error* is a list of protections separated by semicolons (;). They are specified in either hex or string form. Note that if *val* contains a space, it must be protected by double quotes.

The two keywords, *ripso_only* and *cipso_only*, do not have values. They indicate that a route can only forward packets having RIPSO or CIPSO labels. They must be specified if a SUN_RIPSO or SUN_CIPSO gateway is involved in a route.

Some keywords are necessary, and others are optional. The following rules apply when specifying the extended metric information.

- *metric*, *min_sl*, and *max_sl* must be specified.
- *ripso_label* and *ripso_error* must both be present or both be absent.
- If *cipso_only* is specified, *doi* must be specified; and no *ripso_label*, *ripso_error*, or *ripso_only* can be specified.
- If *ripso_only* is specified, *ripso_label* and *ripso_error* must be specified; and no *doi* or *cipso_only* can be specified.

Note – When the `-e` option is used, emetrics are generated for a route. These emetrics are used for accreditation checks when selecting a route.

Without the `-e` option, no emetric is generated. If the command adds a remote route, the template of the gateway will be used for accreditation checks when selecting a route, since no emetric is available.

OPTIONS

- | | |
|-----------------|---|
| <code>-f</code> | Flush the routing tables of all gateway entries. If this is used in conjunction with one of the commands described above, <code>route</code> flushes the gateways before performing the command. |
| <code>-n</code> | Prevent attempts to print host and network names symbolically when reporting actions. This is useful, for example, when all name servers are down on your local net, and you need a route before you can contact the name server. |
| <code>-v</code> | (Verbose) Print additional details. |
| <code>-q</code> | Suppress all output. |

	<p><code>-t <i>file1</i></code> Obtain extended metric information from <i>file1</i>, where <i>file1</i> has the same format as <code>/etc/tsolgateways</code>. This option always implies <code>add</code> command, i.e. add a route.</p>												
Commands	<p><code>route</code> executes one of four <i>commands</i> on a route to a <i>destination</i>. Two additional <i>commands</i> operate globally on all routing information. The (six) commands are:</p> <table> <tr> <td><code>add</code></td><td>Add a route.</td></tr> <tr> <td><code>change</code></td><td>Change aspects of a route (such as its gateway).</td></tr> <tr> <td><code>delete</code></td><td>Delete a specific route.</td></tr> <tr> <td><code>flush</code></td><td>Remove all gateway entries from the routing table.</td></tr> <tr> <td><code>get</code></td><td>Look up and display the route for a destination.</td></tr> <tr> <td><code>monitor</code></td><td>Continuously report any changes to the routing information base, routing lookup misses, or suspected network partitionings.</td></tr> </table> <p>The <code>add</code>, <code>delete</code>, and <code>change</code> commands have the following syntax:</p> <pre>route [-fnvq] add delete [-net -host] dest gate [args] [ext_metric]</pre> <p>or</p> <pre>route [-fnvq] change get [-net -host] dest gate [args] [ext_metric]</pre> <p>where <i>dest</i> is the destination host or network, <i>gate</i> is the next-hop intermediary through which packets should be routed, and <i>ext_metric</i> is one of</p> <p><code>-e <i>file</i></code></p> <p>or</p> <p><code>-m <i>emetric_val</i> ... -m <i>emetric_val</i></code></p> <p>The <code>-e</code> or <code>-m</code> option is required for <code>add</code> commands, and must be nonzero if the route utilizes one or more gateways. These options are used to specify extended metric information associated with a route. See the explanation in the section <i>Trusted Solaris Routing</i> under <i>DESCRIPTION</i>.</p>	<code>add</code>	Add a route.	<code>change</code>	Change aspects of a route (such as its gateway).	<code>delete</code>	Delete a specific route.	<code>flush</code>	Remove all gateway entries from the routing table.	<code>get</code>	Look up and display the route for a destination.	<code>monitor</code>	Continuously report any changes to the routing information base, routing lookup misses, or suspected network partitionings.
<code>add</code>	Add a route.												
<code>change</code>	Change aspects of a route (such as its gateway).												
<code>delete</code>	Delete a specific route.												
<code>flush</code>	Remove all gateway entries from the routing table.												
<code>get</code>	Look up and display the route for a destination.												
<code>monitor</code>	Continuously report any changes to the routing information base, routing lookup misses, or suspected network partitionings.												
OPERANDS	<code>route</code> executes its commands on routes to destinations.												
Destinations	<p>By default, a destination is looked up under the <code>AF_INET</code> address family or as an IPv4 address. All symbolic names specified for a destination or gateway are looked up first as a host name, using <code>getipnodebyname(3SOCKET)</code>. If this lookup fails in the <code>AF_INET</code> case, <code>getnetbyname(3SOCKET)</code> is used to interpret the name as that of a network.</p> <p>An optional modifier may be included on the command line before a <i>destination</i>, to force how <code>route</code> interprets a destination:</p>												

route(1M)

- host Forces the destination to be interpreted as a host.
- net Forces the destination to be interpreted as a network.
- inet Forces the destination to be interpreted under the AF_INET address family or as an IPv4 address.
- inet6 Forces the destination to be interpreted under the AF_INET6 address family or as an IPv6 address.

In the case of the AF_INET address family or an IPv4 address, routes to a particular host may be distinguished from those to a network by interpreting the Internet address specified as the *destination*. If the *destination* has a “local address part” of INADDR_ANY, or if the *destination* is the symbolic name of a network, then the route is assumed to be to a network; otherwise, it is presumed to be a route to a host.

For example:

The following route:	Is interpreted as:
192.168	-host 192.0.0.32
192.168.130	-host 192.168.0.130
-net 192.168	192.168.0.0
-net 192.168.130	192.168.130.0

If the destination is directly reachable by way of an interface requiring no intermediary system to act as a gateway, this can be indicated by including one of two optional modifiers after the destination: The *interface* modifier can be included or a *metric* of 0 can be specified. These modifiers are illustrated in the following alternative examples:

```
example% route add default hostname -interface
example% route add default hostname 0
```

hostname is the name or IP address associated with the network interface all packets should be sent over. On a host with a single network interface, *hostname* is normally the same as the nodename returned by **uname -n** (see `uname(1)`).

In the above examples, the route does not refer to a gateway, but rather to one of the machine’s interfaces. Destinations matching such a route are sent out on the interface identified by the *gateway* address. For interfaces using the ARP protocol, this type of route is used to specify *all destinations are local*. That is, a host should use ARP for all addresses by adding a default route using one of the two commands listed above.

With the AF_INET address family or an IPv4 address, the optional *-netmask* qualifier is intended to manually add subnet routes with netmasks different from that of the implied network interface. The implicit network mask generated in the AF_INET case

can be overridden by making sure this option, and an ensuing address parameter (to be interpreted as a network mask), follows the destination parameter.

Alternatively, the length of the netmask may be supplied by appending a slash character and the length immediately after the destination. For example:

```
example% route add 192.0.2.32/27 somegateway
```

will create an IPv4 route to the destination 192.0.2.32 with a netmask of 255.255.255.224, and

```
example% route add -inet6 fec0::/16 somegateway
```

will create an IPv6 route to the destination fec0:: with a netmask of 16 one-bits followed by 112 zero-bits.

Routing Flags

Routes have associated flags which influence operation of the protocols when sending to destinations matched by the routes. These flags may be set (or sometimes cleared) by including the following corresponding modifiers on the command line:

Modifier	Flag	Description
-cloning	RTF_CLONING	generates a new route on use
-xresolve	RTF_XRESOLVE	emit mesg on use (for external lookup)
-iface	~RTF_GATEWAY	destination is directly reachable
-static	RTF_STATIC	manually added route
-nostatic	~RTF_STATIC	pretend route added by kernel or daemon
-reject	RTF_REJECT	emit an ICMP unreachable when matched
-blackhole	RTF_BLACKHOLE	silently discard pkts (during updates)
-proto1	RTF_PROTO1	set protocol specific routing flag #1
-proto2	RTF_PROTO2	set protocol specific routing flag #2
-private	RTF_PRIVATE	do not advertise this route

The optional modifiers `-rtt`, `-rttvar`, `-sendpipe`, `-recvpipe`, `-mtu`, `-hopcount`, `-expire`, and `-ssthresh` provide initial values to quantities maintained in the routing entry by transport level protocols, such as TCP. These may be individually locked either by preceding each modifier to be locked by the `-lock` meta-modifier, or by specifying that all ensuing metrics may be locked by the `-lockrest` meta-modifier.

The optional modifiers are defined as follows:

route(1M)

-expire	Lifetime for the entry. This optional modifier is not currently supported.
-hopcount	Maximum hop count. This optional modifier is not currently supported.
-mtu	Maximum MTU in bytes.
-recvpipe	Receive pipe size in bytes.
-rtt	Round trip time in microseconds.
-rttvar	Round trip time variance in microseconds.
-sendpipe	Send pipe size in bytes.
-ssthresh	Send pipe size threshold in bytes.

Some transport layer protocols may support only some of these metrics.

In a change or add command where the destination and gateway are not sufficient to specify the route (for example, , when several interfaces have the same address), the -ifp or -ifa modifiers may be used to determine the interface or interface address.

FILES

/etc/hosts	List of host names and net addresses.
/etc/networks	List of network names and addresses.
/etc/defaultrouter	List of default routes.
/etc/tsolgateways	List of trusted gateways and metrics.
/etc/security/tsol/device_policy	Policy for trusted devices.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

route must inherit the sys_net_config privilege to operate directly on the routing table for the specific host or network indicated by *destination*. It also must run with a uid of 0, or have the file_dac_read and file_dac_write privileges. The file /etc/security/tsol/device_policy specifies the access policy for the device special files used by route.

The Trusted Solaris environment adds the file /etc/tsolgateways and the -t option to read it. It also adds the *extended_metric* arguments to handle security routing information.

Trusted Solaris 8 4/01 Reference Manual

uname(1), in.rdisc(1M), netstat(1M), routed(1M), device_policy(4), tsolgateways(4)

**SunOS 5.8
Reference Manual**

DIAGNOSTICS

Trusted Solaris Administrator's Procedures

get(1), ioctl(2), getipnodebyname(3SOCKET), getnetbyname(3SOCKET),
hosts(4), networks(4), attributes(5), ARP(7P), route(7P), routing(7P)

add [host | network] destination:gateway flags

The specified route is being added to the tables. The values printed are from the routing table entry supplied in the ioctl(2) call. If the gateway address used was not the primary address of the gateway (the first one returned by gethostbyname(3NSL)) the gateway address is printed numerically as well as symbolically.

delete [host | network] destination:gateway flags

As above, but when deleting an entry.

destination done

When the -f flag is specified, or in the flush command, each routing table entry deleted is indicated with a message of this form.

Network is unreachable

An attempt to add a route failed because the gateway listed was not on a directly-connected network. Give the next-hop gateway instead.

not in table

A delete operation was attempted for an entry that is not in the table.

routing table overflow

An add operation was attempted, but the system was unable to allocate memory to create the new entry.

NOTES

All destinations are local assumes that the routers implement the protocol, proxy arp. Normally, using router discovery (see in.rdisc(1M)) is more reliable than using proxy arp.

Combining the *all destinations are local* route with subnet or network routes can lead to unpredictable results: the search order as it relates to the *all destinations are local* route are undefined and may vary from release to release.

routed(1M)

NAME	<code>in.routed</code> , <code>routed</code> – Network routing daemon
SYNOPSIS	<code>/usr/sbin/in.routed</code> [-s] [-q] [-t] [-g] [-S] [-v] [<i>logfile</i>]
DESCRIPTION	<p><code>in.routed</code> is invoked at boot time to manage the network routing tables. The routing daemon uses a variant of the Xerox NS Routing Information Protocol in maintaining up-to-date kernel routing table entries.</p> <p>In normal operation, <code>in.routed</code> listens on udp(7P) socket 520 (decimal) for routing information packets. If the host is an internetwork router, it periodically supplies copies of its routing tables to any directly connected hosts and networks.</p> <p>When <code>in.routed</code> is started, it uses the <code>SIOCGIFCONF</code> <code>ioctl</code>(2) to find those directly connected interfaces configured into the system and marked “up” (the software loopback interface is ignored). If multiple interfaces are present, it is assumed the host will forward packets between networks. <code>in.routed</code> then transmits a <i>request</i> packet on each interface (using a broadcast packet if the interface supports it) and enters a loop, listening for <i>request</i> and <i>response</i> packets from other hosts.</p> <p>For trusted routing, extended security attributes must be associated with a route along with the simple metric that indicates the number of hops to the destination. The additional security routing information (SRI) includes a sensitivity label range, and can include a CIPSO domain of interpretation, a RIPSO label, and a RIPSO error, and some additional keywords: <code>ripso_only</code>, <code>cipso_only</code>, and <code>msix_only</code>. The SRI combined with the simple metric is called the extended metric, or <i>emetric</i>.</p> <p>For Trusted Solaris 7 and later systems, two additional types of packets are exchanged. The first one is <i>sec_response</i>, which is like the <i>response</i> packet but also carries the SRI for the routes. Similar to the <i>response</i> packet, the <i>sec_response</i> packet propagates a route while adjusting its metric and SRI one hop at a time. The SRI that is carried in <i>sec_response</i> packets cannot be propagated through non-Trusted Solaris gateways.</p> <p>The second additional packet type is <i>sec_t_response</i>, which has the exact format as <i>sec_response</i> but with a different command number. The <i>sec_t_response</i> packets are used for tunneling. Every time a <i>response</i> is sent, a <i>sec_response</i> and a <i>sec_t_response</i> packet are also sent.</p> <p>Tunneling can be set up for trusted routing between Trusted Solaris 7 and later gateways when non-Trusted Solaris gateways exist between the Trusted Solaris 7 and later gateways. For tunneling to work, all Trusted Solaris gateways must be running Trusted Solaris 2.5.1 or 7 or 8, and they must be using the extended <code>in.routed(1M)</code> for dynamic routing. Also, the non-Trusted Solaris gateways must be using the standard <code>in.routed(1M)</code> for dynamic routing. All gateways must be in the same Intranet. To forward SRIs through non-Trusted Solaris gateways to a target (sub)network, a Trusted Solaris system sends an unlabeled <i>sec_t_response</i> packet in a (sub)network directed broadcast to the target (sub)network on behalf of the non-Trusted Solaris gateway connected to that (sub)network. Trusted Solaris systems on the (sub)network can use the SRI to configure their routing tables correctly, and Trusted Solaris 7 gateways on that (sub)network can propagate the SRI to other</p>

(sub)networks. A machine that does tunneling is called the forwarding machine; any Trusted Solaris gateway can be a forwarding machine.

Tunneling is enabled by the existence of the file `/etc/security/tsol/tunnel`, and the target (sub)network addresses are obtained from this file. A Trusted Solaris gateway can be responsible for tunneling to more than one (sub)network. The file is composed of a series of lines, each in the following format:

```
broadcast_addr
```

A Trusted Solaris gateway can be responsible for tunneling to more than one (sub)network.

A Trusted Solaris system ignores a *response* packet if it is sent by another Trusted Solaris gateway, because in this case, *sec_response* packets should be used in place of *response* packets. A Trusted Solaris system processes a *response* packet if it is sent by a non-Trusted Solaris gateway. If tunneling is done on behalf of that non-Trusted Solaris gateway, it will process both the *response* packets sent by the non-Trusted Solaris gateway and the *sec_response* packets sent by a remote Trusted Solaris gateway on behalf of the non-Trusted Solaris gateway.

When a *request* packet is received, `in.routed` formulates a reply based on the information maintained in its internal tables. The *response* packet contains a list of known routes, each marked with a “hop count” metric (a count of 16, or greater, is considered “infinite”). The metric associated with each route returned, provides a metric relative to the sender.

sec_response and *sec_t_response* packets are formulated by ANDing the metric of the route with the metric derived from the outgoing interface. Before the *response* packet is sent, a *sec_response* and a *sec_t_response* packet are sent to the same destination with the same metric and additional SRI.

response, *sec_response*, and *request* packets received by `in.routed` are used to update the routing tables if one of the following conditions is satisfied:

- No routing table entry exists for the destination network or host, and the metric indicates the destination is “reachable” (that is, the hop count is not infinite). For *sec_response* and *sec-t_response* packets, a destination is also unreachable if its SRI restricts all possible packets.
- The source host of the packet is the same as the router in the existing routing table entry. That is, updated information is being received from the very internetwork router through which packets for the destination are being routed. The only exception occurs when `in.routed` is supposed to process both the *response* packet from a non-Trusted Solaris gateway and the *sec_response* packet tunneled on behalf of that non-Trusted Solaris gateway. In this situation, if both packets carry routing information for the same route, the SRI from the tunneled *sec_response* packet and the metric from the *response* packet are used.

routed(1M)

- The existing entry in the routing table has not been updated for some time (defined to be 90 seconds) and the route is at least as cost effective as the current route.
- The new route describes a shorter route to the destination than the one currently stored in the routing tables; the metric of the new route is compared against the one stored in the table to decide this.

For *sec_response* and *sec_t_response* packets, the last rule above is changed to compare the SRIs as well as the metrics. One route is better than another if (a) its metric is smaller; and (b) its SRI is more relaxed than or equal to that of the other. Note that when comparing the SRIs of two routes, one route cannot always serve as a substitute for the other. For example, if the SRIs of two routes have different sensitivity labels, one SRI cannot be said to be more restrictive, because they restrict different sensitivity label ranges.

If two routes cannot be compared, both routes are kept in the routing table, because they represent two routes to the same destination although with different security characteristics; and both routes are needed.

When an update is applied, *in.routed* records the change in its internal tables and generates a *sec_response* packet and a *response* packet to all directly connected hosts and networks. *in.routed* waits a short period of time (no more than 30 seconds) before modifying the kernel's routing tables to allow possible unstable situations to settle.

In addition to processing incoming packets, *in.routed* also periodically checks the routing table entries. If an entry has not been updated for 3 minutes, the entry's metric is set to infinity and marked for deletion. Deletions are delayed an additional 60 seconds to insure the invalidation is propagated throughout the internet.

Hosts acting as internetwork routers gratuitously supply their routing tables every 30 seconds to all directly connected hosts and networks.

In addition to the facilities described above, *in.routed* supports the notion of "distant" passive and active gateways. When *in.routed* is started up, it reads the file *gateways* to find gateways which may not be identified using the *SIOCGIFCONF* ioctl. Gateways specified in this manner should be marked passive if they are not expected to exchange routing information, while gateways marked active should be willing to exchange routing information (that is, they should have a *in.routed* process running on the machine). Passive gateways are maintained in the routing tables forever. Information regarding their existence is not included in any routing information transmitted. Active gateways are treated equally to network interfaces. Routing information is distributed to the gateway and if no routing information is received for a period of time, the associated route is deleted.

The gateways is comprised of a series of lines, each in the following format:

```
< net | host> filename1 gateway filename2 metric value < passive | active >
```

The `net` or `host` keyword indicates if the route is to a network or specific host.

filename1 is the name of the destination network or host. This may be a symbolic name located in `networks` or `hosts`, or an Internet address specified in “dot” notation; see `inet(3SOCKET)`.

filename2 is the name or address of the gateway to which messages should be forwarded.

value is a metric indicating the hop count to the destination host or network.

The keyword `passive` or `active` indicates if the gateway should be treated as passive or active (as described above).

For both the passive and active gateway, the SRI of their routes are obtained initially from their remote host template. For an active gateway, further routing information will be exchanged with this machine. If later a `sec_response` packet is received from the active gateway or a `sec_t_response` tunneled on its behalf is received, the initial SRI will be updated. If no `sec_response` packet is ever received for this active gateway, use of the initial SRI is continued. For a passive gateway, no further routing information will be exchanged; therefore, the initial SRI is continuously used.

`in.routed` must be started from the Trusted path at `ADMIN_HIGH`. It must inherit the `net_mac_read`, `net_privaddr`, `net_broadcast`, and `sys_net_config` privileges. If a log file is specified, `in.routed` must also inherit the `file_mac_write` privilege.

OPTIONS

- g Is used on internetwork routers to offer a route to the “default” destination. This is typically used on a gateway to the Internet, or on a gateway that uses another routing protocol whose routes are not reported to other local routers.
- q Is the opposite of the `-s` option.
- s Forces `in.routed` to supply routing information whether it is acting as an internetwork router or not.
- S If `in.routed` is not acting as an internetwork router it will, instead of entering the whole routing table in the kernel, only enter a default route for each internetwork router. This reduces the the memory requirements without losing any routing reliability.
- t All packets sent or received are printed on standard output. In addition, `in.routed` will not divorce itself from the controlling terminal so that interrupts from the keyboard will kill the process. Any other argument supplied is interpreted as the name of the file in which `in.routed`’s actions should be logged. This log contains information about any changes to the routing tables and a history of recent messages sent and received which are related to the changed route.

routed(1M)

	-v	Allows a logfile (whose name must be supplied) to be created showing the changes made to the routing tables with a timestamp.				
FILES	/etc/gateways	For distant gateways				
	/etc/networks	Associations of Internet Protocol network numbers with network names				
	/etc/hosts	Internet host table				
	/etc/security/tsolgateways	For trusted routing through listed gateways				
	/etc/security/tsol/tunnel	Tunneling information table for Trusted Solaris hosts				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:					
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWcsu</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE					
Availability	SUNWcsu					
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>in.routed should be started at ADMIN_HIGH. It must inherit the net_mac_read, net_privaddr, net_broadcast, and sys_net_config privileges. If a log file is specified, in.routed must also inherit the file_mac_write privilege. Because trusted routing considers the security of the route along with the route's metric when making routing decisions, in.routed sends two additional types of response packets containing security information for routes: <i>sec_response</i> packets for communications with connected Trusted Solaris gateways, and <i>sec_t_response</i> packets for tunneling to Trusted Solaris gateways on the other side of non-Trusted Solaris gateways.</p>					
Trusted Solaris 8 4/01 Reference Manual NOTES	route(1M)					
	ioctl(2), inet(3SOCKET), attributes(5), inet(7P), udp(7P)					
	<p>The kernel's routing tables may not correspond to those of in.routed for short periods of time while processes that utilize existing routes exit; the only remedy for this is to place the routing process in the kernel.</p> <p>in.routed should listen to intelligent interfaces, such as an IMP, and to error protocols, such as ICMP, to gather more information.</p> <p>in.routed initially obtains a routing table by examining the interfaces configured on a machine and the gateways file. It then sends a request on all directly connected networks for more routing information. in.routed does not recognize or use any routing information already established on the machine prior to startup. With the exception of interface changes, in.routed does not see any routing table changes that have been done by other programs on the machine, for example, routes added, deleted or flushed by way of the route(1M) command. Therefore, these types of changes should not be done while in.routed is running. Rather, shut down in.routed, make the changes required, and then restart in.routed.</p>					

NAME	rpcbind – Universal addresses to RPC program number mapper
-------------	--

SYNOPSIS **rpcbind** [-d] [-w]

DESCRIPTION	rpcbind is a server that converts RPC program numbers into universal addresses. It must be running on the host to be able to make RPC calls on a server on that machine.
--------------------	--

When an RPC service is started, it tells `rpcbind` the address at which it is listening, and the RPC program numbers it is prepared to serve. When a client wishes to make an RPC call to a given program number, it first contacts `rpcbind` on the server machine to determine the address where RPC requests should be sent.

rpcbind should be started before any other RPC service. Normally, standard RPC servers are started by port monitors, so rpcbind must be started before port monitors are invoked.

When `rpcbind` is started, it checks that certain name-to-address translation-calls function correctly. If they fail, the network configuration databases may be corrupt. Since RPC services cannot function correctly in this situation, `rpcbind` reports the condition and terminates.

rpcbind should be run at a sensitivity label of ADMIN_HIGH with all privileges, and must be run from the trusted path. Note, however, that these privileges are made effective only when required for rpcbind's operation. Most are used only when rpcbind makes an RPC call on behalf of a privileged RPCBPROC CALLIT client.

OPTIONS The following options are supported:

- d Run in debug mode. In this mode, `rpcbind` will not fork when it starts, will print additional information during operation, and will abort on certain errors. With this option, the name-to-address translation consistency checks are shown in detail.

-w Do a warm start. If `rpcbind` aborts or terminates on `SIGINT` or `SIGTERM`, it will write the current list of registered services to `/tmp/portmap.file` and `/tmp/rpcbind.file`. Starting `rpcbind` with the `-w` option instructs it to look for these files and start operation with the registrations found in them. This allows `rpcbind` to resume operation without requiring all RPC services to be restarted.

FILES	/tmp/portmap.file	File of registered services used during a warm start.
--------------	-------------------	---

<code>/tmp/rpcbind.file</code>	File of registered services used during a warm start.
--------------------------------	---

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

rpcbind(1M)

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**

rpcbind should be run at a sensitivity label of ADMIN_HIGH and must be run from the trusted path. rpcbind should be run with all privileges. Note, however, that these privileges are made effective only when required for rpcbind's operation. Most are used only when rpcbind makes an RPC call on behalf of a privileged RPCBPROC_CALLIT client.

**Trusted Solaris 8
4/01 Reference
Manual
NOTES**

rpcinfo(1M), rpcbind(3NSL)

attributes(5)

Terminating rpcbind with SIGKILL will prevent the warm-start files from being written.

All RPC servers must be restarted if the following occurs: rpcbind crashes (or is killed with SIGKILL) and is unable to write the warm-start files; rpcbind is started without the -w option after a graceful termination; or, the warm-start files are not found by rpcbind.

NAME	rpc.bootparamd, bootparamd – Boot parameter server				
SYNOPSIS	/usr/sbin/rpc.bootparamd [-d]				
DESCRIPTION	<p>rpc.bootparamd is a server process that provides information from a bootparams database to diskless clients at boot time. See bootparams(4)</p> <p>The source for the bootparams database is determined by the nsswitch.conf(4) file (on the machine running the rpc.bootparamd process).</p> <p>The rpc.bootparamd program can be invoked either by inetd(1M) or directly from the command line.</p>				
OPTIONS	-d Display debugging information.				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>rpc.bootparamd requires the trust path attribute with a UID of 0, and the sensitivity label ADMIN_LOW.</p> <p>/etc/bootparams Boot parameter database.</p> <p>/etc/nsswitch.conf Configuration file for the name-service switch.</p>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
Trusted Solaris 8 4/01 Reference Manual	inetd(1M), nsswitch.conf(4)				
Notes	bootparams(4), attributes(5)				
NOTES	<p>A diskless client requires service from at least one rpc.bootparamd process running on a server that is on the same IP subnetwork as the diskless client.</p> <p>Some routines that compare hostnames use case-sensitive string comparisons; some do not. If an incoming request fails, verify that the case of the hostname in the file to be parsed matches the case of the hostname called for, and attempt the request again.</p>				

rpc.getpeerinfod(1M)

NAME	rpc.getpeerinfod – Getpeerinfo service daemon
SYNOPSIS	/usr/sbin/rpc.getpeerinfod
DESCRIPTION	rpc.getpeerinfo is an RPC server that returns process attributes for peer processes. It is used to obtain values used for process creation and audit characteristic propagation. The rpc.getpeerinfo daemon is normally started through rc scripts.

NAME	rpcinfo – Report RPC information
SYNOPSIS	<pre> rpcinfo [-m -s] [<i>host</i>] rpcinfo -p [<i>host</i>] rpcinfo -T <i>transport</i> <i>host</i> <i>prognum</i> [<i>versnum</i>] rpcinfo [-M] [-s] [<i>host</i>] rpcinfo -l [-T <i>transport</i>] <i>host</i> <i>prognum</i> <i>versnum</i> rpcinfo [-n <i>portnum</i>] -u <i>host</i> <i>prognum</i> [<i>versnum</i>] rpcinfo [-n <i>portnum</i>] -t <i>host</i> <i>prognum</i> [<i>versnum</i>] rpcinfo -a <i>serv_address</i> -T <i>transport</i> <i>prognum</i> [<i>versnum</i>] rpcinfo -b [-T <i>transport</i>] <i>prognum</i> <i>versnum</i> rpcinfo -d [-T <i>transport</i>] <i>prognum</i> <i>versnum</i> </pre>
DESCRIPTION	<p>rpcinfo makes an RPC call to an RPC server and reports what it finds.</p> <p>In the first synopsis, <i>rpcinfo</i> lists all the registered RPC services with <i>rpcbind</i> on <i>host</i>. If <i>host</i> is not specified, the local host is the default. If <i>-s</i> is used, the information is displayed in a concise format.</p> <p>In the second synopsis, <i>rpcinfo</i> lists all the RPC services registered with <i>rpcbind</i>, version 2. Note that the format of the information is different in the first and the second synopsis. This is because the second synopsis is an older protocol used to collect the information displayed (version 2 of the <i>rpcbind</i> protocol).</p> <p>The third synopsis makes an RPC call to procedure 0 of <i>prognum</i> and <i>versnum</i> on the specified <i>host</i> and reports whether a response was received. <i>transport</i> is the transport which has to be used for contacting the given service. The remote address of the service is obtained by making a call to the remote <i>rpcbind</i>.</p> <p>The fourth synopsis is an extended version of the first. While the default report lists the RPC services that are registered for the user's sensitivity label (including multilevel services), the <i>-M</i> option lists all RPC services that are registered at or below the sensitivity label of the user. If the process has the <i>net_mac_read</i> privilege, the list includes all RPC services. These reports include the same information as that produced by the default report plus a multilevel mapping indicator or the sensitivity label at which the RPC service is registered.</p> <p>The <i>prognum</i> argument is a number that represents an RPC program number (see <i>rpc(4)</i>).</p> <p>If a <i>versnum</i> is specified, <i>rpcinfo</i> attempts to call that version of the specified <i>prognum</i>. Otherwise, <i>rpcinfo</i> attempts to find all the registered version numbers for the specified <i>prognum</i> by calling version 0, which is presumed not to exist; if it does exist, <i>rpcinfo</i> attempts to obtain this information by calling an extremely high</p>

rpcinfo(1M)

version number instead, and attempts to call each registered version. Note that the version number is required for `-b` and `-d` options.

The **EXAMPLES** section describe other ways of using `rpcinfo`.

OPTIONS

- | | |
|------------------------------|---|
| <code>-T transport</code> | Specify the transport on which the service is required. If this option is not specified, <code>rpcinfo</code> uses the transport specified in the <code>NETPATH</code> environment variable, or if that is unset or <code>NULL</code> , the transport in the <code>netconfig(4)</code> database is used. This is a generic option, and can be used in conjunction with other options as shown in the SYNOPSIS . |
| <code>-a serv_address</code> | Use <i>serv_address</i> as the (universal) address for the service on <i>transport</i> to ping procedure 0 of the specified <i>prognum</i> and report whether a response was received. The <code>-T</code> option is required with the <code>-a</code> option. If <i>versnum</i> is not specified, <code>rpcinfo</code> tries to ping all available version numbers for that program number. This option avoids calls to remote <code>rpcbind</code> to find the address of the service. The <i>serv_address</i> is specified in universal address format of the given transport. |
| <code>-b</code> | Make an RPC broadcast to procedure 0 of the specified <i>prognum</i> and <i>versnum</i> and report all hosts that respond. If <i>transport</i> is specified, it broadcasts its request only on the specified transport. If broadcasting is not supported by any transport, an error message is printed. Use of broadcasting requires the <code>net_broadcast</code> privilege. |
| <code>-d</code> | Delete registration for the RPC service of the specified <i>prognum</i> and <i>versnum</i> . If <i>transport</i> is specified, unregister the service on only that transport, otherwise unregister the service on all the transports on which it was registered. Only the owner of a service or a process with the <code>net_setid</code> privilege can delete a registration. The <code>net_mac_read</code> privilege is required to delete a multilevel mapping. The <code>net_privaddr</code> privilege is required to delete a mapping to a transport that uses a privileged address. |
| <code>-l</code> | Display a list of entries with a given <i>prognum</i> and <i>versnum</i> on the specified <i>host</i> . Entries are returned for all transports in the same protocol family as that used to contact the remote <code>rpcbind</code> . |
| <code>-m</code> | Display a table of statistics of <code>rpcbind</code> operations on the given <i>host</i> . The table shows statistics for each version of <code>rpcbind</code> (versions 2, 3 and 4), giving the number of times each procedure was requested and successfully serviced, the number and type of remote call requests that were made, and information about RPC address lookups that were handled. This is useful for monitoring RPC activities on <i>host</i> . |

-M	This extended reporting option lists all RPC services that are registered at or below the sensitivity label of the process. If the process has the <code>net_mac_read</code> privilege, the list includes all RPC services regardless of sensitivity label. These reports include the same information as that produced by the default report plus a multilevel mapping indicator or the sensitivity label of the RPC service. Note that the process will require the <code>sys_trans_label</code> privilege in order to display the names of sensitivity labels not dominated by the process.
-n <i>portnum</i>	Use <i>portnum</i> as the port number for the -t and -u options instead of the port number given by <code>rpcbind</code> . Use of this option avoids a call to the remote <code>rpcbind</code> to find out the address of the service. This option is made obsolete by the -a option.
-p	Probe <code>rpcbind</code> on <i>host</i> using version 2 of the <code>rpcbind</code> protocol, and display a list of all registered RPC programs. If <i>host</i> is not specified, it defaults to the local host. Note that version 2 of the <code>rpcbind</code> protocol was previously known as the portmapper protocol.
-s	Display a concise list of all registered RPC programs on <i>host</i> . If <i>host</i> is not specified, it defaults to the local host.
-t	Make an RPC call to procedure 0 of <i>prognum</i> on the specified <i>host</i> using TCP, and report whether a response was received. This option is made obsolete by the -T option as shown in the third synopsis.
-u	Make an RPC call to procedure 0 of <i>prognum</i> on the specified <i>host</i> using UDP, and report whether a response was received. This option is made obsolete by the -T option as shown in the third synopsis.

EXAMPLES**EXAMPLE 1** RPC services.

To show all of the RPC services registered on the local machine use:

```
example% rpcinfo
```

To show all of the RPC services registered with `rpcbind` on the machine named `klaxon` use:

```
example% rpcinfo klaxon
```

The information displayed by the above commands can be quite lengthy. Use the -s option to display a more concise list:

```
example% rpcinfo -s klaxon
```

programversion	netid(s)	service	owner
----------------	----------	---------	-------

rpcinfo(1M)

EXAMPLE 1 RPC services. (Continued)

100000	2,3,4	tcp,udp,ticlts,ticots,ticotsord	rpcbind	superuser
100008	1	ticotsord,ticots,ticlts,udp,tcp	walld	superuser
100002	2,1	ticotsord,ticots,ticlts,udp,tcp	rusersd	superuser
100001	2,3,4	ticotsord,ticots,tcp,ticlts,udp	rstatd	superuser
100012	1	ticotsord,ticots,ticlts,udp,tcp	sprayd	superuser
100007	3	ticotsord,ticots,ticlts,udp,tcp	ypbind	superuser
100029	1	ticotsord,ticots,ticlts	keyserv	superuser
100078	4	ticotsord,ticots,ticlts	kerbd	superuser
100024	1	ticotsord,ticots,ticlts,udp,tcp	status	superuser
100021	2,1	ticotsord,ticots,ticlts,udp,tcp	nlockmgr	superuser
100020	1	ticotsord,ticots,ticlts,udp,tcp	llockmgr	superuser

To show whether the RPC service with program number *prognum* and version *versnum* is registered on the machine named *klaxon* for the transport TCP use:

```
example% rpcinfo -T tcp klaxon prognum versnum
```

To show all RPC services registered with version 2 of the rpcbind protocol on the local machine use:

```
example% rpcinfo -p
```

To delete the registration for version 1 of the walld (program number 100008) service for all transports use:

```
example# rpcinfo -d 100008 1
```

or

```
example# rpcinfo -d walld 1
```

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

The -M option is added, and it requires privilege. The -b and -d options require privilege. See the OPTIONS definitions for details.

rpcbind(1M), rpc(3NSL)
netconfig(4), rpc(4), attributes(5)

rpcinfo(1M)

rpc.nisd(1M)

NAME	rpc.nisd, nisd – NIS+ service daemon
SYNOPSIS	/usr/sbin/rpc.nisd [-ACDFhluv] [-Y [-B [-t <i>netid</i>]]] [-d <i>dictionary</i>] [-L <i>load</i>] [-S <i>level</i>]
DESCRIPTION	<p>The <code>rpc.nisd</code> daemon is an RPC service that implements the NIS+ service. This daemon must be running on all machines which serve a portion of the NIS+ namespace. A Trusted Solaris system must be the root master in the NIS+ configuration.</p> <p><code>rpc.nisd</code> is usually started from a system startup script. It must be started through a role that has a UID of 0 and run with a sensitivity label of ADMIN_LOW. (For example, the role might be assigned the predefined NIS+ security administration and NIS+ administration profiles.) <code>rpc.nisd</code> must be run from the Trusted Path and inherit the <code>net_mac_read</code>, <code>net_upgrade_sl</code>, and <code>proc_setsl</code> privileges.</p> <p>The <code>-B</code> option causes <code>rpc.nisd</code> to start an auxiliary process, <code>rpc.nisd_resolv</code>, which provides ypserv compatible DNS forwarding for NIS host requests. <code>rpc.nisd_resolv</code> can also be started independently. See <code>rpc.nisd_resolv(1M)</code> for more information on using <code>rpc.nisd_resolv</code> independently.</p>
OPTIONS	<p>-A Authentication verbose mode. The daemon logs all the authentication related activities to <code>syslogd(1M)</code> with LOG_INFO priority.</p> <p>-C Open diagnostic channel on <code>/dev/console</code>.</p> <p>-D Debug mode (don't fork).</p> <p>-F Force the server to do a checkpoint of the database when it starts up. Forced checkpoints may be required when the server is low on disk space. This option removes updates from the transaction log that have propagated to all of the replicas.</p> <p>-h Print list of options.</p> <p>-u Allow updates from non-Trusted Solaris TCB clients.</p> <p>-v Verbose. With this option, the daemon sends a running narration of what it is doing to the syslog daemon (see <code>syslogd(1M)</code>) at LOG_INFO priority. This option is most useful for debugging problems with the service (see also <code>-A</code> option).</p> <p>-Y Put the server into NIS (YP) compatibility mode. When operating in this mode, the NIS+ server will respond to NIS Version 2 requests using the version 2 protocol. Because the YP protocol is not authenticated, only those items that have read access to nobody (the unauthenticated request) will be visible through the V2 protocol. It supports only the standard Version 2 maps in this mode (see <code>-B</code> option and NOTES in <code>ypfiles(4)</code>).</p> <p>-B Provide ypserv compatible DNS forwarding for NIS host requests. The DNS resolving process, <code>rpc.nisd_resolv</code>, is started and controlled by <code>rpc.nisd</code>. This option requires that the <code>/etc/resolv.conf</code> file be setup</p>

for communication with a DNS nameserver. The `nslookup` utility can be used to verify communication with a DNS nameserver. See `resolv.conf(4)` and `nslookup(1M)`.

<code>-t netid</code>	Use <i>netid</i> as the transport for communication between <code>rpc.nisd</code> and <code>rpc.nisd_resolv</code> . The default transport is <code>ticots(7D)</code> (<code>tcp</code> on SunOS 4.x systems).
<code>-d dictionary</code>	Specify an alternate dictionary for the NIS+ database. The primary use of this option is for testing. Note that the string is not interpreted, rather it is simply passed to the <code>db_initialize()</code> function.
<code>-L number</code>	Specify the “load” the NIS+ service is allowed to place on the server. The load is specified in terms of the <i>number</i> of child processes that the server may spawn. This <i>number</i> must be at least 1 for the callback functions to work correctly. The default is 128.
<code>-S level</code>	Set the authorization security level of the service. The argument is a number between 0 and 2. By default, the daemon runs at security level 2.
0	Security level 0 is designed to be used for testing and initial setup of the NIS+ namespace. When running at level 0, the daemon does not enforce any access controls. Any client is allowed to perform any operation, including updates and deletions.
1	At security level 1, the daemon accepts both <code>AUTH_SYS</code> and <code>AUTH_DES</code> credentials for authenticating clients and authorizing them to perform NIS+ operations. This is not a secure mode of operation since <code>AUTH_SYS</code> credentials are easily forged. It should not be used on networks in which any untrusted users may potentially have access.
2	At security level 2, the daemon only accepts authentication using the security mechanisms configured by <code>nisauthconf(1M)</code> . The default security mechanism is <code>AUTH_DES</code> . Security level 2 is the default if the <code>-S</code> option is not used.

EXAMPLES **EXAMPLE 1** Setting up the NIS+ service.

The following example sets up the NIS+ service.

```
example% rpc.nisd
```

rpc.nisd(1M)

EXAMPLE 1 Setting up the NIS+ service. *(Continued)*

EXAMPLE 2 Setting Up NIS+ Service Emulating YP With DNS Forwarding

The following example sets up the NIS+ service, emulating YP with DNS forwarding.

```
example% rpc.nisd -YB
```

**ENVIRONMENT
VARIABLES**

NETPATH The transports that the NIS+ service will use can be limited by setting this environment variable (see `netconfig(4)`).

FILES

`/var/nis/data/parent.object`

This file describes the namespace that is logically above the NIS+ namespace. The most common type of parent object is a DNS object. This object contains contact information for a server of that domain.

`/var/nis/data/root.object`

This file describes the root object of the NIS+ namespace. It is a standard XDR-encoded NIS+ directory object that can be modified by authorized clients using the `nis_modify(3NSL)` interface.

`/etc/init.d/rpc`

Initialization script for NIS+.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWnisu

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**

A Trusted Solaris system must be the root master of the NIS+ configuration. The Trusted Solaris environment adds the `-u`. `rpc.nisd` must be run from the Trusted Path and inherit the `net_mac_read`, `net_upgrade_sl`, and `proc_setsl`. The daemon must be started by a role with a UID of 0 and run with a sensitivity label of `ADMIN_LOW`.

**Trusted Solaris 8
4/01 Reference
Manual**

`nis_cachemgr(1M)`, `nissetup(1M)`, `nslookup(1M)`, `rpc.nisd_resolv(1M)`, `rpc.nispasswd(1M)`, `nis_modify(3NSL)`, `resolv.conf(4)`

**SunOS 5.8
Reference Manual**

`nisauthconf(1M)`, `nisinit(1M)`, `syslogd(1M)`, `netconfig(4)`, `nisfiles(4)`, `attributes(5)`, `ticots(7D)`

NAME	rpc.nisd_resolv, nisd_resolv – NIS+ service daemon				
SYNOPSIS	rpc.nisd_resolv [-v -V] [-F [-C <i>fd</i>]] [-t <i>xx</i>] [-p <i>yy</i>]				
DESCRIPTION	<p>rpc.nisd_resolv is an auxiliary process which provides DNS forwarding service for NIS hosts requests to both ypserve and rpc.nisd that are running in the NIS compatibility mode. It is generally started by invoking rpc.nisd(1M) with the -B option or ypserve(1M) with the -d option. Although it is not recommended, rpc.nisd_resolv can also be started independently with the following options.</p> <p>This command is not supported in the Trusted Solaris environment because ypserve and other NIS(YP) compatibility is unsupported.</p>				
OPTIONS	<p>-F Run in foreground.</p> <p>-C <i>fd</i> Use <i>fd</i> for service xprt (from nisd).</p> <p>-v Verbose. Send output to the syslog daemon.</p> <p>-V Verbose. Send output to stdout.</p> <p>-t <i>xx</i> Use transport <i>xx</i>.</p> <p>-p <i>yy</i> Use transient program# <i>yy</i>.</p>				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>This command is not supported in the Trusted Solaris environment.</p> <p>See attributes(5) for descriptions of the following attributes:</p> <table><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWnisu</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWnisu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWnisu				
Trusted Solaris 8 4/01 Reference Manual	rpc.nisd(1M)				
Notes	nslookup(1M), resolv.conf(4), attributes(5)				
Reference Manual	This command requires that the /etc/resolv.conf file be setup for communication with a DNS nameserver. The nslookup utility can be used to verify communication with a DNS nameserver. See resolv.conf(4) and nslookup(1M).				
NOTES					

rpc.nispasswd(1M)

NAME	rpc.nispasswd, nispasswd – NIS+ password update daemon	
SYNOPSIS	/usr/sbin/rpc.nispasswd [-a <i>attempts</i>] [-c <i>minutes</i>] [-D] [-g] [-v]	
DESCRIPTION	<p>rpc.nispasswd daemon is an ONC+ RPC service that services password update requests from nispasswd(1). It updates password entries in the NIS+ passwd table.</p> <p>rpc.nispasswd is normally started from a system startup script after the NIS+ server (rpc.nisd(1M)) has been started. rpc.nispasswd will determine whether it is running on a machine that is a master server for one or more NIS+ directories. If it discovers that the host is not a master server, then it will promptly exit. It will also determine if rpc.nisd(1M) is running in NIS(YP) compatibility mode (the -Yoption) and will register as yppasswd for NIS(YP) clients as well.</p> <p>ypserv and other NIS (YP) compatibility is not supported.</p> <p>rpc.nispasswd will syslog all failed password update attempts, which will allow an administrator to determine whether someone was trying to "crack" the passwords.</p> <p>rpc.nispasswd has to be run by a superuser.</p>	
OPTIONS	-a <i>attempts</i>	Set the maximum number of attempts allowed to authenticate the caller within a password update request session. Failed attempts are syslogd(1M) and the request is cached by the daemon. After the maximum number of allowed attempts the daemon severs the connection to the client. The default value is set to 3.
	-c <i>minutes</i>	Set the number of minutes a failed password update request should be cached by the daemon. This is the time during which if the daemon receives further password update requests for the same user and authentication of the caller fails, then the daemon will simply not respond. The default value is set to 30 minutes.
	-D	Debug. Run in debugging mode.
	-g	Generate DES credential. By default the DES credential is not generated for a user if who does not have one. By specifying this option, if a user does not have a credential, then one will be generated and stored in the NIS+ cred table.
	-v	Verbose. With this option, the daemon sends a running narration of what it is doing to the syslog daemon. This option is useful for debugging problems.
EXIT STATUS	0	success
	1	an error has occurred.
SUMMARY OF TRUSTED SOLARIS CHANGES	rpc.nispasswd must be run with a UID of 0 and with a sensitivity label of ADMIN_LOW. On startup, rpc.nispasswd must inherit the net_mac_read and net_upgrade_sl privileges. ypserv and other NIS (YP) compatibility is not supported.	

rpc.nispasswd(1M)

FILES /etc/init.d/rpc initialization script for NIS+

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWnisu

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

rpc.nisd(1M), nsswitch.conf(4)
nispasswd(1), passwd(1), syslogd(1M), attributes(5)

rpc.tbootparamd(1M)

NAME	rpc.tbootparamd – Trusted Solaris boot parameter server
SYNOPSIS	<code>/usr/sbin/rpc.tbootparamd</code>
DESCRIPTION	<p>rpc.tbootparamd is a server process that monitors when clients change their state from the booting state to normal state and back.</p> <p>During booting, a diskless client changes from an unlabeled to a labeled machine. When the change occurs, the client sends out an RPC broadcast message informing its server of the change. Upon receipt of the message, the rpc.tbootparamd process running on the server calls <code>chstate()</code> to inform the kernel of the change.</p> <p>rpc.tbootparamd should be started with a uid 0 and a sensitivity label of ADMIN_LOW; and it must inherit the <code>sys_net_config</code> and <code>net_mac_read</code> privileges.</p>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWtsu

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

tbootparam(1M), `chstate(2)`

`attributes(5)`

NAME	rpc.yppasswdd, yppasswdd – Server for modifying NIS password file
SYNOPSIS	<pre> /usr/lib/netsvc/yp/rpc.yppasswdd [-D <i>directory</i>] [-nogecos] [-noshell] [-nopw] [-m <i>argument1 argument2...</i>] [-u] /usr/lib/netsvc/yp/rpc.yppasswdd [<i>passwordfile adjunctfile</i>] [-nogecos] [-noshell] [-nopw] [-m <i>argument1 argument2...</i>] [-u] </pre>
DESCRIPTION	<p>rpc.yppasswdd is a server that handles password change requests from yppasswd(1). It changes a password entry in the passwd, shadow, and security/passwd.adjunct files. The passwd and shadow files provide the basis for the passwd.byname and passwd.byuid maps. The passwd.adjunct file provides the basis for the passwd.adjunct.byname and passwd.adjunct.byuid maps. Entries in the passwd, shadow or passwd.adjunct files are changed only if the password presented by yppasswd(1) matches the encrypted password of the entry. All password files are located in the PWDIR directory. rpc.yppasswdd must be run from the Trusted Path and inherit the net_broadcast, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl privileges.</p> <p>If the -D option is given, the passwd, shadow, or passwd.adjunct files are placed under the directory path that is the argument to -D.</p> <p>If the -noshell, -nogecos or -nopw options are given, these fields cannot be changed remotely using chfn, chsh, or passwd(1).</p> <p>If the -m option is given, a make(1) is performed in /var/yp after any of the passwd, shadow, or passwd.adjunct files are modified. All arguments following the flag are passed to make.</p> <p>If the -u option is given, updates from non-Trusted Solaris TCB clients are allowed.</p> <p>The second of the listed syntaxes is provided only for backward compatibility. If the second syntax is used, the <i>passwordfile</i> is the full pathname of the password file and <i>adjunctfile</i> is the full pathname of the optional passwd.adjunct file. If a shadow file is found in the same directory as <i>passwordfile</i>, the <i>shadowfile</i> is used as described above. Use of this syntax and the discovery of a <i>shadowfile</i> file generates diagnostic output. The daemon, however, starts normally.</p> <p>The first and second syntaxes are mutually exclusive. You cannot specify the full pathname of the passwd, passwd.adjunct files and use the -D option at the same time.</p> <p>The daemon is started automatically on the master server of the passwd map by ypstart(1), which is invoked at boot time by the /etc/init.d/rpc script.</p> <p>The server does not insist on the presence of a shadow file unless there is no -D option present or the directory named with the -D option is /etc. In addition, a passwd.adjunct file is not necessary. If the -D option is given, the server attempts to find a passwd.adjunct file in the security subdirectory of the named directory.</p>

rpc.yppasswdd(1M)

For example, in the presence of “-D /var/yp” the server checks for a “/var/yp/security/passwd.adjunct” file.

If only a passwd file exists, then the encrypted password is expected in the second field. If both a passwd and a passwd.adjunct file exist, the encrypted password is expected in the second field of the adjunct file with ##username in the second field of the passwd file. If all three files are in use, the encrypted password is expected in the shadow file. Any deviation causes a password update to fail.

If you remove or add a shadow or passwd.adjunct file after rpc.yppasswdd has started, you must stop and restart the daemon to enable it to recognize the change. See ypstart(1) for information on restarting the daemon.

SUMMARY OF TRUSTED SOLARIS CHANGES

rpc.yppasswdd must be run from the Trusted Path and inherit the net_broadcast, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl privileges.

If the -u option is given, updates from non-Trusted Solaris TCB clients are allowed.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWypu

Trusted Solaris 8 4/01 Reference Manual NOTES

passwd(1), inetd(1M), shadow(4)

make(1), yppasswd(1), ypmake(1), passwd(4), ypfiles(4), attributes(5)

If make has not been installed and the -m option is given, the daemon outputs a warning and proceeds, effectively ignoring the -m flag.

When using the -D option, you should make sure that the PWDIR of the /var/yp/Makefile is set accordingly.

The second listed syntax is supplied only for backward compatibility and might be removed in a future release of this daemon.

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications PLC, and cannot be used without permission.

NAME	rpc.yppupdated, yppupdated – server for changing NIS information				
SYNOPSIS	/usr/lib/netsvc/yp/rpc.yppupdated [-isu]				
DESCRIPTION	<p>yppupdated is a daemon that updates information in the Network Information Service (NIS). yppupdated consults the updaters(4) file in the /var/yp directory to determine which NIS maps should be updated and how to change them. yppupdated must be run from the Trusted Path and inherit the net_broadcast, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl privileges.</p> <p>By default, the daemon requires the most secure method of authentication available to it, either DES (secure) or UNIX (insecure).</p>				
OPTIONS	<p>-i Accept RPC calls with the insecure AUTH_UNIX credentials. This allows programmatic updating of the NIS maps in all networks.</p> <p>-s Accept only calls authenticated using the secure RPC mechanism (AUTH_DES authentication). This disables programmatic updating of the NIS maps unless the network supports these calls.</p> <p>-u Allow updates from non-Trusted Solaris TCB clients.</p>				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>yppupdated must be run from the Trusted Path and inherit the net_broadcast, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl privileges.</p> <p>-u allows updates from non-Trusted Solaris TCB clients.</p>				
FILES	/var/yp/updaters Configuration file for rpc.updated command.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWypu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWypu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWypu				
SunOS 5.8 Reference Manual	<p>keyserv(1M), updaters(4), attributes(5)</p> <p><i>System Administration Guide, Volume I</i></p>				
NOTES	The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two services remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications PLC, and must not be used without permission.				

rshd(1M)

NAME	in.rshd, rshd – Remote shell server
SYNOPSIS	in.rshd <i>host.port</i>
DESCRIPTION	<p>in.rshd is the server for the rsh(1) program. The server provides remote execution facilities with authentication based on privileged port numbers.</p> <p>in.rshd is invoked by inetd(1M) each time a shell service is requested, and executes the following protocol:</p> <ol style="list-style-type: none">1. The server checks the client's source port. If the port is not in the range 0-1023, the server aborts the connection. The client's host address (in hex) and port number (in decimal) are the arguments passed to in.rshd.2. The server reads characters from the socket up to a null (0) byte. The resultant string is interpreted as an ASCII number, base 10.3. If the number received in step 1 is non-zero, it is interpreted as the port number of a secondary stream to be used for the stderr. A second connection is then created to the specified port on the client's machine. The source port of this second connection is also in the range 0-1023.4. The server checks the client's source address. If the address is associated with a host for which no corresponding entry exists in the host name data base (see hosts(4)), the server aborts the connection. Please refer to the SECURITY section below for more details.5. A null terminated user name of at most 16 characters is retrieved on the initial socket. This user name is interpreted as a user identity to use on the <i>server's</i> machine.6. A null terminated user name of at most 16 characters is retrieved on the initial socket. This user name is interpreted as the user identity on the <i>client's</i> machine.7. A null terminated command to be passed to a shell is retrieved on the initial socket. The length of the command is limited by the upper bound on the size of the system's argument list.8. in.rshd checks whether logins are currently allowed by looking for an /etc/nologin file. If the file exists, the connection is terminated. If logins are allowed, the user is validated according to the following steps. The remote user name is looked up in the password file and a chdir is performed to the user's home directory. If the lookup fails, the connection is terminated. If the chdir fails, it does a chdir to / (root). If the user is not the superuser, (user ID 0), and if the pam_rhosts_auth PAM module is configured for authentication, the file /etc/hosts.equiv is consulted for a list of hosts considered "equivalent". If the client's host name is present in this file, the authentication is considered successful. See the SECURITY section below for a discussion of PAM authentication. <p>If the lookup fails, or the user is root, then the file .rhosts in the home directory of the remote user is checked for the machine name and identity of the user on the client's machine. If this lookup fails, the connection is terminated</p>

9. A null byte is returned on the connection associated with the `stderr` and the command line is passed to the normal login shell of the user. (The `PATH` variable is set to `/usr/bin`.) The shell inherits the network connections established by `in.rshd`.

USAGE `rshd` and `in.rshd` are IPv6-enabled. See `ip6(7P)`.

SECURITY `in.rshd` uses `pam(3PAM)` for authentication, account management, and session management. The PAM configuration policy, listed through `/etc/pam.conf`, specifies the modules to be used for `in.rshd`. Here is a partial `pam.conf` file with entries for the `rsh` command using `rhosts` authentication, UNIX account management, and session management module.

<code>rsh</code>	<code>auth</code>	<code>required</code>	<code>/usr/lib/security/pam_rhosts_auth.so.1</code>
<code>rsh</code>	<code>account</code>	<code>required</code>	<code>/usr/lib/security/pam_unix.so.1</code>
<code>rsh</code>	<code>session</code>	<code>required</code>	<code>/usr/lib/security/pam_unix.so.1</code>

If there are no entries for the `rsh` service, then the entries for the "other" service will be used. To maintain the authentication requirement for `in.rshd`, the `rsh` entry must always be configured with the `pam_rhosts_auth.so.1` module.

SUMMARY OF TRUSTED SOLARIS CHANGES If the `/etc/nologin` file exists, the server will not allow connections. The values of the trusted path, label view, and label-translation process attributes from the client process are propagated to the remote shell.

FILES `/etc/hosts.equiv`

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8 `inetd(1M)`, `inetd.conf(4)`, `pam_unix(5)`

4/01 Reference Manual `rsh(1)`, `pam(3PAM)`, `hosts(4)`, `pam.conf(4)`, `attributes(5)`, `pam_rhosts_auth(5)`, `ip6(7P)`

DIAGNOSTICS The following diagnostic messages are returned on the connection associated with `stderr`, after which any network connections are closed. An error is indicated by a leading byte with a value of 1 in step 9 above (0 is returned above upon successful completion of all the steps prior to the command execution).

`locuser too long`

The name of the user on the client's machine is longer than 16 characters.

rshd(1M)

remuser too long

The name of the user on the remote machine is longer than 16 characters.

command too long

The command line passed exceeds the size of the argument list (as configured into the system).

Hostname for your address unknown.

No entry in the host name database existed for the client's machine.

Login incorrect.

No password file entry for the user name existed.

Permission denied.

The authentication procedure described above failed.

Can't make pipe.

The pipe needed for the stderr was not created.

Try again.

A fork by the server failed.

NOTES

The authentication procedure used here assumes the integrity of each client machine and the connecting medium. This is insecure, but it is useful in an "open" environment.

A facility to allow all data exchanges to be encrypted should be present.

NAME	runpd – Run a command for privilege debugging				
SYNOPSIS	/usr/sbin/runpd [-p] [-a -f] <i>command</i> [<i>args</i>]				
DESCRIPTION	<p>The runpd command is a debugging utility intended for use by administrators and developers. runpd turns on the priv_debug process attribute and executes the program specified by <i>command</i>. The <i>command</i> process inherits the priv_debug process attribute from runpd, and privilege-checking logs are generated for it. The logs list privileges that <i>command</i> needed to succeed, but lacked. <i>args</i> is the optional set of arguments passed as input to <i>command</i>.</p> <p>runpd must be invoked from the Trusted Path.</p> <p>To enable privilege debugging with runpd, the tsol_privs_debug kernel variable in /etc/system must be set to 1, and entries for kern.debug, daemon.debug, and local0.debug must be uncommented in the /etc/syslog.conf file, as in:</p> <pre>kern.debug;daemon.debug;local0.debug /var/log/privdebug.log</pre> <p>The string kern.debug enables privilege debugging of an application's use of system calls. The local0.debug and daemon.debug strings enable debugging of privileges interpreted by system daemons (for example, the sys_trans_label privilege and X window calls). Multiple strings are separated by semicolons.</p>				
OPTIONS	<p>-p Execute <i>command</i> with the trusted_path process attribute. This option is useful when testing a program (<i>command</i>) that requires the attribute.</p> <p>-a The log will include all privilege debugging records from this and previous executions of runpd.</p> <p>-f The log will include any privilege debugging records generated by <i>command</i> or its descendants. runpd looks for all process IDs that are greater than or equal to that of <i>command</i>. Since process IDs can wrap and child processes may not terminate before <i>command</i> terminates, some entries may not be displayed. Use -a to display all records.</p>				
EXIT STATUS	runpd returns the exit code it receives from <i>command</i> .				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWtsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWtsu				
Trusted Solaris 8 4/01 Reference Manual	pattr(1)				
SUNWtsu Reference Manual	syslog.conf(4), system(4), attributes(5)				
NOTES	These interfaces are uncommitted. Although they are not expected to change between minor releases of the Trusted Solaris environment, they may.				

rwall(1M)

NAME	rwall – Write to all users over a network				
SYNOPSIS	<pre>/usr/sbin/rwall hostname... /usr/sbin/rwall -n netgroup... /usr/sbin/rwall -h hostname -n netgroup</pre>				
DESCRIPTION	<p>rwall reads a message from standard input until EOF. It then sends this message, preceded by the line:</p> <p>Broadcast Message . . . to all users logged in on the specified host machines. With the -n option, it sends to the specified network groups.</p>				
OPTIONS	<p>-n <i>netgroup</i> Send the broadcast message to the specified network groups.</p> <p>-h <i>hostname</i> Specify the hostname, the name of the host machine.</p>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWcsu</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>When it is used to send messages to broadcast addresses rather than to specific hosts, this program needs to inherit the net_broadcast privilege to run properly.</p>				
Trusted Solaris 8 4/01 Reference Manual	<p>inetd(1M)</p> <p>listen(1M), pmadm(1M), sacadm(1M), wall(1M), attributes(5)</p>				
NOTES	<p>The timeout is fairly short to allow transmission to a large group of machines (some of which may be down) in a reasonable amount of time. Thus the message may not get through to a heavily loaded machine.</p>				

NAME	sendmail – send mail over the internet						
SYNOPSIS	<pre> /usr/lib/sendmail [-ba] [-bD] [-bd] [-bi] [-bm] [-bp] [-bs] [-bt] [-bv] [-B type] [-C file] [-d X] [-F fullname] [-f name] [-h N] [-M xvalue] [-Nnotifications] [-n] [-Ooption =value] [-o xvalue] [-p protocol] [-q [time]] [-q Xstring] [-R ret] [-r name] [-t] [-V envid] [-v] [-X logfile] [address...] </pre>						
DESCRIPTION	<p>sendmail sends a message to one or more people, routing the message over whatever networks are necessary. sendmail does internetwork forwarding as necessary to deliver the message to the correct place.</p> <p>sendmail is not intended as a user interface routine; other programs provide user-friendly front ends. sendmail is used only to deliver pre-formatted messages.</p> <p>With no flags, sendmail reads its standard input up to an EOF, or a line with a single dot, and sends a copy of the letter found there to all of the addresses listed. It determines the network to use based on the syntax and contents of the addresses.</p> <p>Local addresses are looked up in the local <code>aliases(4)</code> file, or in a name service as defined by the <code>nsswitch.conf(4)</code> file, and aliased appropriately. In addition, if there is a <code>.forward</code> file in a recipient's home directory, sendmail forwards a copy of each message to the list of recipients that file contains. Refer to the NOTES section for more information about <code>.forward</code> files. Aliasing can be prevented by preceding the address with a backslash. Normally the sender is not included in alias expansions. For example, if "john" sends to "group", and "group" includes "john" in the expansion, then the message will not be delivered to "john". See the <code>MeToo</code> Processing Option for more information.</p> <p>There are several conditions under which the expected behavior is for the alias database to be either built or rebuilt. It is important to note that this cannot occur under any circumstances unless root owns <i>and</i> has exclusive write permission to the <code>/etc/mail/aliases*</code> files.</p> <p>If a message is found to be undeliverable, it is returned to the sender with diagnostics that indicate the location and nature of the failure; or, the message is placed in a <code>dead.letter</code> file in the sender's home directory.</p>						
OPTIONS	<table> <tr> <td>-ba</td><td>Go into ARPANET mode. All input lines must end with a RETURN-LINEFEED, and all messages will be generated with a RETURN-LINEFEED at the end. Also, the From: and Sender: fields are examined for the name of the sender.</td></tr> <tr> <td>-bd</td><td>Run as a daemon in the background, waiting for incoming SMTP connections.</td></tr> <tr> <td>-bD</td><td>Run as a daemon in the foreground, waiting for incoming SMTP connections.</td></tr> </table>	-ba	Go into ARPANET mode. All input lines must end with a RETURN-LINEFEED, and all messages will be generated with a RETURN-LINEFEED at the end. Also, the From: and Sender: fields are examined for the name of the sender.	-bd	Run as a daemon in the background, waiting for incoming SMTP connections.	-bD	Run as a daemon in the foreground, waiting for incoming SMTP connections.
-ba	Go into ARPANET mode. All input lines must end with a RETURN-LINEFEED, and all messages will be generated with a RETURN-LINEFEED at the end. Also, the From: and Sender: fields are examined for the name of the sender.						
-bd	Run as a daemon in the background, waiting for incoming SMTP connections.						
-bD	Run as a daemon in the foreground, waiting for incoming SMTP connections.						

sendmail(1M)

-bi	Initialize the <code>aliases(4)</code> database. Root must own <i>and</i> have exclusive write permission to the <code>/etc/mail/aliases*</code> files for successful use of this option.
-bm	Deliver mail in the usual way (default).
-bp	Print a summary of the mail queue.
-bs	Use the SMTP protocol as described in RFC 821. This flag implies all the operations of the <code>-ba</code> flag that are compatible with SMTP.
-bt	Run in address test mode. This mode reads addresses and shows the steps in parsing; it is used for debugging configuration tables.
-bv	Verify names only; do not try to collect or deliver a message. Verify mode is normally used for validating users or mailing lists.
-B <i>type</i>	Indicate body <i>type</i> (7BIT or 8BITMIME).
-C <i>file</i>	Use alternate configuration file.
-d <i>X</i>	Set debugging value to <i>X</i> .
-F <i>fullname</i>	Set the full name of the sender.
-f <i>name</i>	Sets the name of the "from" person (that is, the sender of the mail).
-h <i>N</i>	Set the hop count to <i>N</i> . The hop count is incremented every time the mail is processed. When it reaches a limit, the mail is returned with an error message, the victim of an aliasing loop.
-M <i>xvalue</i>	Set macro <i>x</i> to the specified <i>value</i> .
-n	Do not do aliasing.
-N <i>notifications</i>	Tag all addresses being sent as wanting the indicated <i>notifications</i> , which consists of the word "NEVER" or a comma-separated list of "SUCCESS", "FAILURE", and "DELAY" for successful delivery, failure and a message that is stuck in a queue somewhere. The default is "FAILURE, DELAY".
-o <i>xvalue</i>	Set option <i>x</i> to the specified <i>value</i> . Processing Options are described below.
-O <i>option=value</i>	Set <i>option</i> to the specified <i>value</i> (for long from names). Processing Options are described below.
-p <i>protocol</i>	Set the sending protocol. The <i>protocol</i> field can be in form <i>protocol: host</i> to set both the sending protocol and the sending host. For example: <code>-pUUCP:uunet</code> sets the sending <i>protocol</i> to UUCP and the sending host to uunet. (Some existing programs use <code>-oM</code> to set the <i>r</i> and <i>s</i> macros; this is equivalent to using <code>-p</code>).
-q[<i>time</i>]	Process saved messages in the queue at given intervals. If <i>time</i> is omitted, process the queue once. <i>time</i> is given as a tagged number, with <i>s</i> being seconds, <i>m</i> being minutes, <i>h</i> being hours, <i>d</i> being

	<p>days, and <i>w</i> being weeks. For example, <code>-q1h30m</code> or <code>-q90m</code> would both set the timeout to one hour thirty minutes.</p>
<code>-q Xstring</code>	<p>Run the queue once, limiting the jobs to those matching <i>Xstring</i>. The key letter <i>X</i> can be:</p> <p>I to limit based on queue identifier.</p> <p>R to limit based on recipient.</p> <p>S to limit based on sender.</p> <p>A particular queued job is accepted if one of the corresponding addresses contains the indicated <i>string</i>.</p>
<code>-r name</code>	An alternate and obsolete form of the <code>-f</code> flag.
<code>-R ret</code>	Identify the information you want returned if the message bounces; <i>ret</i> can be "HDRS" for headers only or "FULL" for headers plus body.
<code>-t</code>	Read message for recipients. <code>To:</code> , <code>Cc:</code> , and <code>Bcc:</code> lines will be scanned for people to send to. The <code>Bcc:</code> line will be deleted before transmission. Any addresses in the argument list will be suppressed. The NoRecipientAction Processing Option can be used to change the behaviour when no legal recipients are included in the message.
<code>-v</code>	Go into verbose mode. Alias expansions will be announced, and so forth.
<code>-V envid</code>	The indicated <i>envid</i> is passed with the envelope of the message and returned if the message bounces.
<code>-X logfile</code>	Log all traffic in and out of sendmail in the indicated <i>logfile</i> for debugging mailer problems. This produces a lot of data very quickly and should be used sparingly.
Processing Options	<p>There are a number of "random" options that can be set from a configuration file. Options are represented by a single character or by multiple character names. The syntax for the single character names of is:</p> <p><code>-Oxvalue</code></p> <p>This sets option <i>x</i> to be <i>value</i>. Depending on the option, <i>value</i> may be a string, an integer, a boolean (with legal values <code>t</code>, <code>T</code>, <code>f</code>, or <code>F</code>; the default is <code>TRUE</code>), or a time interval.</p> <p>The multiple character or long names use this syntax:</p> <p><code>-O Longname=argument</code></p>

sendmail(1M)

This sets the option *Longname* to be *argument*. The long names are beneficial because they are easier to interpret than the single character names.

Not all processing options have single character names associated with them. In the list below the multiple character name is presented first followed by the single character syntax enclosed in parentheses.

AliasFile (*Afile*)

Specify possible alias file(s).

AliasWait (*a N*)

If set, wait up to *N* minutes for an "@:@" entry to exist in the `aliases(4)` database before starting up. If it does not appear in *N* minutes, rebuild the database (if the `AutoRebuildAliases` option is also set) or issue a warning. Defaults to 10 minutes.

AllowBogusHELO

Allow a HELO SMTP command that does not include a host name. By default this option is disabled.

AutoRebuildAliases (*D*)

If set, rebuild the `/etc/mail/aliases` database if necessary and possible. If this option is not set, `sendmail` will never rebuild the `aliases` database unless explicitly requested using `-bi`, or `newaliases(1)` is invoked. Note that in order for the database to be rebuilt, root must own *and* have exclusive write permission to the `/etc/mail/aliases*` files. This option is not supported in the Trusted Solaris environment.

BlankSub (*Bc*)

Set the blank substitution character to *c*. Unquoted spaces in addresses are replaced by this character. Defaults to `SPACE` (that is, no change is made).

CheckAliases (*n*)

Validate the RHS of aliases when rebuilding the `aliases(4)` database.

CheckpointInterval (*CN*)

Checkpoints the queue every *N* (default 10) addresses sent. If your system crashes during delivery to a large list, this prevents retransmission to any but the last *N* recipients.

ClassFactor (*zfact*)

The indicated factor *fact* is multiplied by the message class (determined by the `Precedence:` field in the user header and the `P` lines in the configuration file) and subtracted from the priority. Thus, messages with a higher `Priority:` will be favored. Defaults to 1800.

ColonOkInAddr

If set, colons are treated as a regular character in addresses. If not set, they are treated as the introducer to the RFC 822 "group" syntax. This option is on for version 5 and lower configuration files.

ConnectionCacheSize (*kN*)

The maximum number of open connections that will be cached at a time. The default is 1. This delays closing the current connection until either this invocation of `sendmail` needs to connect to another host or it terminates. Setting it to 0 defaults to the old behavior, that is, connections are closed immediately.

ConnectionCacheTimeout (*Ktimeout*)

The maximum amount of time a cached connection will be permitted to idle without activity. If this time is exceeded, the connection is immediately closed. This value should be small (on the order of ten minutes). Before `sendmail` uses a cached connection, it always sends a NOOP (no operation) command to check the connection; if this fails, it reopens the connection. This keeps your end from failing if the other end times out. The point of this option is to be a good network neighbor and avoid using up excessive resources on the other end. The default is five minutes.

ConnectionRateThrottle

The maximum number of connections permitted per second. After this many connections are accepted, further connections will be delayed. If not set or ≤ 0 , there is no limit.

DaemonPortOptions (*Options*)

Set server SMTP options. The options are *key=value* pairs. Known keys are:

Addr	Address mask (defaults INADDR_ANY)
	The address mask may be a numeric address in dot notation or a network name.
Family	Address family (defaults to INET)
Listen	Size of listen queue (defaults to 10)
Port	Name/number of listening port (defaults to smtp)
ReceiveSize	The size of the TCP/IP receive buffer.
SendSize	The size of the TCP/IP send buffer.

DefaultCharSet

Set the default character set to use when converting unlabeled 8 bit input to MIME.

DefaultUser (*gid*) or (*uid*)

Set the default group ID for mailers to run in to *gid* or set the default userid for mailers to *uid*. Defaults to 1. The value can also be given as a symbolic group or user name.

DeliveryMode (*dx*)

Deliver in mode *x*. Legal modes are:

- i Deliver interactively (synchronously).
- b Deliver in background (asynchronously).

sendmail(1M)

<code>d</code>	Deferred mode — database lookups are deferred until the actual queue run.
<code>q</code>	Just queue the message (deliver during queue run).
Defaults to <code>b</code> if no option is specified, <code>i</code> if it is specified but given no argument (that is, <code>Od</code> is equivalent to <code>Odi</code>).	
DialDelay If a connection fails, wait this many seconds and try again. Zero means “do not retry”.	
DontBlameSendmail If set, override the file safety checks. This compromises system security and should not be used. See http://www.sendmail.org/tips/DontBlameSendmail.html for more information.	
DontExpandCnames If set, <code>\$(... \$)</code> lookups that do DNS-based lookups do not expand CNAME records.	
DontInitGroups If set, the <code>initgroups(3C)</code> routine will never be invoked. If you set this, agents run on behalf of users will only have their primary (<code>/etc/passwd</code>) group permissions.	
DontProbeInterfaces If set, <code>sendmail</code> will not insert the names and addresses of any local interfaces into the <code>\$(=w</code> class. If set, you must also include support for these addresses, otherwise mail to addresses in this list will bounce with a configuration error.	
DontPruneRoutes (R) If set, do not prune route-addr syntax addresses to the minimum possible.	
DoubleBounceAddress If an error occurs when sending an error message, send that “double bounce” error message to this address.	
EightBitMode (8) Use 8-bit data handling. This option requires one of the following keys. The key can be selected by using just the first character, but using the full word is better for clarity.	
<code>mimify</code>	Do any necessary conversion of 8BITMIME to 7-bit.
<code>pass</code>	Pass unlabeled 8-bit input through as is.
<code>strict</code>	Reject unlabeled 8-bit input.
ErrorHeader (Efile/message) Append error messages with the indicated message. If it begins with a slash, it is assumed to be the pathname of a file containing a message (this is the recommended setting). Otherwise, it is a literal message. The error file might contain the name, email address, and/or phone number of a local postmaster who	

could provide assistance to end users. If the option is missing or `NULL`, or if it names a file which does not exist or which is not readable, no message is printed.

ErrorMode (*ex*)

Dispose of errors using mode *x*. The values for *x* are:

- e** Mail back errors and give 0 exit status always.
- m** Mail back errors.
- p** Print error messages (default).
- q** No messages, just give exit status.
- w** Write back errors (mail if user not logged in).

FallbackMXhost (*Vfallbackhost*)

If specified, the *fallbackhost* acts like a very low priority MX on every host. This is intended to be used by sites with poor network connectivity.

ForkEachJob (*Y*)

If set, deliver each job that is run from the queue in a separate process. Use this option if you are short of memory, since the default tends to consume considerable amounts of memory while the queue is being processed.

ForwardPath (*Jpath*)

Set the path for searching for users' `.forward` files. The default is `$z/.forward`. Some sites that use the automounter may prefer to change this to `/var/forward/$u` to search a file with the same name as the user in a system directory. It can also be set to a sequence of paths separated by colons; `sendmail` stops at the first file it can successfully and safely open. For example, `/var/forward/$u:$z/.forward` will search first in `/var/forward/username` and then in `~username/.forward` (but only if the first file does not exist). Refer to the **NOTES** section for more information.

HelpFile (*Hfile*)

Specify the help file for SMTP.

HoldExpensive (*c*)

If an outgoing mailer is marked as being expensive, don't connect immediately.

HostsFile

Set the file to use when doing "file" type access of host names.

HostStatusDirectory

If set, host status is kept on disk between `sendmail` runs in the named directory tree. If a full path is not used, then the path is interpreted relative to the queue directory.

IgnoreDots (*i*)

Ignore dots in incoming messages. This is always disabled (that is, dots are always accepted) when reading SMTP mail.

sendmail(1M)

LabelAdminLow

Specifies what to do with mail sent at ADMIN_LOW when that label is below the recipient's minimum label. Choices for the argument are *return* to the sender, *upgrade* to the recipient's minimum label, or *accept* at the mail's label. If not set, mail sent at the ADMIN_LOW label is upgraded and delivered. (See also LabelTooLow.)

LabelTooLow

Specifies what to do with mail sent at any label below a user's minimum label except for ADMIN_LOW when the mail's label is below the recipient's minimum label. Choices for the argument are *return* to the sender, *upgrade* to the recipient's minimum label, or *accept* at the mail's label. If not set, mail sent at a label below the user's minimum label other than ADMIN_LOW is returned. (See also LabelAdminLow.)

LogLevel (*Ln*)

Set the default log level to *n*. Defaults to 9.

(*Mx value*)

Set the macro *x* to *value*. This is intended only for use from the command line.

MatchGECOS (*G*)

Try to match recipient names using the GECOS field. This allows for mail to be delivered using names defined in the GECOS field in */etc/passwd* as well as the login name.

MaxDaemonChildren

The maximum number of children the daemon will permit. After this number, connections are rejected. If not set or ≤ 0 , there is no limit.

MaxHopCount (*hN*)

The maximum hop count. Messages that have been processed more than *N* times are assumed to be in a loop and are rejected. Defaults to 25.

MaxMessageSize

The maximum size of messages that will be accepted (in bytes).

MaxMimeHeaderLength=*M* [*/N*]

Sets the maximum length of certain MIME header field values to *M* characters. For some of these headers which take parameters, the maximum length of each parameter is set to *N* if specified. If */N* is not specified, one half of *M* will be used. By default, these values are 0, meaning no checks are done.

MaxQueueRunSize

If set, limit the maximum size of any given queue run to this number of entries. This stops reading the queue directory after this number of entries is reached; job priority is not used. If not set, there is no limit.

MeToo (*M*)

Send to me too, even if I am in an alias expansion.

MaxRecipientsPerMessage

If set, allow no more than the specified number of recipients in an SMTP envelope. Further recipients receive a 452 error code and are deferred for the next delivery attempt.

MinFreeBlocks (bN/M)

Insist on at least *N* blocks free on the file system that holds the queue files before accepting email via SMTP. If there is insufficient space, `sendmail` gives a 452 response to the `MAIL` command. This invites the sender to try again later. The optional *M* is a maximum message size advertised in the ESMTP EHLO response. It is currently otherwise unused.

MinQueueAge

The amount of time a job must sit in the queue between queue runs. This allows you to set the queue run interval low for better responsiveness without trying all jobs in each run. The default value is 0.

MustQuoteChars

Characters to be quoted in a full name phrase. `&`, `;`, `\`, `()`, `[]` are quoted automatically.

NoRecipientAction

Set action if there are no legal recipient files in the message. The legal values are:

<code>add-apparently-to</code>	Add an <code>Apparently-to:</code> header with all the known recipients (which may expose blind recipients).
<code>add-bcc</code>	Add an empty <code>Bcc:</code> header.
<code>add-to</code>	Add a <code>To:</code> header with all the known recipients (which may expose blind recipients).
<code>add-to-undisclosed</code>	Add a <code>To: undisclosed-recipients:</code> header.
<code>none</code>	Do nothing, leave the message as it is.

OldStyleHeaders (o)

Assume that the headers may be in old format, that is, spaces delimit names. This actually turns on an adaptive algorithm: if any recipient address contains a comma, parenthesis, or angle bracket, it will be assumed that commas already exist. If this flag is not on, only commas delimit names. Headers are always output with commas between the names.

OperatorChars or \$o

Defines the list of characters that can be used to separate the components of an address into tokens.

PostmasterCopy (Ppostmaster)

If set, copies of error messages will be sent to the named *postmaster*. Only the header of the failed message is sent. Since most errors are user problems, this is probably not a good idea on large sites, and arguably contains all sorts of privacy violations, but it seems to be popular with certain operating systems vendors.

sendmail(1M)

PrivacyOptions (*popt,opt,...*)

Set privacy options. Privacy is really a misnomer; many of these are just a way of insisting on stricter adherence to the SMTP protocol.

The goaway pseudo-flag sets all flags except `restrictmailq` and `restrictqrun`. If `mailq` is restricted, only people in the same group as the queue directory can print the queue. If queue runs are restricted, only root and the owner of the queue directory can run the queue. `authwarnings` add warnings about various conditions that may indicate attempts to spoof the mail system, such as using a non-standard queue directory.

The options can be selected from:

<code>authwarnings</code>	Put X-Authentication-Warning: headers in messages.
<code>goaway</code>	Disallow essentially all SMTP status queries.
<code>needexpnhelo</code>	Insist on HELO or EHLO command before EXPN.
<code>needmailhelo</code>	Insist on HELO or EHLO command before MAIL.
<code>needvrfyhelo</code>	Insist on HELO or EHLO command before VRFY.
<code>noetrn</code>	Disallow ETRN entirely.
<code>noexpn</code>	Disallow EXPN entirely.
<code>noreceipts</code>	Prevent return receipts.
<code>novrfy</code>	Disallow VRFY entirely.
<code>public</code>	Allow open access.
<code>restrictmailq</code>	Restrict <code>mailq</code> command.
<code>restrictqrun</code>	Restrict <code>-q</code> command line flag.

QueueDirectory (*Qdir*)

Use the named *dir* as the queue directory.

QueueFactor (*qfactor*)

Use *factor* as the multiplier in the map function to decide when to just queue up jobs rather than run them. This value is divided by the difference between the current load average and the load average limit (`xflag`) to determine the maximum message priority that will be sent. Defaults to 600000.

QueueLA (*xLA*)

When the system load average exceeds *LA*, just queue messages (that is, do not try to send them). Defaults to 8.

QueueSortOrder

Select the queue sort algorithm. The default value is `Priority`. Other values are `Host` or `Time`.

QueueTimeout (*Trtime/wtime*)

Set the queue timeout to *rtime*. After this interval, messages that have not been successfully sent will be returned to the sender. Defaults to five days (5d). The optional *wtime* is the time after which a warning message is sent. If it is missing or 0, then no warning messages are sent.

RecipientFactor (*yfact*)

The indicated factor *fact* is added to the priority (thus *lowering* the priority of the job) for each recipient, that is, this value penalizes jobs with large numbers of recipients. Defaults to 30000.

RefuseLA (*XLA*)

When the system load average exceeds *LA*, refuse incoming SMTP connections. Defaults to 12.

RemoteMode (> [*RemoteMboxHost*])

If *RemoteMboxHost* is specified, then *remote-mode* is enabled using this host. If *RemoteMboxHost* is not specified, and if /var/mail is remotely mounted, then *remote-mode* is enabled using the remote mount host. If *RemoteMboxHost* is not specified and /var/mail is locally mounted, then *remote-mode* is disabled.

When *remote-mode* is enabled, all outgoing messages are sent through that server.

ResolverOptions (*I*)

Tune DNS lookups.

RetryFactor (*Zfact*)

The indicated factor *fact* is added to the priority every time a job is processed. Thus, each time a job is processed, its priority will be decreased by the indicated value. In most environments this should be positive, since hosts that are down are all too often down for a long time. Defaults to 90000.

RunAsUser

If set, become this user when reading and delivering mail. Intended for use of firewalls where users do not have accounts.

SafeFileEnvironment

If set, sendmail will do a chroot into this directory before writing files.

SaveFromLine (*f*)

Save Unix-style From lines at the front of headers. Normally they are assumed redundant and discarded.

SendMimeErrors (*j*)

If set, send error messages in MIME format (see RFC 1341 and RFC 1344 for details).

ServiceSwitchFile

Defines the path to the service-switch file. Since the service-switch file is defined in the Solaris operating environment this option is ignored.

sendmail(1M)

SevenBitInput (7)

Strip input to seven bits for compatibility with old systems. This should not be necessary.

SingleLineFromHeader

If set, From: lines that have embedded newlines are unwrapped onto one line.

SingleThreadDelivery

If this option and the HostStatusDirectory option are both set, use single thread deliveries to other hosts.

SmtgGreetingMessage or \$e

The initial SMTP greeting message.

StatusFile (*Sfile*)

Log statistics in the named file.

SuperSafe (s)

Be super-safe when running things, that is, always instantiate the queue file, even if you are going to attempt immediate delivery. sendmail always instantiates the queue file before returning control to the client under any circumstances.

TempFileMode (*Fmode*)

The file mode for queue files.

Timeout (*rtimeouts*)

Timeout reads after *time* interval. The *timeouts* argument is a list of *keyword=value* pairs. All but *command* apply to client SMTP. For backward compatibility, a timeout with no *keyword=* part will set all of the longer values. The recognized timeouts and their default values, and their minimum values specified in RFC 1123 section 5.3.2 are:

command

command read [1h, 5m]

connect

initial connect [0, unspecified]

datablock

data block read [1h, 3m]

datafinal

reply to final "." in data [1h, 10m]

datainit

reply to DATA command [5m, 2m]

fileopen

file open [60sec, none]

helo

reply to HELO or EHLO command [5m, none]

`hoststatus`
 `host retry` [30m, unspecified]

`iconnect`
 first attempt to connect to a host [0, unspecified]

`ident`
 IDENT protocol timeout [30s, none]

`initial`
 wait for initial greeting message [5m, 5m]

`mail`
 reply to MAIL command [10m, 5m]

`misc`
 reply to NOOP and VERB commands [2m, none]

`queuereturn`
 undeliverable message returned [5d]

`queuwarn`
 deferred warning [4h]

`quit`
 reply to QUIT command [2m, none]

`rcpt`
 reply to RCPT command [1h, 5m]

`rset`
 reply to RSET command [5m, none]

TimeZoneSpec (*tzinfo*)

Set the local time zone info to *tzinfo*, for example, "PST8PDT". Actually, if this is not set, the TZ environment variable is cleared (so the system default is used); if set but null, the user's TZ variable is used, and if set and non-null, the TZ variable is set to this value.

TryNullMXList (*w*)

If you are the "best" (that is, lowest preference) MX for a given host, you should normally detect this situation and treat that condition specially, by forwarding the mail to a UUCP feed, treating it as local, or whatever. However, in some cases (such as Internet firewalls) you may want to try to connect directly to that host as though it had no MX records at all. Setting this option causes `sendmail` to try this. The downside is that errors in your configuration are likely to be diagnosed as "host unknown" or "message timed out" instead of something more meaningful. This option is deprecated.

UnixFromLine or `$1`

The "From " line used when sending to files or programs.

sendmail(1M)

	<p>UnsafeGroupWrites If set, group-writable <code>:include:</code> and <code>.forward</code> files are considered “unsafe”, that is, programs and files cannot be directly referenced from such files.</p> <p>UseErrorsTo (l) If there is an <code>Errors-To:</code> header, send error messages to the addresses listed there. They normally go to the envelope sender. Use of this option causes <code>sendmail</code> to violate RFC 1123.</p> <p>UserDatabaseSpec (U) Defines the name and location of the file containing User Database information.</p> <p>Verbose (v) Run in verbose mode. If this is set, <code>sendmail</code> adjusts the <code>HoldExpensive</code> and <code>DeliveryMode</code> options so that all mail is delivered completely in a single job so that you can see the entire delivery process. The <code>Verbose</code> option should <i>never</i> be set in the configuration file; it is intended for command line use only.</p> <p>All options can be specified on the command line using the <code>-o</code> flag, but most will cause <code>sendmail</code> to relinquish its <code>setuid</code> permissions. The options that will not cause this are <code>b</code>, <code>d</code>, <code>e</code>, <code>E</code>, <code>i</code>, <code>L</code>, <code>m</code>, <code>o</code>, <code>p</code>, <code>r</code>, <code>s</code>, <code>v</code>, <code>C</code>, and <code>7</code>. Also considered “safe” is <code>M</code> (define macro) when defining the <code>r</code> or <code>s</code> macros.</p> <p>If the first character of the user name is a vertical bar, the rest of the user name is used as the name of a program to pipe the mail to. It may be necessary to quote the name of the user to keep <code>sendmail</code> from suppressing the blanks from between arguments.</p> <p>If invoked as <code>newaliases</code>, <code>sendmail</code> rebuilds the alias database, so long as the <code>/etc/mail/aliases*</code> files are owned by <code>root</code> <i>and</i> <code>root</code> has exclusive write permission. If invoked as <code>mailq</code>, <code>sendmail</code> prints the contents of the mail queue.</p>														
OPERANDS	<p><i>address</i> address of an intended recipient of the message being sent.</p>														
USAGE	<p>See <code>largefile(5)</code> for the description of the behavior of <code>sendmail</code> when encountering files greater than or equal to 2 Gbyte (2³¹ bytes).</p>														
EXIT STATUS	<p><code>sendmail</code> returns an exit status describing what it did. The codes are defined in <code></usr/include/sys/exits.h></code>.</p> <table><tr><td><code>EX_OK</code></td><td>Successful completion on all addresses.</td></tr><tr><td><code>EX_NOUSER</code></td><td>User name not recognized.</td></tr><tr><td><code>EX_UNAVAILABLE</code></td><td>Catchall. Necessary resources were not available.</td></tr><tr><td><code>EX_SYNTAX</code></td><td>Syntax error in address.</td></tr><tr><td><code>EX_SOFTWARE</code></td><td>Internal software error, including bad arguments.</td></tr><tr><td><code>EX_OSERR</code></td><td>Temporary operating system error, such as “cannot fork”.</td></tr><tr><td><code>EX_NOHOST</code></td><td>Host name not recognized.</td></tr></table>	<code>EX_OK</code>	Successful completion on all addresses.	<code>EX_NOUSER</code>	User name not recognized.	<code>EX_UNAVAILABLE</code>	Catchall. Necessary resources were not available.	<code>EX_SYNTAX</code>	Syntax error in address.	<code>EX_SOFTWARE</code>	Internal software error, including bad arguments.	<code>EX_OSERR</code>	Temporary operating system error, such as “cannot fork”.	<code>EX_NOHOST</code>	Host name not recognized.
<code>EX_OK</code>	Successful completion on all addresses.														
<code>EX_NOUSER</code>	User name not recognized.														
<code>EX_UNAVAILABLE</code>	Catchall. Necessary resources were not available.														
<code>EX_SYNTAX</code>	Syntax error in address.														
<code>EX_SOFTWARE</code>	Internal software error, including bad arguments.														
<code>EX_OSERR</code>	Temporary operating system error, such as “cannot fork”.														
<code>EX_NOHOST</code>	Host name not recognized.														

	EX_TEMPFAIL	Message could not be sent immediately, but was queued.				
FILES	dead.letter	unmailable text				
	/etc/mail/aliases	mail aliases file (plain text)				
	/etc/mail/aliases.dir	database of mail aliases (binary)				
	/etc/mail/aliases.pag	database of mail aliases (binary)				
	/etc/mail/sendmail.cf	defines environment for sendmail				
	/var/spool/mqueue/*	temp files and queued mail				
	~/ .forward	list of recipients for forwarding messages				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:					
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWsndmu</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWsndmu
ATTRIBUTE TYPE	ATTRIBUTE VALUE					
Availability	SUNWsndmu					
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>The -ba, -bd , -bi , -bs , -bt , -bv , -M , and -q options require that sendmail be invoked from the trusted path with UID of 0 and that needed privileges be inherited. The -d and -X options are ignored if sendmail is not invoked from the trusted path. The -bp option will list only queued messages that are dominated by the process. The -p processing option in the configuration file specifies actions to take for mail received at a sensitivity label that is below the recipient’s minimum label. The -D option is not supported in the Trusted Solaris environment.</p> <p>The LabelAdminLow and LabelTooLow processing options may be specified.</p>					
Trusted Solaris 8 4/01 Reference Manual SunOS 5.8 Reference Manual	initgroups(3C), resolver(3RESOLV), nsswitch.conf(4)					
	biff(1B), mail(1), mailx(1), newaliases(1), check-hostname(1M), check-permissions(1M), getusershell(3C), aliases(4), hosts(4), attributes(5), largefile(5)					
	Postel, Jon, <i>Simple Mail Transfer Protocol</i> , RFC 821, Network Information Center, SRI International, Menlo Park, Calif., August 1982.					
	Crocker, Dave, <i>Standard for the Format of ARPA-Internet Text Messages</i> , RFC 822, Network Information Center, SRI International, Menlo Park, Calif., August 1982.					
	Costales, Bryan with Eric Allman, <i>sendmail, Second Edition</i> , O’Reilly & Associates, Inc., 1997.					
NOTES	The sendmail program requires a fully qualified host name when starting. A script has been included to help verify if the host name is defined properly (see check-hostname(1M)).					

sendmail(1M)

The permissions and the ownership of several directories have been changed in order to increase security. In particular, access to `/etc/mail` and `/var/spool/mqueue` has been restricted.

Security restrictions have been placed users using `.forward` files to pipe mail to a program or redirect mail to a file. The default shell (as listed in `/etc/passwd`) of these users must be listed in `/etc/shells`. This restriction does not affect mail that is being redirected to another alias.

Additional restrictions have been put in place on `.forward` and `:include:` files. These files and the directory structure that they are placed in cannot be group- or world-writable (see `check-permissions(1M)`).

NAME	setaudit – Run a command with the audit mask set				
SYNOPSIS	setaudit [-u <i>username</i>] <i>command</i> <i>command_args</i>				
DESCRIPTION	<p>setaudit invokes a command using the audit characteristics of the specified user, rather than the audit characteristics of the effective uid of the process executing the setaudit command. The command can be used to selectively turn on auditing for daemons and commands that are run from the <i>/etc/rc</i> scripts. If the -u option is not used, setaudit sets the audit characteristics to the context of the user invoking the command; if the option is present, setaudit sets the audit characteristics to the context of the specified <i>username</i>. Within the set context, setaudit then executes the specified <i>command</i> with its arguments (<i>command_args</i>).</p> <p>-u <i>username</i> Use the audit characteristics of <i>username</i> rather than the audit characteristics of the effective uid of the process executing the setaudit command.</p> <p><i>command</i> <i>command_args</i> The command to execute and its arguments.</p> <p>To succeed, setaudit must have the <i>file_dac_read</i> and <i>sys_audit</i> privileges in its set of effective privileges.</p>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWtsu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWtsu				
EXAMPLES	<p>EXAMPLE 1 To setaudit with the characteristics of a user</p> <p>To execute the <i>cat</i> command on the file <i>/etc/system</i> as the user <i>maverick</i>, use this:</p> <pre>setaudit -u maverick /usr/bin/cat /etc/system</pre> <p>EXAMPLE 2 To setaudit with the characteristics of the invoking shell</p> <p>To execute the <i>ls</i> command on the current working directory from the system shell, use the following command:</p> <pre>setaudit /sbin/sysh -c ls</pre>				
Trusted Solaris 8 4/01 Reference Manual	<p>audit_control(4), audit_user(4)</p> <p>attributes(5)</p>				

setfsattr(1M)

NAME	setfsattr, newsecfs – Set security attributes on an existing or newly created file system														
SYNOPSIS	<pre> /usr/sbin/setfsattr { [-l sensitivity-level-range] [-m MLD-prefix] [-p allowed-privilege-set] [-P forced-privilege-set] [-s CMW-Label] ... } {special filesystem} /usr/sbin/newsecfs { [-l sensitivity-level-range] [-M] [-m MLD-prefix] [-o newfs options] [-p allowed-privilege-set] [-P forced-privilege-set] [-s CMW-Label] ... } {special filesystem} </pre>														
DESCRIPTION	<p>setfsattr changes the security attributes of a file system. The file system may be specified either as a <i>filesystem</i> or as <i>special</i>, the device on which the file system resides. <i>filesystem</i> must be in <i>/etc/vfstab</i>, and it must be unmounted before setfsattr is invoked on it. setfsattr requires at least one option be specified; if not, an error is returned.</p> <p>newsecfs works similarly to setfsattr except that it runs newfs(1M) on the file system prior to setting the security attributes, then sets the label on the <i>lost+found</i> directory to [ADMIN_HIGH].</p>														
OPTIONS	<table> <tr> <td>-l <i>sensitivity-level-range</i></td><td>Set the filesystem sensitivity level range, a semicolon-separated pair of sensitivity labels. The labels must be valid sensitivity labels for the system. The first in the pair is the minimum sensitivity label, and it must be dominated by the second label, the maximum sensitivity label. The default is ADMIN_LOW;ADMIN_HIGH.</td></tr> <tr> <td>-M</td><td>Create the root directory of the file system as a multilevel directory (MLD). This option is available only with the newsecfs command.</td></tr> <tr> <td>-m <i>MLD-prefix</i></td><td>Set the file system MLD prefix. The default is ".MLD.". The MLD prefix is the string that disables multilevel directory translation in pathname lookup.</td></tr> <tr> <td>-o <i>newfs options</i></td><td>Set the file system newfs options. The options must be exactly the same as those expected by the newfs(1M) command. This option is available only with newsecfs.</td></tr> <tr> <td>-p <i>allowed-privileges</i></td><td>Set the file system allowed-privilege set, specified as a text-string of comma-separated privilege names. The privileges in the allowed set must include all privileges in the forced set, or the operation fails.</td></tr> <tr> <td>-P <i>forced-privileges</i></td><td>Set the filesystem forced-privilege set, specified as a text string of comma-separated privilege names. All privileges in the forced set must also be in the allowed set, or the operation fails.</td></tr> <tr> <td>-s <i>CMW-Label</i></td><td>Set the filesystem CMW label.</td></tr> </table>	-l <i>sensitivity-level-range</i>	Set the filesystem sensitivity level range, a semicolon-separated pair of sensitivity labels. The labels must be valid sensitivity labels for the system. The first in the pair is the minimum sensitivity label, and it must be dominated by the second label, the maximum sensitivity label. The default is ADMIN_LOW;ADMIN_HIGH.	-M	Create the root directory of the file system as a multilevel directory (MLD). This option is available only with the newsecfs command.	-m <i>MLD-prefix</i>	Set the file system MLD prefix. The default is ".MLD.". The MLD prefix is the string that disables multilevel directory translation in pathname lookup.	-o <i>newfs options</i>	Set the file system newfs options. The options must be exactly the same as those expected by the newfs(1M) command. This option is available only with newsecfs.	-p <i>allowed-privileges</i>	Set the file system allowed-privilege set, specified as a text-string of comma-separated privilege names. The privileges in the allowed set must include all privileges in the forced set, or the operation fails.	-P <i>forced-privileges</i>	Set the filesystem forced-privilege set, specified as a text string of comma-separated privilege names. All privileges in the forced set must also be in the allowed set, or the operation fails.	-s <i>CMW-Label</i>	Set the filesystem CMW label.
-l <i>sensitivity-level-range</i>	Set the filesystem sensitivity level range, a semicolon-separated pair of sensitivity labels. The labels must be valid sensitivity labels for the system. The first in the pair is the minimum sensitivity label, and it must be dominated by the second label, the maximum sensitivity label. The default is ADMIN_LOW;ADMIN_HIGH.														
-M	Create the root directory of the file system as a multilevel directory (MLD). This option is available only with the newsecfs command.														
-m <i>MLD-prefix</i>	Set the file system MLD prefix. The default is ".MLD.". The MLD prefix is the string that disables multilevel directory translation in pathname lookup.														
-o <i>newfs options</i>	Set the file system newfs options. The options must be exactly the same as those expected by the newfs(1M) command. This option is available only with newsecfs.														
-p <i>allowed-privileges</i>	Set the file system allowed-privilege set, specified as a text-string of comma-separated privilege names. The privileges in the allowed set must include all privileges in the forced set, or the operation fails.														
-P <i>forced-privileges</i>	Set the filesystem forced-privilege set, specified as a text string of comma-separated privilege names. All privileges in the forced set must also be in the allowed set, or the operation fails.														
-s <i>CMW-Label</i>	Set the filesystem CMW label.														

setfsattr(1M)

USAGE To specify arguments that include semicolons or embedded spaces (such as for the -l and -o options), use quotes to enclose the arguments.

EXAMPLES **EXAMPLE 1** To create a new file system with a limited label range
To create a new file system with an allowable label range of Confidential to Secret, use this command:

```
$ newsecfs -l 'confidential;secret' raw_device
```

EXIT STATUS setfsattr exits with one of these values:

0 Success.

1 Failure.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWtsu

Trusted Solaris 8
4/01 Reference
Manual
SunOS 8
Reference Manual

fork(2)

mkfs(1M), newfs(1M), terminfo(4), attributes(5)

setuname(1M)

NAME	setuname – Change machine information						
SYNOPSIS	setuname [-t] [-n <i>node</i>] [-s <i>name</i>]						
DESCRIPTION	<p>The setuname utility changes the parameter value for the system name and node name. Each parameter can be changed using setuname and the appropriate option.</p> <p>Either or both the -s and -n options must be given when invoking setuname.</p> <p>The system architecture may place requirements on the size of the system and network node name. The command will issue a fatal warning message and an error message if the name entered is incompatible with the system requirements.</p>						
OPTIONS	<p>The following options are supported:</p> <table><tr><td>-t</td><td>Temporary change. No attempt will be made to create a permanent change.</td></tr><tr><td>-n <i>node</i></td><td>Changes the node name. <i>node</i> specifies the new network node name and can consist of alphanumeric characters and the special characters dash, underbar, and dollar sign.</td></tr><tr><td>-s <i>name</i></td><td>Changes the system name. <i>name</i> specifies new system name and can consist of alphanumeric characters and the special characters dash, underbar, and dollar sign.</td></tr></table>	-t	Temporary change. No attempt will be made to create a permanent change.	-n <i>node</i>	Changes the node name. <i>node</i> specifies the new network node name and can consist of alphanumeric characters and the special characters dash, underbar, and dollar sign.	-s <i>name</i>	Changes the system name. <i>name</i> specifies new system name and can consist of alphanumeric characters and the special characters dash, underbar, and dollar sign.
-t	Temporary change. No attempt will be made to create a permanent change.						
-n <i>node</i>	Changes the node name. <i>node</i> specifies the new network node name and can consist of alphanumeric characters and the special characters dash, underbar, and dollar sign.						
-s <i>name</i>	Changes the system name. <i>name</i> specifies new system name and can consist of alphanumeric characters and the special characters dash, underbar, and dollar sign.						
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td rowspan="2">Availability</td><td>SUNWcsu (32-bit)</td></tr><tr><td>SUNWcsxu (64-bit)</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu (32-bit)	SUNWcsxu (64-bit)	
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWcsu (32-bit)						
	SUNWcsxu (64-bit)						
SUMMARY OF TRUSTED SOLARIS CHANGES SunOS 5.8 Reference Manual	<p>The setuname command must have the following privileges: file_dac_read, file_dac_write, file_mac_read, and file_mac_write.</p> <p>attributes(5)</p> <p>setuname attempts to change the parameter values in two places: the running kernel and, as necessary per implementation, to cross system reboots. A temporary change changes only the running kernel.</p>						
NOTES							

NAME	share – Make local resource available for mounting by remote systems
SYNOPSIS	share [-F <i>FSType</i>] [-o <i>specific_options</i>] [-d <i>description</i>] [<i>pathname</i>]
DESCRIPTION	The share command exports, or makes a resource available for mounting, through a remote file system of type <i>FSType</i> . If the option -F <i>FSType</i> is omitted, the first file system type listed in /etc/dfs/fstypes is used as default. For a description of NFS specific options, see share_nfs(1M). <i>pathname</i> is the pathname of the directory to be shared. When invoked with no arguments, share displays all shared file systems.
OPTIONS	<p>-F <i>FSType</i> Specify the filesystem type.</p> <p>-o <i>specific_options</i> The <i>specific_options</i> are used to control access of the shared resource. (See share_nfs(1M) for the NFS specific options.) They may be any of the following:</p> <p>devices nodevices Allow (disallow) opens on character and block devices. The default is devices.</p> <p>Note: In the Trusted Solaris environment, device special files are typically located only in the /dev and /devices directories in the root file system. All other file systems should be mounted with the nodevices option to prevent recognition of devices that may reside in any other directories.</p> <p>priv nopriv Forced privileges on executables are allowed or disallowed. The default is priv.</p> <p>rw <i>pathname</i> is shared read/write to all clients. This is also the default behavior.</p> <p>rw=client[:client]... <i>pathname</i> is shared read/write only to the listed clients. No other systems can access <i>pathname</i>.</p> <p>ro <i>pathname</i> is shared read-only to all clients.</p> <p>ro=client[:client]... <i>pathname</i> is shared read-only only to the listed clients. No other systems can access <i>pathname</i>.</p> <p>-d <i>description</i> The -d flag may be used to provide a description of the resource being shared.</p>
EXAMPLES	<p>EXAMPLE 1 A sample of using share command.</p> <p>This line will share the /disk file system read-only at boot time.</p> <pre>share -F nfs -o ro /disk</pre>
SUMMARY OF TRUSTED SOLARIS CHANGES	When invoked with no option or with only the -F <i>FSType</i> option, the share command displays shared file systems. For all other cases, the command must be run

share(1M)

with an effective UID of 0. If the shared file is of the type NFS, then the `sys_nfs` privilege is also required. If the file `/etc/dfs/sharetab` does not exist, this command will create it; therefore this command must be run at the sensitivity label `ADMIN_LOW`. If the file `/etc/dfs/sharetab` exists, this command can be run at any other sensitivity label if it has the `file_mac_write` privilege. To succeed in all cases, this command needs the `file_mac_read` and `file_mac_search` privileges.

FILES	<code>/etc/dfs/dfstab</code>	list of share commands to be executed at boot time
	<code>/etc/dfs/fstypes</code>	list of file system types, NFS by default
	<code>/etc/dfs/sharetab</code>	system record of shared file systems

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual
NOTES

`mountd(1M)`, `nfsd(1M)`, `share_nfs(1M)`, `shareall(1M)`, `unshare(1M)`
`attributes(5)`

Export (old terminology): file system sharing used to be called exporting on SunOS 4.x, so the `share` command used to be invoked as `exportfs(1B)` or `/usr/sbin/exportfs`.

If `share` commands are invoked multiple times on the same file system, the last `share` invocation supersedes the previous—the options set by the last `share` command replace the old options. For example, if read-write permission was given to `usera` on `/somefs`, then to give read-write permission also to `userb` on `/somefs`:

```
example% share -F nfs -o rw=usera:userb /somefs
```

This behavior is not limited to sharing the root file system, but applies to all file systems.

NAME	shareall, unshareall – Share, unshare multiple resources				
SYNOPSIS	shareall [-F <i>FSType</i> [, <i>FSType</i> ...]] [- <i>file</i>] unshareall [-F <i>FSType</i> [, <i>FSType</i> ...]]				
DESCRIPTION	<p>When used with no arguments, shareall shares all resources from <i>file</i>, which contains a list of share command lines. If the operand is a hyphen (-), then the share command lines are obtained from the standard input. Otherwise, if neither a <i>file</i> nor a hyphen is specified, then the file <i>/etc/dfs/dfstab</i> is used as the default.</p> <p>Resources may be shared by specific file system types by specifying the file systems in a comma-separated list as an argument to -F.</p> <p>unshareall unshares all currently shared resources. Without a -F flag, it unshares resources for all distributed file system types.</p>				
OPTIONS	-F <i>FSType</i> Specify file system type. Defaults to the first entry in <i>/etc/dfs/fstypes</i> .				
FILES	<i>/etc/dfs/dfstab</i> List of share commands to be executed at boot time.				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>The shareall and unshareall commands must be run with an effective UID of 0. If any file being shared or unshared is of the type NFS, then the command requires the <i>sys_nfs</i> privilege [see share_nfs(1M)]. If the file <i>/etc/dfs/sharetab</i> does not exist, the shareall command will create the file; thus, the shareall command must be run at the sensitivity level of <i>ADMIN_LOW</i>. If the file <i>/etc/dfs/sharetab</i> exists, then the shareall and unshareall commands can be run at any other sensitivity level if they have the <i>file_mac_write</i> privilege. To succeed in all cases, the commands need the <i>file_mac_read</i> and <i>file_mac_search</i> privileges.</p>				
Trusted Solaris 8 4/01 Reference Manual	share(1M), unshare(1M) attributes(5)				

share_nfs(1M)

NAME	share_nfs – Make local NFS file systems available for mounting by remote systems
SYNOPSIS	share [-d <i>description</i>] [-F nfs] [-o <i>specific_options</i>] <i>pathname</i>
DESCRIPTION	<p>The share utility makes local file systems available for mounting by remote systems.</p> <p>If no argument is specified, then share displays all file systems currently shared, including NFS file systems and file systems shared through other distributed file system packages.</p>
OPTIONS	<p>The following options are supported:</p> <ul style="list-style-type: none"> -d <i>description</i> Provide a comment that describes the file system to be shared. -F nfs Share NFS file system type. -o <i>specific_options</i> Specify <i>specific_options</i> in a comma-separated list of keywords and attribute-value-assertions for interpretation by the file-system-type-specific command. If <i>specific_options</i> is not specified, then by default sharing will be read-write to all clients. <p><i>specific_options</i> can be any combination of the following:</p> <ul style="list-style-type: none"> aclok Allows the NFS server to do access control for NFS Version 2 clients (running SunOS 2.4 or earlier). When aclok is set on the server, maximal access is given to all clients. For example, with aclok set, if anyone has read permissions, then everyone does. If aclok is not set, minimal access is given to all clients. anon=<i>uid</i> Set <i>uid</i> to be the effective user ID of unknown users. By default, unknown users are given the effective user ID <code>UID_NOBODY</code>. If <i>uid</i> is set to -1, access is denied. index=file Load <i>file</i> rather than a listing of the directory containing this file when the directory is referenced by an NFS URL. kerberos This option has been deprecated in favor of the <i>sec=krb4</i> option. log=tag Enables NFS server logging for the specified file system. The optional tag determines the location of

share_nfs(1M)

the related log files. The tag is defined in `etc/nfs/nfslog.conf`. If no tag is specified, the default values associated with the “global” tag in `etc/nfs/nfslog.conf` will be used.

nosub

Prevents clients from mounting subdirectories of shared directories. For example, if `/exportF` is shared with the `nosub` option on server *foeey* then a NFS client will not be able to do:

```
mount -F nfs foeey:/export/home/mnt
```

nosuid

By default, clients are allowed to create files on the shared file system with the `setuid` or `setgid` mode enabled. Specifying `nosuid` causes the server file system to silently ignore any attempt to enable the `setuid` or `setgid` mode bits.

nodevices

By default, clients are allowed to create block and character special devices on the shared file system. Specifying `nodevices` causes the server file system to prevent the creation of such devices.

nopriv

By default, clients are allowed to set forced privileges on files on the shared file system. Specifying `nopriv` causes the server file system to prevent the setting of forced privileges.

public

Moves the location of the public file handle from root (`/`) to the exported directory for WebNFS-enabled browsers and clients. This option does not enable WebNFS service; WebNFS is always on. Only one file system per server may use this option. Any other option, including the `-ro=list` and `-rw=list` options can be included with the `public` option.

ro

Sharing will be read-only to all clients.

ro=access_list

Sharing will be read-only to the clients listed in *access_list*; overrides the `rw` suboption for the clients specified. See *access_list* below.

share_nfs(1M)

root=access_list

Only root users from the hosts specified in *access_list* will have root access. See *access_list* below. By default, no host has root access, so root users are mapped to an anonymous user ID (see the *anon=uid* option described above). Netgroups can be used if the file system shared is using UNIX authentication (AUTH_SYS).

rw

Sharing will be read-write to all clients.

rw=access_list

Sharing will be read-write to the clients listed in *access_list*; overrides the *ro* suboption for the clients specified. See *access_list* below.

sec=mode[:mode]. . .

Sharing will use one or more of the specified security modes. The *mode* in the *sec=mode* option must be a node name supported on the client. If the *sec=* option is not specified, the default security mode used is AUTH_SYS. Multiple *sec=* options can be specified on the command line, although each mode can appear only once. The security modes are defined in *nfssec(5)*.

Each *sec=* option specifies modes that apply to any subsequent *window=*, *rw*, *ro*, *rw=*, *ro=* and *root=* options that are provided before another *sec=* option. Each additional *sec=* resets the security mode context, so that more *window=*, *rw*, *ro*, *rw=*, *ro=* and *root=* options can be supplied for additional modes.

sec=none

If the option *sec=none* is specified when the client uses AUTH_NONE, or if the client uses a security mode that is not one that the file system is shared with, then the credential of each NFS request is treated as unauthenticated. See the *anon=uid* option for a description of how unauthenticated requests are handled.

secure

This option has been deprecated in favor of the *sec=dh* option.

share_nfs(1M)

window=*value*

When sharing with `sec=dh` or `sec=krb4` set the maximum life time (in seconds) of the RPC request's credential (in the authentication header) that the NFS server will allow. If a credential arrives with a life time larger than what is allowed, the NFS server will reject the request. The default value is 30000 seconds (8.3 hours).

access_list The *access_list* argument is a colon-separated list whose components may be any number of the following:

hostname The name of a host. With a server configured for DNS or LDAP naming in the `nsswitch` "hosts" entry, any hostname must be represented as a fully qualified DNS or LDAP name.

netgroup A netgroup contains a number of hostnames. With a server configured for DNS or LDAP naming in the `nsswitch` "hosts" entry, any hostname in a netgroup must be represented as a fully qualified DNS or LDAP name.

domain name
suffix To use domain membership the server must use DNS or LDAP to resolve hostnames to IP addresses; that is, the "hosts" entry in the `/etc/nsswitch.conf` must specify "dns" or "ldap" ahead of "nis" or "nisplus", since only DNS and LDAP return the full domain name of the host. Other name services like NIS or NIS+ cannot be used to resolve hostnames on the server because when mapping an IP address to a hostname they do not return domain information. For example,

NIS or NIS+ 129.144.45.9 --> "myhost

DNS or LDAP 129.144.45.9 -->
"myhost.mydomain.mycompany.com"

The domain name suffix is distinguished from hostnames and netgroups by a prefixed dot. For example,

`rw=.mydomain.mycompany.com`

A single dot can be used to match a hostname with no suffix. For example,

`rw=.`

will match "mydomain" but not "mydomain.mycompany.com". This feature can be used to match hosts resolved through NIS and NIS+ rather than DNS and LDAP.

share_nfs(1M)

network

The network or subnet component is preceded by an at-sign (@). It can be either a name or a dotted address. If a name, it will be converted to a dotted address by `getnetbyname(3SOCKET)`. For example,

`=@mynet` would be equivalent to:

`=@129.144` or `=@129.144.0.0` The network prefix assumes an octet aligned netmask determined from the zero octets in the low-order part of the address. In the case where network prefixes are not byte-aligned, the syntax will allow a mask length to be specified explicitly following a slash (/) delimiter. For example,

`=@mynet/17` or `rw=@129.144.132/17` where the mask is the number of leftmost contiguous significant bits in the corresponding IP address.

domain name suffix

To use domain membership the server must use DNS or LDAP to resolve hostnames to IP addresses; that is, the "hosts" entry in the `/etc/nsswitch.conf` must specify "dns" or "ldap" ahead of "nis" or "nisplus", since only DNS and LDAP return the full domain name of the host. Other name services like NIS or NIS+ cannot be used to resolve hostnames on the server because when mapping an IP address to a hostname they do not return domain information. For example,

NIS or NIS+ 129.144.45.9 --> "myhost

DNS or LDAP 129.144.45.9 -->
 "myhost.mydomain.mycompany.com"

The domain name suffix is distinguished from hostnames and netgroups by a prefixed dot. For example,

`rw=.mydomain.mycompany.com`

A single dot can be used to match a hostname with no suffix. For example,

`rw=.`

share_nfs(1M)

will match "mydomain" but not "mydomain.mycompany.com". This feature can be used to match hosts resolved through NIS and NIS+ rather than DNS and LDAP.

network The network or subnet component is preceded by an at-sign (@). It can be either a name or a dotted address. If a name, it will be converted to a dotted address by `getnetbyname(3SOCKET)`. For example,

`=@mynet` would be equivalent to:

`=@129.144` or `=@129.144.0.0` The network prefix assumes an octet aligned netmask determined from the zero octets in the low-order part of the address. In the case where network prefixes are not byte-aligned, the syntax will allow a mask length to be specified explicitly following a slash (/) delimiter. For example,

`=@mynet/17` or `rw=@129.144.132/17` where the mask is the number of leftmost contiguous significant bits in the corresponding IP address.

A prefixed minus sign (-) denies access to that component of *access_list*. The list is searched sequentially until a match is found that either grants or denies access, or until the end of the list is reached. For example, if host "terra" is in the "engineering" netgroup, then

`rw=-terra:engineering` will deny access to terra but

`rw=engineering:-terra` will grant access to terra.

OPERANDS The following operands are supported:

pathname The pathname of the file system to be shared.

EXAMPLES **EXAMPLE 1** Sharing A File System With Logging Enabled

The following example shows the `/export` file system shared with logging enabled:

```
example% share -o log /export
```

The default global logging parameters are used since no tag identifier is specified. The location of the log file, as well as the necessary logging work files, is specified by the global entry in `/etc/nfs/nfslog.conf`. Note that the `nfslogd(1M)` daemon will run only if at least one file system entry in `/etc/dfs/dfstab` is shared with logging enabled upon starting or rebooting the system. Simply sharing a file system with logging enabled from the command line will not start the `nfslogd(1M)`.

share_nfs(1M)

EXIT STATUS

The following exit values are returned:

0 Successful completion.
>0 An error occurred.

FILES

/etc/dfs/dfstab List of share commands to be executed at boot time.
/etc/dfs/fstypes List of system types, NFS by default.
/etc/dfs/sharetab System record of shared file systems.
/etc/nfs/nfslogtab system record of logged file systems
/etc/nfs/nfslog.conf logging configuration file

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

The nodevices and nopriv options have been added.

When invoked with no option or with only the option `-F FSType`, the `share_nfs` command displays shared file systems. For all other cases, the command must be run with an effective UID of 0. If the shared file is of the type NFS, then the `sys_nfs` privilege is also required. If the file `/etc/dfs/sharetab` does not exist, this command will create it; therefore this command must be run at the sensitivity label `ADMIN_LOW`. If the file `/etc/dfs/sharetab` exists, this command can be run at any other sensitivity label if it has the `file_mac_write` privilege. To succeed in all cases, this command needs the `file_mac_read` and `file_mac_search` privileges.

Trusted Solaris 8 4/01 Reference Manual SunOS 5.8 Reference Manual

mount(1M), mountd(1M), nfsd(1M), share(1M), unshare(1M)

nfslogd(1M), getnetbyname(3SOCKET), netgroup(4), nfslog.conf(4), attributes(5), nfssec(5)

NOTES

If the `sec=` option is presented at least once, all uses of the `window=`, `rw`, `ro`, `rw=`, `ro=`, and `root=` options must come *after* the first `sec=` option. If the `sec=` option is not presented, then `sec=sys` is implied.

If one or more explicit `sec=` options are presented, `sys` must appear in one of the options mode lists for accessing using the `AUTH_SYS` security mode to be allowed. For example:

```
share -F nfs /var
share -F nfs -o sec=sys /var will grant read-write access to any host using AUTH_SYS, but
share -F nfs -o sec=dh /var will grant no access to clients that use AUTH_SYS.
```

Unlike previous implementations of `share_nfs(1M)`, access checking for the `window=`, `rw`, `ro`, `rw=`, and `ro=` options is done per NFS request, instead of per mount request.

Combining multiple security modes can be a security hole in situations where the `ro=` and `rw=` options are used to control access to weaker security modes. In this example,

```
share -F nfs -o sec=dh,rw,sec=sys,rw=hosta /var
```

an intruder can forge the IP address for `hosta` (albeit on each NFS request) to side-step the stronger controls of `AUTH_DES`. Something like:

```
share -F nfs -o sec=dh,rw,sec=sys,ro /var
```

is safer, because any client (intruder or legitimate) that avoids `AUTH_DES` will only get read-only access. In general, multiple security modes per `share` command should only be used in situations where the clients using more secure modes get stronger access than clients using less secure modes.

If `rw=` and `ro=` options are specified in the same `sec=` clause, and a client is in both lists, the order of the two options determines the access the client gets. If client `hosta` is in two netgroups - `group1` and `group2` - in this example, the client would get read-only access:

```
share -F nfs -o ro=group1,rw=group2 /var
```

In this example `hosta` would get read-write access:

```
share -F nfs -o rw=group2,ro=group1 /var
```

If within a `sec=` clause, both the `ro` and `rw=` options are specified, for compatibility, the order of the options rule is not enforced. All hosts would get read-only access, with the exception to those in the read-write list. Likewise, if the `ro=` and `rw` options are specified, all hosts get read-write access with the exceptions of those in the read-only list.

The `ro=` and `rw=` options are guaranteed to work over UDP and TCP but may not work over other transport providers.

The `root=` option with `AUTH_SYS` is guaranteed to work over UDP and TCP but may not work over other transport providers.

The `root=` option with `AUTH_DES` and `AUTH_KERB` is guaranteed to work over any transport provider.

There are no interactions between the `root=` option and the `rw`, `ro`, `rw=`, and `ro=` options. Putting a host in the `root` list does not override the semantics of the other options. The access the host gets is the same as when the `root=` option is absent. For example, the following `share` command will deny access to `hostb`:

```
share -F nfs -o ro=hosta,root=hostb /var
```

share_nfs(1M)

The following will give read-only permissions to hostb:

```
share -F nfs -o ro=hostb,root=hostb /var
```

The following will give read-write permissions to hostb:

```
share -F nfs -o ro=hosta,rw=hostb,root=hostb /var
```

If the file system being shared is a symbolic link to a valid pathname, the canonical path (the path which the symbolic link follows) will be shared. For example, if `/export/foo` is a symbolic link to `/export/bar` (`/export/foo -> /export/bar`), the following `share` command will result in `/export/bar` as the shared pathname (and not `/export/foo`).

```
example# share -F nfs /export/foo
```

Note that an NFS mount of `server:/export/foo` will result in `server:/export/bar` really being mounted.

This line in the `/etc/dfs/dfstab` file will share the `/disk` file system read-only at boot time:

```
share -F nfs -o ro /disk
```

Note that the same command entered from the command line will not share the `/disk` file system unless there is at least one file system entry in the `/etc/dfs/dfstab` file. The `mountd(1M)` and `nfsd(1M)` daemons only run if there is a file system entry in `/etc/dfs/dfstab` when starting or rebooting the system.

NAME	showmount – Show all remote mounts				
SYNOPSIS	/usr/sbin/showmount [-ade] [<i>hostname</i>]				
DESCRIPTION	showmount lists all the clients that have remotely mounted a filesystem from <i>hostname</i> . This information is maintained by the mountd(1M) server on <i>hostname</i> , and is saved across crashes in the file <i>/etc/rmtab</i> . The default value for <i>hostname</i> is the value returned by <i>hostname(1)</i> .				
OPTIONS	<p>-a Print all remote mounts in the format:</p> <p style="padding-left: 40px;"><i>hostname : directory</i></p> <p style="padding-left: 40px;">where <i>hostname</i> is the name of the client, and <i>directory</i> is the root of the file system that has been mounted.</p> <p>-d List directories that have been remotely mounted by clients.</p> <p>-e Print the list of shared file systems.</p>				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>The list of clients that have remotely mounted a filesystem is maintained by the mountd(1M) server.</p> <p><i>/etc/rmtab</i> List of clients that have remotely mounted a filesystem from this machine.</p>				
ATTRIBUTES	<p>See <i>attributes(5)</i> for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
Trusted Solaris 8 4/01 Reference Manual	<p>mountd(1M)</p> <p><i>hostname(1)</i>, <i>attributes(5)</i></p>				
BUGS	If a client crashes, its entry will not be removed from the list of remote mounts on the server.				

smc(1M)

NAME	<code>smc</code> – start the Solaris Management Console (SMC)
SYNOPSIS	<p>smc [<i>subcommand</i>] [<i>args</i>]</p> <p>smc [<i>subcommand</i>] [<i>args</i>] -T <i>tool_name</i> [- - <i>tool_args</i>]</p>
DESCRIPTION	<p>The <code>smc</code> command starts the Solaris Management Console (SMC). The SMC is a graphical user interface that provides access to Solaris system administration tools. It relies on SMC servers running on one or more computers to perform modifications and report data. Each of these servers is a repository for code which the console can retrieve after the user of the console has authenticated himself or herself to the server.</p> <p>The console can also retrieve toolboxes from the server. These toolboxes are descriptions of organized collections of tools available on that and possibly other servers. Once one of these toolboxes is loaded, the console will display it and the tools referenced in it.</p> <p>The console can also run in a terminal (non-graphically), for use over remote connections or non-interactively from a script.</p> <p>For information on the use of the graphical console, and for more detailed explanations of authentication, tools, and toolboxes, please refer to the SMC online help available under the "Help" menu in the SMC console.</p>
<i>subcommands</i>	<p><code>smc</code> <i>subcommands</i> are:</p> <p><code>open</code> The default subcommand for the SMC is <code>open</code>. This will launch the console and allow you to run tools from the toolboxes you load. It does not need to be specified explicitly on the command line.</p> <p><code>edit</code> The <code>edit</code> subcommand will also launch the console, like the <code>open</code> subcommand. However, after loading a toolbox, you will not be able to run the referenced tools. Instead, you will be able to edit that toolbox, that is, add, remove, or modify any tools or folders in that toolbox.</p>
OPTIONS	<p>The following options are supported. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either <code>-D</code> or <code>- --domain</code> with the <i>domain</i> argument.</p> <p>If <i>tool_args</i> are specified, they must be preceded by the <code>- -</code> option and separated from the double dashes by a space.</p> <p>- <code>--auth-data <i>file</i></code> Specifies a file which the console can read to collect authentication data. When running the SMC non-interactively, the console will still need to authenticate itself with the server to retrieve tools. This data can either be passed on the command line using the <code>-u</code>, <code>-p</code>, <code>-r</code>, and <code>-l</code> options (which is insecure, because any user can see this data), or it can be placed in a file for the console to read. For security reasons, this file should be readable only by the user running the console, although the console does not enforce this restriction.</p>

The format of *file* is:

```
hostname=host name
username=user name
password=password for user name
rolename=role name
rolepassword=password for role name
```

Only one set of *hostname-username-password-rolename-rolepassword* may be specified in any one file. If the *rolename* is not specified, no role will be assumed.

-B | **-** *-toolbox toolbox*

Loads the specified toolbox. *toolbox* can be either a fully-qualified URL or a filename. If you specify an HTTP URL as, for example,

```
http://host_name:port/ . . .
```

it must point to a *host_name* and *port* on which an SMC server is running. If you omit *port*, the default port, 898, is used. This option overrides the **-H** option.

-D | **-** *-domain domain*

Specifies the default domain that you want to manage. The syntax of *domain* is *type:/host_name/domain_name*, where *type* is *nis*, *nisplus*, *dns*, *ldap*, or *file*; *host_name* is the name of the machine that serves the domain; and *domain_name* is the name of the domain you want to manage. (*Note: Do not use nis+ for nisplus.*) This option applies only to a single tool run in the terminal console.

If you do not specify this option, the SMC assumes the *file* default domain on whatever server you choose to manage, meaning that changes are local to the server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.

-h | **-** *-help*

Prints a usage statement about the `smc` command and its subcommands to the terminal window. To print a usage statement for one of the subcommands, enter **-h** after the subcommand.

-H | **-** *-hostname host_name:port*

Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the SMC connects to the local host on port 898. You may still have to choose a toolbox to load into the console. To override this behavior, use the **-B** option (see above), or set your console preferences to load a “home toolbox” by default.

-J*java_option*

Specifies an option that can be passed directly to the Java runtime (see `java(1)`). Do not enter a space between **-J** and the argument. This option is most useful for developers.

smc(1M)

-l | **-rolepassword** *role_password*

Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

When the SMC client and server are both running the Trusted Solaris operating environment, role assumption relies on trusted networking. The Trusted Path status and the user role identities are retrieved as network attributes. Therefore, the user must first assume the role using the CDE Trusted Path menu before invoking the SMC. Only the role password is required for authentication in this case; neither the **-u**, **-p**, nor **-r** options are required because they are obtained as network attributes.

If the SMC server is running the Trusted Solaris operating environment and the SMC client is not, role assumption is not allowed by default. This policy can be changed by including the **-u** option in the `smcwbemserver` script on the SMC server. See `init.wbem(1M)`.

-p | **-password** *password*

Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-r | **-rolename** *role_name*

Specifies a role name for authentication. If you are running the SMC in a terminal and you do not specify this option, no role is assumed. The GUI console may prompt you for a role name, although you may not need to assume a role.

If the SMC server is running the Trusted Solaris operating environment and the SMC client is not, this option is subject to the same restrictions as the **-l** option.

-s | **-silent**

Disables informational messages printed to the terminal.

-t

Runs the SMC in terminal mode. If this option is not given, the SMC will automatically run in terminal mode if it cannot find a graphical display.

-trust

Trusts all downloaded code implicitly. Use this option when running the terminal console non-interactively and you cannot let the console wait for user input.

-T | **-tool** *tool_name*

Runs the tool with the Java class name that corresponds to *tool_name*. If you do not specify this option and the SMC is running in terminal mode, the system prompts you. If the SMC is running in graphical mode, the system either loads a toolbox or prompts you for one (see options **-H** and **-B**).

-u | - --username *user_name*
 Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.

-v | - --version
 Prints the version of the SMC to the terminal. In the graphical console, this information can be found in the About box, available from the Help menu.

-y | - --yes
 Answers yes to all yes/no questions. Use this option when running the terminal console non-interactively and you cannot let the console wait for user input.

EXAMPLES **EXAMPLE 1** Printing a usage statement

The following prints a usage statement about the `smc` command to the terminal window:

```
smc --help
```

EXAMPLE 2 Passing an option to Java

The following passes an option through to the Java VM, which sets the `com.example.boolean` system property to `true`. This system property is only an example; the SMC does not use it.

```
smc -J-Dcom.example.boolean=true
```

ENVIRONMENT VARIABLES

See `environ(5)` for descriptions of the following environment variables that affect the execution of the `smc` command:

<code>JAVA_HOME</code>	If you do not specify this environment variable, the <code>/usr/java1.2</code> location is used.
<code>DISPLAY</code>	If you do not set this environment variable, set it to null, or set it to an <code>X(7)</code> display to which you are not authorized to connect, the SMC starts in terminal mode instead of graphical mode.

EXIT STATUS

The following exit values are returned. Other error codes may be returned if you specify a tool (using `-T tool_name`) that has its own error codes. See the documentation for the appropriate tool.

0	Successful completion.
1	An error occurred.

SUMMARY OF TRUSTED SOLARIS CHANGES

The `-l` and `-r` options are limited when the SMC client is not running the Trusted Solaris operating environment. The `-u`, `-p`, and `-r` options are obtained as network attributes.

ATTRIBUTES

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

smc(1M)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmcc

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

auths(1), init.wbem(1M), smprofile(1), smrole(1)
java(1), smcconf(1M), attributes(5), environ(5), x(7)

NAME	smcron – manage jobs in the crontab database								
SYNOPSIS	<code>/usr/sadm/bin/smcron subcommand [auth_args] - - [subcommand_args]</code>								
DESCRIPTION	The smcron command manages jobs in the crontab(1) database.								
subcommands	<p>smcron subcommands are:</p> <table> <tr> <td>add</td><td>Adds a job to the crontab(1) database. To add a job, the administrator must have the <code>solaris.jobs.admin</code> authorization. To add a job to another user's crontab file, the administrator must have the <code>solaris.jobs.user</code> authorization.</td></tr> <tr> <td>delete</td><td>Deletes a job from the crontab(1) database. To delete a job, the administrator must have the <code>solaris.jobs.admin</code> authorization. To delete a job from another user's crontab file, the administrator must have the <code>solaris.jobs.user</code> authorization.</td></tr> <tr> <td>list</td><td>Lists one or more jobs in the crontab(1) database. To list all jobs, the administrator must have the <code>solaris.jobs.admin</code> authorization. To list a job in another user's crontab file, the administrator must have the <code>solaris.jobs.user</code> authorization. No authorization is needed to list a user's own jobs.</td></tr> <tr> <td>modify</td><td>Modifies a job in the crontab(1) database. To modify a job, the administrator must have the <code>solaris.jobs.admin</code> authorization. To modify a job in another user's crontab file, the administrator must have the <code>solaris.jobs.user</code> authorization.</td></tr> </table>	add	Adds a job to the crontab(1) database. To add a job, the administrator must have the <code>solaris.jobs.admin</code> authorization. To add a job to another user's crontab file, the administrator must have the <code>solaris.jobs.user</code> authorization.	delete	Deletes a job from the crontab(1) database. To delete a job, the administrator must have the <code>solaris.jobs.admin</code> authorization. To delete a job from another user's crontab file, the administrator must have the <code>solaris.jobs.user</code> authorization.	list	Lists one or more jobs in the crontab(1) database. To list all jobs, the administrator must have the <code>solaris.jobs.admin</code> authorization. To list a job in another user's crontab file, the administrator must have the <code>solaris.jobs.user</code> authorization. No authorization is needed to list a user's own jobs.	modify	Modifies a job in the crontab(1) database. To modify a job, the administrator must have the <code>solaris.jobs.admin</code> authorization. To modify a job in another user's crontab file, the administrator must have the <code>solaris.jobs.user</code> authorization.
add	Adds a job to the crontab(1) database. To add a job, the administrator must have the <code>solaris.jobs.admin</code> authorization. To add a job to another user's crontab file, the administrator must have the <code>solaris.jobs.user</code> authorization.								
delete	Deletes a job from the crontab(1) database. To delete a job, the administrator must have the <code>solaris.jobs.admin</code> authorization. To delete a job from another user's crontab file, the administrator must have the <code>solaris.jobs.user</code> authorization.								
list	Lists one or more jobs in the crontab(1) database. To list all jobs, the administrator must have the <code>solaris.jobs.admin</code> authorization. To list a job in another user's crontab file, the administrator must have the <code>solaris.jobs.user</code> authorization. No authorization is needed to list a user's own jobs.								
modify	Modifies a job in the crontab(1) database. To modify a job, the administrator must have the <code>solaris.jobs.admin</code> authorization. To modify a job in another user's crontab file, the administrator must have the <code>solaris.jobs.user</code> authorization.								
OPTIONS	<p>The smcron authentication arguments, <i>auth_args</i>, are derived from the smc(1M) arg set and are the same regardless of which subcommand you use. The smcron command requires the SMC to be initialized for the command to succeed (see smc(1M)). After rebooting the SMC server, the first smc connection may time out, so you may need to retry the command.</p> <p>The subcommand-specific options, <i>subcommand_args</i>, must come after the <i>auth_args</i> and must be separated from them by the - - option.</p>								
auth_args	<p>The valid <i>auth_args</i> are -D, -H, -l, -p, -r, and -u; they are all optional. If no <i>auth_args</i> are specified, certain defaults will be assumed and the user may be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either -D or - --domain with the <i>domain</i> argument.</p> <p>-D - --domain <i>domain</i></p> <p>Specifies the default domain that you want to manage. The syntax of <i>domain</i> is <i>type</i>:/<i>host_name</i>/<i>domain_name</i>, where <i>type</i> is nis, nisplus, dns, ldap, or file; <i>host_name</i> is the name of the machine that serves the domain; and <i>domain_name</i> is the name of the domain you want to manage. (Note: Do not use nis+ for nisplus.)</p>								

smcron(1M)

If you do not specify this option, the SMC assumes the `file` default domain on whatever server you choose to manage, meaning that changes are local to the server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.

`-H | - -hostname host_name:port`

Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the SMC connects to the local host on port 898. You may still have to choose a toolbox to load into the console. To override this behavior, use the `smc(1M) -B` option, or set your console preferences to load a “home toolbox” by default.

`-l | - -rolepassword role_password`

Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

`-p | - -password password`

Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

`-r | - -rolename role_name`

Specifies a role name for authentication. If you do not specify this option, no role is assumed.

`-u | - -username user_name`

Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.

`- -`

This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the `- -` option.

subcommand_args

Note: Descriptions and other arg options that contain white spaces must be enclosed in double quotes.

■ For subcommand `add`:

`-c command` Specifies the command that you want to run.

`-h` (Optional) Displays the command’s usage statement.

`-m day_of_month` (Optional) Specifies the day of the month you want to run the job. Valid values are 1–31. If you specify both `-t` and `-m` options, the job executes one day per month at the time specified by `-t`.

- | | |
|-----------------------|---|
| -M <i>month</i> | (Optional) Specifies the month that you want to run the job. Valid values are 1–12. If you specify both -t and -M options, the job executes during the specified month at the time specified by -t . |
| -n <i>name</i> | Specifies the unique name of the job. |
| -o <i>owner</i> | (Optional) Specifies the user name that is the owner of the job. If you do not specify this option, the user name specified by the -U option is assumed. |
| -t <i>time_of_day</i> | Specifies the time (in <i>hh:mm</i>) that you want to execute the command. If no other time-related options are specified (-m, -M , or -w), the job executes every day at the time specified by -t . If you specify both -t and -w options, the job executes one day per week at the time specified by -t . If you specify both -t and -m options, the job executes one day per month at the time specified by -t . If you specify both -t and -M options, the job executes each day during the specified month at the time specified by -t . |
| -w <i>day_of_week</i> | <p>(Optional) Specifies the day of the week you want to execute the command. Valid values are as follows:</p> <ul style="list-style-type: none"> ■ 0=Sunday ■ 1=Monday ■ 2=Tuesday ■ 3=Wednesday ■ 4=Thursday ■ 5=Friday ■ 6=Saturday <p>If you specify both -t and -w options, the job executes one day per week at the time specified by -t .</p> |
- For subcommand **delete**:
- | | |
|-----------------|--|
| -h | (Optional) Displays the command's usage statement. |
| -n <i>name</i> | Specifies the unique name of the job. |
| -o <i>owner</i> | (Optional) Specifies the user name that is the owner of the job. If you do not specify this option, the user name specified by the -U option is assumed. |
- For subcommand **list**:
- | | |
|-----------------|---|
| -f <i>n s v</i> | <p>(Optional) Specifies the format of the output. See EXAMPLES for examples of each output type.</p> <ul style="list-style-type: none"> ■ <i>n</i> — Displays the data in native format, as it appears in the crontab(1) database. ■ <i>s</i> — Default format. Displays the data in summary format. ■ <i>v</i> — Displays the data in verbose format. |
|-----------------|---|

smcron(1M)

-h	(Optional) Displays the command's usage statement.
-o <i>owner</i>	(Optional) Lists all jobs for the specified owner (user name). If you do not specify this option, all jobs in the <code>crontab(1)</code> database are listed.
■ For subcommand <code>modify</code> :	
-c <i>command</i>	(Optional) Specifies the command that you want to run.
-h	(Optional) Displays the command's usage statement.
-m <i>day_of_month</i>	(Optional) Specifies the day of the month you want to run the job. Valid values are 1–31. If you specify both <code>-t</code> and <code>-m</code> options, the job executes one day per month at the time specified by <code>-t</code> .
-M <i>month</i>	(Optional) Specifies the month that you want to run the job. Valid values are 1–12. If you specify both <code>-t</code> and <code>-M</code> options, the job executes during the specified month at the time specified by <code>-t</code> .
-n <i>name</i>	Specifies the current unique name of the job.
-N <i>new_name</i>	(Optional) Specifies the new unique name of the job.
-o <i>owner</i>	(Optional) Specifies the user name that is the owner of the job. If you do not specify this option, the user name specified by the <code>-U</code> option is assumed.
-O <i>new_owner</i>	(Optional) Specifies the new owner of the job.
-t <i>time_of_day</i>	(Optional) Specifies the time (in <i>hh:mm</i>) that you want to execute the command. If no other time-related options are specified (<code>-m</code> , <code>-M</code> , or <code>-w</code>), then the job executes every day at the time specified by <code>-t</code> . If you specify both <code>-t</code> and <code>-w</code> options, the job executes one day per week at the time specified by <code>-t</code> . If you specify both <code>-t</code> and <code>-m</code> options, the job executes one day per month at the time specified by <code>-t</code> . If you specify both <code>-t</code> and <code>-M</code> , then the job executes each day during the specified month at the time specified by <code>-t</code> .
-w <i>day_of_week</i>	(Optional) Specifies the day of the week you want to execute the command. Valid values are as follows: <ul style="list-style-type: none">■ 0=Sunday■ 1=Monday■ 2=Tuesday■ 3=Wednesday■ 4=Thursday■ 5=Friday■ 6=Saturday If you specify both <code>-t</code> and <code>-w</code> options, the job executes one day per week at the time specified by <code>-t</code> .

EXAMPLES**EXAMPLE 1** Adding a job

The admin role adds a new job owned by kmochida, with the name, Remove old logs, that removes the old log files from /tmp every Monday at 1:30 AM. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smcron add -- -o kmochida -n "Remove old logs" \
-c "rm /tmp/*.log" -w 1 -t 1:30
```

EXAMPLE 2 Deleting a job

The admin role deletes the job named Remove old logs owned by kmochida. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smcron delete -- -n "Remove old logs" -o kmochida
```

EXAMPLE 3 Listing jobs in native format

The admin role lists all jobs in native, or crontab(1), format. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smcron list -- -f n
MINUTE HOUR DATE MONTH DAY COMMAND

10 3 * * 0,4 /etc/cron.d/logchecker
10 3 * * 0 /usr/lib/newsyslog
15 3 * * 0 /usr/lib/fs/nfs/nfsfind
1 2 * * * [ -x /usr/sbin/rtc ] && /usr/sbin/rtc -c > /dev/null 2>&1
```

EXAMPLE 4 Listing jobs in standard format

The admin role lists all jobs owned by kmochida in standard format. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smcron list -- -f s -o kmochida
NAME::OWNER::SCHEDULE::COMMAND

NoName_1765663371::lp::Weekly on Sundays at 3:13 AM::cd /var/lp/logs;
    if [ -f requests ]; then if [ -f requests.1 ]; then /bin/mv requests.1
    requests.2; fi; /usr/bin/cp requests requests.1; > requests; fi
NoName_512822673::lp::Weekly on Sundays at 4:15 AM::cd /var/lp/logs;
    if [ -f lpsched ]; then if [ -f lpsched.1 ]; then /bin/mv lpsched.1
    lpsched.2; fi; /usr/bin/cp lpsched lpsched.1; >lpsched; fi
```

EXAMPLE 5 Listing jobs in verbose format

The admin role lists all jobs in verbose format. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smcron list -- -f v
NAME::OWNER::SCHEDULE::NEXT_RUN::STATUS::COMMAND
```

smcron(1M)

EXAMPLE 5 Listing jobs in verbose format (Continued)

```
NoName_1075488942::root::Advanced:::Finished on Feb 10 3:10 with code 1
::/etc/cron.d/logchecker
databackup::root::Weekly on Sundays at 3:10 AM::3/19/00 3:10 AM
::Finished on Sep 19 3:10::/usr/lib/newsyslog
runlog::root::Daily at 2:01 AM::3/14/00 2:01 AM::Finished on Feb 11
2:01 AM::/usr/sbin/rtc
```

EXAMPLE 6 Modifying a job

The admin role modifies the job Remove old logs owned by kmochida to execute daily at 2:00 AM. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smcron modify -- -n "Remove old logs" \
-o kmochida -t 2:00
```

ENVIRONMENT VARIABLES

See environ(5) for a description of the JAVA_HOME environment variable, which affects the execution of the smcron command. If this environment variable is not specified, the /usr/java location is used. See smc(1M).

EXIT STATUS

The following exit values are returned:

- | | |
|---|---|
| 0 | Successful completion. |
| 1 | Invalid command syntax. A usage message displays. |
| 2 | An error occurred while executing the command. An error message displays. |

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmga

SUMMARY OF TRUSTED SOLARIS CHANGES

To add, delete, or modify an entry or to list all entries, the administrator must have the solaris.jobs.admin authorization. To add, delete, modify, or list another user's crontab entry, the administrator must have the solaris.jobs.user authorization. No authorization is needed to list a user's own jobs.

Trusted Solaris 8 4/01 Reference Manual

crontab(1), cron(1M), smc(1M), smprofile(1M), smrole(1M)
attributes(5), environ(5)

NAME	smexec – manage entries in the exec_attr database						
SYNOPSIS	<code>/usr/sadm/bin/smexec subcommand [auth_args] - - [subcommand_args]</code>						
DESCRIPTION	<p>The smexec command manages an entry in the exec_attr(4) database in the local /etc files name service or a NIS or NIS+ name service.</p> <p>Symlinked commands should not be used as an argument to smexec. If a non-existent command is passed, the smexec command accepts it, but it does not work.</p> <p><i>subcommands</i> smexec subcommands are:</p> <table> <tr> <td>add</td><td>Adds a new entry to the exec_attr(4) database. To add an entry to the exec_attr database, the administrator must have the solaris.profmgr.execattr.write authorization.</td></tr> <tr> <td>delete</td><td>Deletes an entry from the exec_attr(4) database. To delete an entry from the exec_attr database, the administrator must have the solaris.profmgr.execattr.write authorization.</td></tr> <tr> <td>modify</td><td>Modifies an entry in the exec_attr(4) database. To modify an entry in the exec_attr database, the administrator must have the solaris.profmgr.execattr.write authorization.</td></tr> </table>	add	Adds a new entry to the exec_attr(4) database. To add an entry to the exec_attr database, the administrator must have the solaris.profmgr.execattr.write authorization.	delete	Deletes an entry from the exec_attr(4) database. To delete an entry from the exec_attr database, the administrator must have the solaris.profmgr.execattr.write authorization.	modify	Modifies an entry in the exec_attr(4) database. To modify an entry in the exec_attr database, the administrator must have the solaris.profmgr.execattr.write authorization.
add	Adds a new entry to the exec_attr(4) database. To add an entry to the exec_attr database, the administrator must have the solaris.profmgr.execattr.write authorization.						
delete	Deletes an entry from the exec_attr(4) database. To delete an entry from the exec_attr database, the administrator must have the solaris.profmgr.execattr.write authorization.						
modify	Modifies an entry in the exec_attr(4) database. To modify an entry in the exec_attr database, the administrator must have the solaris.profmgr.execattr.write authorization.						
OPTIONS	<p>The smexec authentication arguments, <i>auth_args</i>, are derived from the smc(1M) arg set and are the same regardless of which subcommand you use. The smexec command requires the SMC to be initialized for the command to succeed (see smc(1M)). After rebooting the SMC server, the first smc connection may time out, so you may need to retry the command.</p> <p>The subcommand-specific options, <i>subcommand_args</i>, must come after the <i>auth_args</i> and must be separated from them by the - - option.</p>						
<i>auth_args</i>	<p>The valid <i>auth_args</i> are -D, -H, -l, -p, -r, and -u; they are all optional. If no <i>auth_args</i> are specified, certain defaults will be assumed and the user may be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either -D or - --domain with the <i>domain</i> argument.</p> <p>-D - --domain <i>domain</i></p> <p>Specifies the default domain that you want to manage. The syntax of <i>domain</i> is <i>type:/host_name/domain_name</i>, where <i>type</i> is nis, nisplus, dns, ldap, or file; <i>host_name</i> is the name of the machine that serves the domain; and <i>domain_name</i> is the name of the domain you want to manage. (Note: Do not use nis+ for nisplus.)</p> <p>If you do not specify this option, the SMC assumes the file default domain on whatever server you choose to manage, meaning that changes are local to the server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.</p>						

smexec(1M)

- H | - -hostname *host_name:port*
Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the SMC connects to the local host on port 898. You may still have to choose a toolbox to load into the console. To override this behavior, use the smc(1M) -B option, or set your console preferences to load a “home toolbox” by default.
- l | - -rolepassword *role_password*
Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.
- p | - -password *password*
Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.
- r | - -rolename *role_name*
Specifies a role name for authentication. If you do not specify this option, no role is assumed.
- u | - -username *user_name*
Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.
- -
This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the - - option.

subcommand_args

Note: Descriptions and other arg options that contain white spaces must be enclosed in double quotes.

■ For subcommand add:

- C *clearance* (Optional) Specifies the human-readable string or hex representation of the clearance. It is a valid option when the `tsol` policy is specified.
- c *command_path* | *CDE_action* Specifies the full path to the command or CDE action associated with the new `exec_attr` entry.
- G *gid* (Optional) Specifies the real group ID that executes with the command or CDE action.
- g *egid* (Optional) Specifies the effective group ID that executes with the command or CDE action.
- h (Optional) Displays the command’s usage statement.

<code>-L label</code>	(Optional) Specifies the the human-readable string or hex representation of the label. It is a valid option when the <code>tsol</code> policy is specified.
<code>-n profile_name</code>	Specifies the name of the profile associated with the new <code>exec_attr</code> entry.
<code>-P priv_to_add1...</code>	Specifies the privilege name(s) or privilege number(s) to add to the new <code>exec_attr</code> entry. Additional privileges may be specified by specifying the <code>-P</code> multiple times. It is a valid option when the <code>tsol</code> policy is specified.
<code>-p policy</code>	Specifies the policy (<code>tsol</code> or <code>suser</code>) associated with the new <code>exec_attr</code> entry. If this option is not specified, the default is <code>suser</code> .
<code>-t type</code>	Specifies the type <code>cmd</code> for command, or type <code>act</code> for CDE action.
<code>-U uid</code>	(Optional) Specifies the real user ID that executes with the command or CDE action.
<code>-u euid</code>	(Optional) Specifies the effective user ID that executes with the command or CDE action.
■ For subcommand <code>delete</code> :	
<code>-c command_path CDE_action</code>	Specifies the full path to the command or CDE action associated with the <code>exec_attr</code> entry.
<code>-h</code>	(Optional) Displays the command's usage statement.
<code>-n profile_name</code>	Specifies the name of the profile associated with the <code>exec_attr</code> entry.
<code>-p policy</code>	(Optional) Specifies the policy (<code>tsol</code> or <code>suser</code>) associated with the new <code>exec_attr</code> entry. If this option is not specified, the default is <code>suser</code> .
<code>-t type</code>	Specifies the type <code>cmd</code> for command, or type <code>act</code> for CDE action.
■ For subcommand <code>modify</code> :	
<code>-C clearance</code>	(Optional) Specifies the human-readable string or hex representation of the clearance. It is a valid option when the <code>tsol</code> policy is specified.
<code>-c command_path CDE_action</code>	Specifies the full path to the command or CDE action associated with the <code>exec_attr</code> entry that you want to modify.
<code>-G gid</code>	(Optional) Specifies the new real group ID that executes with the command or CDE action.

smexec(1M)

<code>-g egid</code>	(Optional) Specifies the new effective group ID that executes with the command or CDE action.
<code>-h</code>	(Optional) Displays the command's usage statement.
<code>-L label</code>	(Optional) Specifies the the human-readable string or hex representation of the label. It is a valid option when the <code>tsol</code> policy is specified.
<code>-n profile_name</code>	Specifies the name of the profile associated with the <code>exec_attr</code> entry.
<code>-P priv_to_add1...</code>	Specifies the privilege name(s) or privilege number(s) to add to the modified <code>exec_attr</code> entry. Additional privileges may be specified by specifying the <code>-P</code> multiple times. It is a valid option when the <code>tsol</code> policy is specified.
<code>-p policy</code>	Specifies the policy (<code>tsol</code> or <code>suser</code>) associated with the new <code>exec_attr</code> entry. If this option is not specified, the default is <code>suser</code> .
<code>-R priv_to_delete1...</code>	Specifies the privilege name(s) or privilege number(s) to delete from the <code>exec_attr</code> entry. Additional privileges may be specified by specifying the <code>-R</code> multiple times. It is a valid option when the <code>tsol</code> policy is specified.
<code>-t type</code>	Specifies the type <code>cmd</code> for command, or type <code>act</code> for CDE action.
<code>-U uid</code>	(Optional) Specifies the new real user ID that executes with the command or CDE action.
<code>-u euid</code>	(Optional) Specifies the new effective user ID that executes with the command or CDE action.

EXAMPLES

EXAMPLE 1 Adding an `exec_attr` database entry

The admin role connects to port 898 (which happens to also be the default) of the aviary server on the `nis:/birds/aves.Sun.COM` domain, and adds a new `exec_attr` entry for the User Manager profile. The entry type is `act` for the CDE action `ReloadApps;*;*;0`. The action has a clearance of Top Secret Able Baker, a label of confidential, and a policy of `tsol`. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smexec add -D nis:/birds/aves.Sun.COM \
-H aviary:898 -- -n "User Manager" -t act -c "ReloadApps;*;*;0" \
-C "TS A B" -L confidential -p tsol
```


EXAMPLE 2 Deleting an exec_attr database entry

The admin role deletes the ReloadResources;*;*;*;0 CDE action entry in the exec_attr database for the User Manager profile. Since no authorization arguments were specified, the administrator connects to port 898 of the local host on the local server with the file domain type, which are the defaults. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smexec delete -- -n "User Manager" -p tsol \
-t act -c "ReloadResources;*;*;*;0"
```

EXAMPLE 3 Modifying an exec_attr database Entry

The admin role modifies the attributes of the exec_attr database entry for the User Manager profile. The ReloadApps;*;*;*;0 CDE action entry is modified to execute with a clearance of Secret Able. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smexec modify -- -n "User Manager" -p tsol \
-t act -c "ReloadApps;*;*;*;0" -C "S A"
```

**ENVIRONMENT
VARIABLES**

See environ(5) for a description of the JAVA_HOME environment variable, which affects the execution of the smexec command. If this environment variable is not specified, the /usr/java location is used. See smc(1M).

EXIT STATUS

The following exit values are returned:

- | | |
|---|---|
| 0 | Successful completion. |
| 1 | Invalid command syntax. A usage message displays. |
| 2 | An error occurred while executing the command. An error message displays. |

FILES

The following file is used by the smexec command:

/etc/security/exec_attr	Execution profiles database. See exec_attr(4).
-------------------------	--

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmga

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**

To add, modify, or delete an entry in the exec_attr database, the administrator must have the solaris.profmgr.execattr.write authorization.

smexec(1M)

The -C, -L, and -P options may be specified for the add and modify subcommands. The -p option may be specified for the add, modify, and delete subcommands. Input for a CDE action may be specified with most options.

**Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual**

smc(1M), smprofile(1M), smrole(1M), exec_attr(4)
attributes(5), environ(5)

NAME	smgroup – manage group entries								
SYNOPSIS	<code>/usr/sadm/bin/smgroup subcommand [auth_args] - - [subcommand_args]</code>								
DESCRIPTION	The smgroup command manages one or more group definitions in the group database for the appropriate files in the local /etc files name service or a NIS or NIS+ name service.								
subcommands	<p>smgroup subcommands are:</p> <table> <tr> <td>add</td><td>Adds a new group entry. To add an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.</td></tr> <tr> <td>delete</td><td>Deletes a group entry. You can delete only one entry at a time. To delete an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization. <i>Note:</i> You cannot delete the system groups with IDs less than 100, or the groups 60001, 60002, or 65534.</td></tr> <tr> <td>list</td><td>Lists one or more group entries in the form of a three-column list, containing the group name, group ID, and group members, separated by colons (:). To list entries, the administrator must have the <code>solaris.admin.usermgr.read</code> authorization.</td></tr> <tr> <td>modify</td><td>Modifies a group entry. To modify an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.</td></tr> </table>	add	Adds a new group entry. To add an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.	delete	Deletes a group entry. You can delete only one entry at a time. To delete an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization. <i>Note:</i> You cannot delete the system groups with IDs less than 100, or the groups 60001, 60002, or 65534.	list	Lists one or more group entries in the form of a three-column list, containing the group name, group ID, and group members, separated by colons (:). To list entries, the administrator must have the <code>solaris.admin.usermgr.read</code> authorization.	modify	Modifies a group entry. To modify an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.
add	Adds a new group entry. To add an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.								
delete	Deletes a group entry. You can delete only one entry at a time. To delete an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization. <i>Note:</i> You cannot delete the system groups with IDs less than 100, or the groups 60001, 60002, or 65534.								
list	Lists one or more group entries in the form of a three-column list, containing the group name, group ID, and group members, separated by colons (:). To list entries, the administrator must have the <code>solaris.admin.usermgr.read</code> authorization.								
modify	Modifies a group entry. To modify an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.								
OPTIONS	<p>The smgroup authentication arguments, <i>auth_args</i>, are derived from the smc(1M) arg set and are the same regardless of which subcommand you use. The smgroup command requires the SMC to be initialized for the command to succeed (see smc(1M)). After rebooting the SMC server, the first SMC connection might time out, so you might need to retry the command.</p> <p>The subcommand-specific options, <i>subcommand_args</i>, must come after the <i>auth_args</i> and must be separated from them by the - - option.</p>								
auth_args	<p>The valid <i>auth_args</i> are -D, -H, -l, -p, -r, and -u; they are all optional. If no <i>auth_args</i> are specified, certain defaults will be assumed and the user may be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either -D or - -domain.</p> <p>-D - -domain <i>domain</i></p> <p>Specifies the default domain that you want to manage. The syntax of <i>domain</i> is <i>type</i>:/<i>host_name</i>/<i>domain_name</i>, where <i>type</i> is nis, nisplus, dns, ldap or file; <i>host_name</i> is the name of the machine that serves the domain; and <i>domain_name</i> is the name of the domain you want to manage. (<i>Note:</i> Do not use nis+ for nisplus.)</p> <p>If you do not specify this option, the SMC assumes the file default domain on whatever server you choose to manage, meaning that changes are local to the</p>								

smgroup(1M)

server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.

-H | - --hostname *host_name:port*

Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the SMC connects to the local host on port 898. You may still have to choose a toolbox to load into the console. To override this behavior, use the smc(1M) -B option, or set your console preferences to load a “home toolbox” by default.

-l | - --rolepassword *role_password*

Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-p | - --password *password*

Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-r | - --rolename *role_name*

Specifies a role name for authentication. If you do not specify this option, no role is assumed.

-u | - --username *user_name*

Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.

- -

This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the - - option.

subcommand_args

Note: Descriptions and other arg options that contain white spaces must be enclosed in double quotes.

■ For subcommand add:

-g *gid*

(Optional) Specifies the group ID for the new group. The group ID must be a non-negative decimal integer with a maximum value of 2MB (2,147,483,647). Group IDs 0–99 are reserved for the system and should be used with care. If you do not specify a *gid*, the system automatically assigns the next available *gid*. To maximize interoperability and compatibility, administrators are recommended to assign groups using the range of GIDs below 60000 where possible.

-h

(Optional) Displays the command’s usage statement.

- m *group_member1* -m *group_member2* . . .
(Optional) Specifies the new members to add to the group.
- n *group_name*
Specifies the name of the new group. The group name must be unique within a domain, contain 2–32 alphanumeric characters, begin with a letter, and contain at least one lowercase letter.
- For subcommand **delete**:
 - h (Optional) Displays the command's usage statement.
 - n *group_name* Specifies the name of the group you want to delete.
- For subcommand **list**:
 - h (Optional) Displays the command's usage statement.
 - n *group_name* (Optional) Specifies the name of the group you want to list. If you do not specify a group name, all groups are listed.
- For subcommand **modify**:
 - h (Optional) Displays the command's usage statement.
 - m *group_member1* -m *group_member2* . . .
(Optional) Specifies the new members to add to the group.
 - n *group_name*
Specifies the name of the group you want to modify.
 - N *new_group*
(Optional) Specifies the new group name. The group name must be unique within a domain, contain 2–32 alphanumeric characters, begin with a letter, and contain at least one lowercase letter.

EXAMPLES**EXAMPLE 1** Creating a test group

The following creates the `test_group` group entry with a group ID of 123 and adds `test_member1` and `test_member2` to the group:

```
./smgroup add -H myhost -p mypasswd -u root -- -n test_group \
-m test_member1 -m test_member2 -g 123
```

EXAMPLE 2 Deleting a group

The following deletes `test_group`:

```
./smgroup delete -H myhost -p mypasswd -u root -- -n test_group
```

EXAMPLE 3 Displaying all groups

The following displays all groups in a three-column list showing the group name, group ID, and group members:

smgroup(1M)

EXAMPLE 3 Displaying all groups (Continued)

```
./smgroup list -H myhost -p mypasswd -u root --
```

EXAMPLE 4 Displaying a group

The following displays the `group_1` data in a three-column list showing the group name, group ID, and group members:

```
./smgroup list -H myhost -p mypasswd -u root -- -n group_1
```

EXAMPLE 5 Renaming a group

The following renames a group from `finance` to `accounting`:

```
./smgroup modify -H myhost -p mypasswd -u root -- \
-n finance -N accounting
```

ENVIRONMENT VARIABLES

See `environ(5)` for a description of the `JAVA_HOME` environment variable, which affects the execution of the `smgroup` command. If this environment variable is not specified, the `/usr/java` location is used. See `smc(1M)`.

EXIT STATUS

The following exit values are returned:

- | | |
|---|---|
| 0 | Successful completion. |
| 1 | Invalid command syntax. A usage message displays. |
| 2 | An error occurred while executing the command. An error message displays. |

FILES

The following files are used by the `smgroup` command:

`/etc/group` Group file. See `group(4)`.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmga

SEE ALSO

`smc(1M)`, `group(4)`, `attributes(5)`, `environ(5)`

NAME	smhost – Manage entries in the hosts database								
SYNOPSIS	<code>/usr/sadm/bin/smhost subcommand [auth_args] - - [subcommand_args]</code>								
DESCRIPTION	<p>The smhost command adds, modifies, deletes, and lists entries in the hosts, ethers, and tnrdhdb databases.</p> <p>smhost subcommands are:</p> <table> <tr> <td>add</td><td>Adds a new entry to the hosts database. To add an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.</td></tr> <tr> <td>modify</td><td>Modifies an entry in the hosts database. To modify an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.</td></tr> <tr> <td>delete</td><td>Deletes an entry from the hosts database. To delete an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.</td></tr> <tr> <td>list</td><td>Lists entries in the specified database. To list an entry, the administrator must have the <code>solaris.network.host.read</code> and <code>solaris.network.security.read</code> authorizations.</td></tr> </table>	add	Adds a new entry to the hosts database. To add an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.	modify	Modifies an entry in the hosts database. To modify an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.	delete	Deletes an entry from the hosts database. To delete an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.	list	Lists entries in the specified database. To list an entry, the administrator must have the <code>solaris.network.host.read</code> and <code>solaris.network.security.read</code> authorizations.
add	Adds a new entry to the hosts database. To add an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.								
modify	Modifies an entry in the hosts database. To modify an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.								
delete	Deletes an entry from the hosts database. To delete an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.								
list	Lists entries in the specified database. To list an entry, the administrator must have the <code>solaris.network.host.read</code> and <code>solaris.network.security.read</code> authorizations.								
OPTIONS	<p>The smhost authentication arguments, <i>auth_args</i>, are derived from the smc(1M) arg set and are the same regardless of which subcommand you use. The smhost command requires the SMC to be initialized for the command to succeed (see smc(1M)). After rebooting the SMC server, the first smc connection may time out, so you may need to retry the command.</p> <p>The subcommand-specific options, <i>subcommand_args</i>, must be <i>preceded</i> by the - - option.</p>								
<i>auth_args</i>	<p>The valid <i>auth_args</i> are -D, -H, -l, -p, -r, and -u; they are all optional. If no <i>auth_args</i> are specified, certain defaults will be assumed and the user may be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either -D or - -domain.</p> <p>-D - -domain <i>domain</i></p> <p>Specifies the default domain that you want to manage. The syntax of <i>domain=type:/host_name/domain_name</i>, where <i>type</i> is nis, nisplus, dns, ldap, or file; <i>host_name</i> is the name of the machine that serves the domain; and <i>domain_name</i> is the name of the domain you want to manage. (Note: Do not use nis+ for nisplus.)</p> <p>If you do not specify this option, the SMC assumes the file default domain on whatever server you choose to manage, meaning that changes are local to the server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.</p>								

smhost(1M)

- H | - -hostname *host_name:port*
Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the SMC connects to the local host on port 898.
- l | - -rolepassword *role_password*
Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.
- p | - -password *password*
Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.
- r | - -rolename *role_name*
Specifies a role name for authentication. If you do not specify this option, no role is assumed.
- u | - -username *user_name*
Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.
- -
This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the - - option.

subcommand_args

Note: Descriptions and other arg options that contain white spaces must be enclosed in double quotes.

- a *aliases* Specifies the aliases by which this hostname may be identified.
- d *description* Specifies the description for this host entry.
- e *ethernetaddress* Specifies the Ethernet address of the host. For the `list` subcommand, the *ethernetaddress* argument is not specified, and the `ethers` database is listed.
- H *hostname* Specifies the name of the host. For the `list` subcommand, the *hostname* argument is not specified, and the `hosts` database is listed.
- h Displays the command's usage statement.
- i *ipaddress* Specifies the IP address of the host.
- n *templatename* Specifies the template name. For the `list` subcommand, the *templatename* argument is not specified, and the `tnrhdb` database is listed.

- p *prefixlen*** Specifies the prefix length (in bits) of a wildcard representation of the IP address. The prefix is the left-most portion of the IP address.
- w *ipaddress_wildcard*** Specifies the IP address of the subnet using a wildcard.
- One of the following sets of arguments must be specified for subcommand **add**:
 - H *hostname* -i *ipaddress* [-a *aliases*] [-d *description*] [-e *ethernetaddress*] [-n *templatename*] |**
 - H *hostname* -n *templatename* |**
 - i *ipaddress* -n *templatename* |**
 - w *ipaddress_wildcard* -n *templatename* [-p *prefixlen*] |**
 - h**
 - One of the following sets of arguments must be specified for subcommand **modify**:
 - H *hostname* { [-i *ipaddress*] [-a *aliases*] [-d *description*] [-e *ethernetaddress*] [-n *templatename*] } |**
 - H *hostname* -n *templatename* |**
 - w *ipaddress_wildcard* { [-n *templatename*] [-p *prefixlen*] } |**
 - h**
 - One of the following sets of arguments must be specified for subcommand **delete**:
 - H *hostname* |**
 - i *ipaddress* |**
 - w *ipaddress_wildcard* [-p *prefixlen*] |**
 - h**
 - One of the following sets of arguments must be specified for subcommand **list**:
 - H |**
 - n |**
 - e |**
 - h**

EXAMPLES**EXAMPLE 1** Adding a new entry to the hosts database

The admin role connects to port 898 (which happens to be the default) of the aviary server on the `nis:/birds/aves.Sun.COM` domain, and creates a new host, `eagle`, with an IP address of `129.150.11.7`. The new host is assigned to the `tsol` template. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smhost add -D nis:/birds/aves.Sun.COM -H aviary:898 -- \
-H eagle -i 129.150.11.7 -n tsol
```

smhost(1M)

EXAMPLE 2 Specifying the template name using an IP address wildcard

The admin role specifies the template name, `tsol`, for series of hosts using the IP address wildcard `129.150.11.0` on the local file system. Since no authorization arguments were specified, the administrator connects to port 898 of the local host on the local server with the `file` domain type, which are the defaults. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smhost add -- -w 129.150.11.0 -n tsol
```

EXAMPLE 3 Modifying an entry in the hosts database

The admin role connects to the `nis:/birds/aves.Sun.COM` domain and modifies the `eagle` entry in the `hosts` database, changing its IP address to `129.150.11.8` and its alias to `falcon`. Since the host and port were not specified, the local host and port 898 are used by default. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smhost modify -D nis:/birds/aves.Sun.COM -- \
-H eagle -i 129.150.11.8 -a falcon
```

EXAMPLE 4 Deleting an entry in the hosts database

The admin role connects to port 898 (which happens to be the default) of the aviary server and deletes the `eagle` entry in the `hosts` database by specifying its IP address, `129.150.11.8`. Since the domain was not specified, the `file` domain type and local server are used by default. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smhost delete -H aviary:898 -- -i 129.150.11.8
```

EXAMPLE 5 Listing the `tnrhdb` database

The admin role connects to the aviary server on the `nis:/birds/aves.Sun.COM` domain and lists the entries in the `tnrhdb` database. Since the port was not specified, port 898 is used by default. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smhost list -D nis:/birds/aves.Sun.COM -H aviary -- \
-n
```

EXIT STATUS

The following exit values are returned:

- | | |
|---|---|
| 0 | Successful completion. |
| 1 | Invalid command syntax. A usage message displays. |
| 2 | An error occurred while executing the command. An error message displays. |

FILES

The following files are used by the `smhost` command:

<code>/etc/ethers</code>	Ethernet address to hostname database or domain. See <code>ethers(4)</code> .
--------------------------	---

smhost(1M)

- /etc/hosts Host name database. See hosts(4).
- /etc/security/tsol/tnrhdb Trusted network remote-host database. See tnrhdb(4).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmgapp

- smc(1M), tnrhdb(4)
- ethers(4), hosts(4), attributes(5)

smmaillist(1M)

NAME	smmaillist – manage email alias entries								
SYNOPSIS	<code>/usr/sadm/bin/smmaillist subcommand [auth_args] - - [subcommand_args]</code>								
DESCRIPTION	The smmaillist command manages one or more email alias entries for the appropriate files in the local /etc files name service or a NIS or NIS+ name service.								
<i>subcommands</i>	<p>smmaillist <i>subcommands</i> are:</p> <table> <tr> <td>add</td><td>Creates a new email alias definition and adds it to the appropriate files. To add an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.</td></tr> <tr> <td>delete</td><td>Deletes an email alias entry. You can delete only one entry at a time. <i>Note:</i> You cannot delete Postmaster or Mailer-Daemon aliases. To delete an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.</td></tr> <tr> <td>list</td><td>Lists one or more email alias entries. To list an entry, the administrator must have the <code>solaris.admin.usermgr.read</code> authorization.</td></tr> <tr> <td>modify</td><td>Modifies an email alias entry. To modify an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.</td></tr> </table>	add	Creates a new email alias definition and adds it to the appropriate files. To add an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.	delete	Deletes an email alias entry. You can delete only one entry at a time. <i>Note:</i> You cannot delete Postmaster or Mailer-Daemon aliases. To delete an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.	list	Lists one or more email alias entries. To list an entry, the administrator must have the <code>solaris.admin.usermgr.read</code> authorization.	modify	Modifies an email alias entry. To modify an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.
add	Creates a new email alias definition and adds it to the appropriate files. To add an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.								
delete	Deletes an email alias entry. You can delete only one entry at a time. <i>Note:</i> You cannot delete Postmaster or Mailer-Daemon aliases. To delete an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.								
list	Lists one or more email alias entries. To list an entry, the administrator must have the <code>solaris.admin.usermgr.read</code> authorization.								
modify	Modifies an email alias entry. To modify an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.								
OPTIONS	<p>The smmaillist authentication arguments, <i>auth_args</i>, are derived from the smc(1M) arg set and are the same regardless of which subcommand you use. The smmaillist command requires the SMC to be initialized for the command to succeed (see smc(1M)). After rebooting the SMC server, the first smc connection may time out, so you may need to retry the command.</p> <p>The subcommand-specific options, <i>subcommand_args</i>, must come after the <i>auth_args</i> and must be separated from them by the - - option.</p>								
<i>auth_args</i>	<p>The valid <i>auth_args</i> are -D, -H, -l, -p, -r, and -u; they are all optional. If no <i>auth_args</i> are specified, certain defaults will be assumed and the user may be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either -D or - --domain with the <i>domain</i> argument.</p> <p>-D - --domain <i>domain</i></p> <p>Specifies the default domain that you want to manage. The syntax of <i>domain</i> is <i>type:/host_name/domain_name</i>, where <i>type</i> is nis, nisplus, dns, ldap, or file; <i>host_name</i> is the name of the machine that serves the domain; and <i>domain_name</i> is the name of the domain you want to manage. (<i>Note:</i> Do not use nis+ for nisplus.)</p> <p>If you do not specify this option, the SMC assumes the file default domain on whatever server you choose to manage, meaning that changes are local to the</p>								

server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.

-H | **-** **-hostname** *host_name:port*

Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the SMC connects to the local host on port 898. You may still have to choose a toolbox to load into the console. To override this behavior, use the `smc(1M) -B` option, or set your console preferences to load a “home toolbox” by default.

-l | **-** **-rolepassword** *role_password*

Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-p | **-** **-password** *password*

Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-r | **-** **-rolename** *role_name*

Specifies a role name for authentication. If you do not specify this option, no role is assumed.

-u | **-** **-username** *user_name*

Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.

- **-**

This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the **- -** option.

subcommand_args

Note: Descriptions and other arg options that contain white spaces must be enclosed in double quotes.

■ For subcommand **add**:

-a *address1* **-a** *address2* . . . (Optional) Specifies the new email address. See `sendmail(1M)`.

-h (Optional) Displays the command’s usage statement.

-n *alias_name* Specifies the name of the alias you want to add. See `sendmail(1M)`.

■ For subcommand **delete**:

-h (Optional) Displays the command’s usage statement.

-n *alias_name* Specifies the alias you want to delete.

■ For subcommand **list**:

smmaillist(1M)

- h (Optional) Displays the command's usage statement.
- n *alias_name* (Optional) Specifies the name of the alias you want to display. If you do not specify an alias, all aliases are listed.
- For subcommand modify:
 - a *address1* -a *address2* ... (Optional) Specifies new email address(es) to replace the existing one(s). See `sendmail(1M)`.
 - h (Optional) Displays the command's usage statement.
 - n *alias_name* (Optional) Specifies the name of the alias you want to modify.
 - N *new_alias_name* Specifies the new alias name. Use only when renaming an alias. See `sendmail(1M)`.

EXAMPLES

EXAMPLE 1 Adding an alias

The admin role adds the `bandmembers` alias and adds the following member list: `mitsuru@sun.com`, `ichiro@sun.com`, and `kaori@sun.com` to the alias. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smmaillist add -- -n bandmembers \  
-a mitsuru@sun.com -a ichiro@sun.com -a kaori@sun.com
```

EXAMPLE 2 Deleting a mail alias

The admin role deletes the `bandmembers` alias. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smmaillist delete -- -n bandmembers
```

EXAMPLE 3 Listing members of a mail alias

The admin role lists all members belonging to the `bandmembers` alias. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smmaillist list -- -n bandmembers
```

EXAMPLE 4 Listing members of all mail aliases

The admin role lists all members belonging to all mail aliases. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smmaillist list --
```

EXAMPLE 5 Renaming a mail alias

The admin role renames the `bandmembers` mail alias to `groupmembers`. The administrator is prompted for the admin password.

EXAMPLE 5 Renaming a mail alias (Continued)

```
$ /usr/sadm/bin/smmaillist modify -- \
-n bandmembers -N groupmembers
```

EXAMPLE 6 Redefining an address list

The admin role changes the recipients of the alias `bandmembers` to `kaori@sun.com`. Any previous recipients are deleted from the alias. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smmaillist modify -- -n bandmembers \
-a kaori@sun.com
```

**ENVIRONMENT
VARIABLES**

See `environ(5)` for a description of the `JAVA_HOME` environment variable, which affects the execution of the `smmaillist` command. If this environment variable is not specified, the `/usr/java` location is used. See `smc(1M)`.

EXIT STATUS

The following exit values are returned:

- 0 Successful completion.
- 1 Invalid command syntax. A usage message displays.
- 2 An error occurred while executing the command. An error message displays.

FILES

The following files are used by the `smmaillist` command:

`/var/mail/aliases` Aliases for `sendmail(1M)`. See `aliases(4)`.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmga

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**

To add, delete, or modify an entry, the administrator must have the `solaris.admin.usermgr.write` authorization. To list an entry, the administrator must have the `solaris.admin.usermgr.read` authorization.

Trusted Solaris 8
4/01 Reference
Manual

`sendmail(1M)`, `smc(1M)`
`aliases(4)`, `attributes(5)`, `environ(5)`

smmultiuser(1M)

NAME	smmultiuser – manage bulk operations on user accounts						
SYNOPSIS	<code>/usr/sadm/bin/smmultiuser subcommand [auth_args] - - [subcommand_args]</code>						
DESCRIPTION	<p>The <code>smmultiuser</code> command allows bulk operations on user entries in the local <code>/etc</code> filesystem or a NIS or NIS+ name service, using either an input file or piped input.</p> <p><i>Note:</i> Both input files and piped input contain a cleartext (non-encrypted) password for each new user entry.</p> <p>Note – The <code>smmultiuser</code> command does not accept Trusted Solaris security attributes, such as labels and clearances. Use <code>smuser(1M)</code> to create a user with Trusted Security attributes or to add such attributes to an existing user.</p>						
<i>subcommands</i>	<p><code>smmultiuser</code> subcommands are:</p> <table> <tr> <td><code>add</code></td><td>Adds multiple user entries to the appropriate files. To add an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.</td></tr> <tr> <td><code>delete</code></td><td>Deletes one or more user entries from the appropriate files. To delete an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.</td></tr> <tr> <td><code>modify</code></td><td> <p>Modifies existing user entries in the user account database. Here is the list of what can be modified using the <code>modify</code> subcommand:</p> <ol style="list-style-type: none"> 1. <code>UserName</code> (only under certain conditions — see Note 2 in NOTES) 2. <code>Password</code> (only under certain conditions — see Note 3 in NOTES) 3. <code>Description</code> 4. <code>Primary Group ID</code> 5. <code>Shell type</code> 6. <code>FullName</code> To modify an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization. </td></tr> </table>	<code>add</code>	Adds multiple user entries to the appropriate files. To add an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.	<code>delete</code>	Deletes one or more user entries from the appropriate files. To delete an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.	<code>modify</code>	<p>Modifies existing user entries in the user account database. Here is the list of what can be modified using the <code>modify</code> subcommand:</p> <ol style="list-style-type: none"> 1. <code>UserName</code> (only under certain conditions — see Note 2 in NOTES) 2. <code>Password</code> (only under certain conditions — see Note 3 in NOTES) 3. <code>Description</code> 4. <code>Primary Group ID</code> 5. <code>Shell type</code> 6. <code>FullName</code> To modify an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.
<code>add</code>	Adds multiple user entries to the appropriate files. To add an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.						
<code>delete</code>	Deletes one or more user entries from the appropriate files. To delete an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.						
<code>modify</code>	<p>Modifies existing user entries in the user account database. Here is the list of what can be modified using the <code>modify</code> subcommand:</p> <ol style="list-style-type: none"> 1. <code>UserName</code> (only under certain conditions — see Note 2 in NOTES) 2. <code>Password</code> (only under certain conditions — see Note 3 in NOTES) 3. <code>Description</code> 4. <code>Primary Group ID</code> 5. <code>Shell type</code> 6. <code>FullName</code> To modify an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization. 						
OPTIONS	<p>The <code>smmultiuser</code> authentication arguments, <i>auth_args</i>, are derived from the <code>smc(1M)</code> arg set and are the same regardless of which subcommand you use. The <code>smmultiuser</code> command requires the SMC to be initialized for the command to succeed (see <code>smc(1M)</code>). After rebooting the SMC server, the first <code>smc</code> connection may time out, so you may need to retry the command.</p> <p>The subcommand-specific options, <i>subcommand_args</i>, must come after the <i>auth_args</i> and must be separated from them by the <code>- -</code> option.</p>						
<i>auth_args</i>	<p>The valid <i>auth_args</i> are <code>-D</code>, <code>-H</code>, <code>-l</code>, <code>-p</code>, <code>-r</code>, <code>- -trust</code>, and <code>-u</code>; they are all optional. If no <i>auth_args</i> are specified, certain defaults will be assumed and the user may be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either <code>-D</code> or <code>- -domain</code>.</p>						

-D | **-** **-domain** *domain*

Specifies the default domain that you want to manage. The syntax of *domain* is *type:/host_name/domain_name*, where *type* is *nis*, *nisplus*, *dns*, *ldap*, or *file*; *host_name* is the name of the machine that serves the domain; and *domain_name* is the name of the domain you want to manage. (Note: Do not use *nis+* for *nisplus*.)

If you do not specify this option, the SMC assumes the *file* default domain on whatever server you choose to manage, meaning that changes are local to the server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.

-H | **-** **-hostname** *host_name:port*

Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the SMC connects to the local host on port 898. You may still have to choose a toolbox to load into the console. To override this behavior, use the *smc(1M)* **-B** option, or set your console preferences to load a "home toolbox" by default.

-l | **-** **-rolepassword** *role_password*

Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-p | **-** **-password** *password*

Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-r | **-** **-rolename** *role_name*

Specifies a role name for authentication. If you do not specify this option, no role is assumed.

- **-trust**

Trusts all downloaded code implicitly. Use this option when running the terminal console non-interactively and you cannot let the console wait for user input.

If using piped input into any of the *smmultiuser* subcommands, it will now be necessary to use the **-** **-trust** option with the **-L** *logfile* option. See EXAMPLES.

-u | **-** **-username** *user_name*

Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.

- **-**

This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the **-** **-** option.

subcommand_args

Note: Descriptions and other arg options that contain white spaces must be enclosed in double quotes.

- For subcommand **add**:
 - h (Optional) Displays the command's usage statement.
 - i *input_file* Specifies the input file containing the user account information. After the command is executed, the input file is removed. The input file must follow the `/etc/passwd` file format. If you do not specify the -i *input_file* option, you must include a *pipedReader* operand immediately before the command. See EXAMPLES.
 - L *logfile* (Optional) Specifies the full pathname to the text file that stores the command's success/failure data. *Note*: This text file is an ASCII—formatted log file; it is different from and unrelated to the output of the normal logging mechanism that also occurs within the Log Viewer tool. The -L *logfile* option is used to dump additional logging information to a text file.
- For subcommand **delete**:
 - h (Optional) Displays the command's usage statement.
 - i *input_file* Specifies the input file containing the user account information. After the command is executed, the input file is removed. The input file must follow the `/etc/passwd` file format. If you do not specify the -i *input_file* option, you must include a *pipedReader* operand immediately before the command. See EXAMPLES.
 - L *logfile* (Optional) Specifies the full pathname to the text file that stores the command's success/failure data.
- For subcommand **modify**:
 - h (Optional) Displays the command's usage statement.
 - i *input_file* Specifies the input file containing the user account information. After the command is executed, the input file is removed. The input file must follow the `/etc/passwd` file format. If you do not specify the -i *input_file* option, you must include a *pipedReader* operand immediately before the command. See EXAMPLES. *Note*: When modifying passwords, use the pipedReader input, since it is more secure than keeping passwords in a file. See Note 1 in NOTES.
 - L *logfile* (Optional) Specifies the full pathname to the text file that stores the command's success/failure data.

OPERANDS

The following operands are supported:

- pipedReader* You must include *pipedReader* if you do not specify an *input_file*. Include the pipedReader input immediately before the command. The pipedReader input must follow the `/etc/passwd` file format. See EXAMPLES. *Note*: Use the - `-trust` option when using pipedReader input with the -L *logfile* option to avoid the user prompt from the Security Alert Manager, which normally asks the user whether the log file should be created. Without the - `-trust` option, the

piped input is improperly taken as the answer to the prompt before the user can answer “Y” or “N”, and the logging operation will probably fail.

EXAMPLES

EXAMPLE 1 Adding multiple user accounts

The admin role reads in user account data from the `/tmp/foo` file and creates new user accounts. The input file is formatted in the `/etc/passwd` format. The administrator is prompted for the admin password.

```

$ /usr/sadm/bin/smmultiuser add -- -i /tmp/foo

```

EXAMPLE 2 Deleting multiple user accounts

The admin role reads in user account data from the `/tmp/foo` file and deletes the named user accounts. The administrator is prompted for the admin password.

```

$ /usr/sadm/bin/smmultiuser delete -- -i /tmp/foo

```

EXAMPLE 3 Creating a log file with piped input

The following example shows the use of the `smc(1M)` - `-trust` option that is required when creating a log file. It is applicable to the `delete` and `modify` subcommands also.

```

$ cat /tmp/users.txt | /usr/sadm/bin/smmultiuser add --trust -- \
-L /tmp/mylog.txt

```

ENVIRONMENT VARIABLES

See `environ(5)` for a description of the `JAVA_HOME` environment variable, which affects the execution of the `smprofile` command. If this environment variable is not specified, the `/usr/java` location is used. See `smc(1M)`.

EXIT STATUS

The following exit values are returned:

- 0 Successful completion.
- 1 Invalid command syntax. A usage message displays.
- 2 An error occurred while executing the command. An error message displays.

FILES

The following files are used by the `smprofile` command:

`/etc/passwd` Contains the file format to use for the *input_file* and *pipd_input*. See `passwd(4)`.

ATTRIBUTES

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

smmultiuser(1M)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmga

SUMMARY OF TRUSTED SOLARIS CHANGES

The `smmultiuser` command does not accept Trusted Solaris security attributes, such as labels and clearances. Use `smuser(1M)` to create a user with Trusted Security attributes or to add such attributes to an existing user.

To add, delete, or modify an entry, the administrator must have the `solaris.admin.usermgr.write` authorization.

Trusted Solaris 8 4/01 Reference Manual NOTES

`smc(1M)`, `smrole(1M)`, `smuser(1M)`
`passwd(4)`, `attributes(5)`, `environ(5)`

1. The file format used by both the `add` and `modify` subcommands is the `/etc/passwd` format. But there is an allowance for a mutated version of this file format that contains an extra field at the end of each line to be used for the Full Name. If the extra field is appended to the end of each line, it will be used for the Full Name value, but if it is omitted, it will be assumed that no FullName modification is being done. The extra field is separated with a colon (:), just like all the other fields.

Example of regulation `/etc/passwd` entry:

`rick2:x:1011:10:description1:/home/rick2:/bin/sh` Example of `/etc/passwd` variant entry:

`rick2:x:1011:10:description1:/home/rick2:/bin/sh:Ricks_fullname`

2. The modifications are all done based on lookups of the user name in the user tables. If a user name cannot be found in this lookup, a secondary check is made to see if the `uid` *and* FullName can be found in the user tables. If they are both found, the user is renamed. If neither can be found, no modification occurs.
3. If no password is supplied, the password is not changed. If a password is being changed, it should be supplied in cleartext as piped input or in the input file. Once read in, the password is changed accordingly.

NAME	smnetidb – Manage entries in the interface database								
SYNOPSIS	<code>/usr/sadm/bin/smnetidb subcommand [auth_args] - - [subcommand_args]</code>								
DESCRIPTION	<p>The smnetidb command adds, modifies, deletes, and lists entries in the tnidb database.</p> <p>smnetidb subcommands are:</p> <table> <tr> <td>add</td><td>Adds a new entry to the tnidb database. To add an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.</td></tr> <tr> <td>modify</td><td>Modifies an entry in the tnidb database. To modify an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.</td></tr> <tr> <td>delete</td><td>Deletes an entry from the tnidb database. To delete an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.</td></tr> <tr> <td>list</td><td>Lists entries in the tnidb database. To list an entry, the administrator must have the <code>solaris.network.host.read</code> and <code>solaris.network.security.read</code> authorizations.</td></tr> </table>	add	Adds a new entry to the tnidb database. To add an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.	modify	Modifies an entry in the tnidb database. To modify an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.	delete	Deletes an entry from the tnidb database. To delete an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.	list	Lists entries in the tnidb database. To list an entry, the administrator must have the <code>solaris.network.host.read</code> and <code>solaris.network.security.read</code> authorizations.
add	Adds a new entry to the tnidb database. To add an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.								
modify	Modifies an entry in the tnidb database. To modify an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.								
delete	Deletes an entry from the tnidb database. To delete an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.								
list	Lists entries in the tnidb database. To list an entry, the administrator must have the <code>solaris.network.host.read</code> and <code>solaris.network.security.read</code> authorizations.								
OPTIONS	<p>The smnetidb authentication arguments, <i>auth_args</i>, are derived from the smc(1M) arg set and are the same regardless of which subcommand you use. The smnetidb command requires the SMC to be initialized for the command to succeed (see smc(1M)). After rebooting the SMC server, the first smc connection may time out, so you may need to retry the command.</p> <p>The subcommand-specific options, <i>subcommand_args</i>, must be <i>preceded</i> by the - - option.</p>								
auth_args	<p>The valid <i>auth_args</i> are -D, -H, -l, -p, -r, and -u; they are all optional. If no <i>auth_args</i> are specified, certain defaults will be assumed and the user may be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either -D or - -domain.</p> <p>-D - -domain <i>domain</i></p> <p>Specifies the default domain that you want to manage. The syntax of <i>domain=type:/host_name/domain_name</i>, where <i>type</i> is <code>nis</code>, <code>nisplus</code>, <code>dns</code>, <code>ldap</code>, or <code>file</code>; <i>host_name</i> is the name of the machine that serves the domain; and <i>domain_name</i> is the name of the domain you want to manage. (Note: Do not use <code>nis+</code> for <code>nisplus</code>.)</p> <p>If you do not specify this option, the SMC assumes the <code>file</code> default domain on whatever server you choose to manage, meaning that changes are local to the server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.</p>								

smnetidb(1M)

- H | - -hostname *host_name:port*
Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the SMC connects to the local host on port 898.
- l | - -rolepassword *role_password*
Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.
- p | - -password *password*
Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.
- r | - -rolename *role_name*
Specifies a role name for authentication. If you do not specify this option, no role is assumed.
- u | - -username *user_name*
Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.
- -
This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the - - option.

subcommand_args

Note: Descriptions and other arg options that contain white spaces must be enclosed in double quotes.

- c *clearance* Specifies the clearance for the interface.
- f *forced_privileges* Specifies the forced privileges for the interface entry. Values can be all, none, empty, or a comma-separated list of privilege names (not privilege numbers).
- h Displays the command's usage statement.
- l *label* Specifies the CMW label for the interface entry.
- n *interfacename* Specifies the name for the interface entry.
- x max=*maximum_label* Specifies the maximum label for the interface entry.
- x min=*minimum_label* Specifies the minimum label for the interface entry.
- One of the following sets of arguments must be specified for subcommand add:
 - n *interfacename* -x min=*minimum_label* -x max=*maximum_label* -l *label* -c *clearance* -f *forced_privileges* |
 - h

- One of the following sets of arguments must be specified for subcommand `modify`:
 - `-n interfacename { [-x min=minimum_label] [-x max=maximum_label] [-l label] [-c clearance] [-f forced_privileges] }` |
 - `-h`
- One of the following arguments must be specified for subcommand `delete`:
 - `-n interfacename |`
 - `-h`
- The following argument may be specified for subcommand `list`:
 - `-h`

EXAMPLES**EXAMPLE 1** Adding a new entry to the interface database

The admin role creates a new interface entry, `1e0`, with a minimum label of `confidential`, maximum label of `top secret`, label of `[secret]`, clearance of `ts a b`, and forced privileges of `all`. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smnetidb add -- -n 1e0 \
-x min=confidential -x max="top secret" -l "[secret]" -c "ts a b" \
-f all
```

EXAMPLE 2 Modifying an entry in the interface database

The user modifies the `1e0` entry in the `tnidb` database, changing its minimum label to `secret` and its forced privileges to `net_mac_read`, `net_reply_equal`, and `net_privaddr`. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smnetidb modify -- -n 1e0 -x min=secret \
-f net_mac_read,net_reply_equal,net_privaddr
```

EXAMPLE 3 Deleting an entry in the interface database

The admin role deletes the `1e0` entry in the `tnidb` database. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smnetidb delete -- -n 1e0
```

EXAMPLE 4 Listing the interface database

The admin role lists the entries in the `tnidb` database. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smnetidb list --
```

EXIT STATUS

The following exit values are returned:

smnetidb(1M)

- 0 Successful completion.
- 1 Invalid command syntax. A usage message displays.
- 2 An error occurred while executing the command. An error message displays.

FILES

The following files are used by the `smnetidb` command:

`/etc/security/tsol/tnidb` Trusted network interface-control database.
See `tnidb(4)`.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmgapp

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

`smc(1M)`, `tnidb(4)`

`attributes(5)`

NAME	smnettpl – manage entries in the network template database								
SYNOPSIS	<code>/usr/sadm/bin/smnettpl subcommand [auth_args] - - [subcommand_args]</code>								
DESCRIPTION	<p>The smnettpl command adds, modifies, deletes, and lists entries in the tnrtcp database.</p> <p>smnettpl subcommands are:</p> <table> <tr> <td>add</td><td>Adds a new entry to the network template database. To add an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.</td></tr> <tr> <td>modify</td><td>Modifies an entry in the network template database. To modify an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.</td></tr> <tr> <td>delete</td><td>Deletes an entry from the network template database. To delete an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.</td></tr> <tr> <td>list</td><td>Lists entries in the network template database. To list an entry, the administrator must have the <code>solaris.network.security.read</code> authorizations.</td></tr> </table>	add	Adds a new entry to the network template database. To add an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.	modify	Modifies an entry in the network template database. To modify an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.	delete	Deletes an entry from the network template database. To delete an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.	list	Lists entries in the network template database. To list an entry, the administrator must have the <code>solaris.network.security.read</code> authorizations.
add	Adds a new entry to the network template database. To add an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.								
modify	Modifies an entry in the network template database. To modify an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.								
delete	Deletes an entry from the network template database. To delete an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.								
list	Lists entries in the network template database. To list an entry, the administrator must have the <code>solaris.network.security.read</code> authorizations.								
OPTIONS	<p>The smnettpl authentication arguments, <i>auth_args</i>, are derived from the smc(1M) arg set and are the same regardless of which subcommand you use. The smnettpl command requires the SMC to be initialized for the command to succeed (see smc(1M)). After rebooting the SMC server, the first smc connection may time out, so you may need to retry the command.</p> <p>The subcommand-specific options, <i>subcommand_args</i>, must be <i>preceded</i> by the - - option.</p>								
<i>auth_args</i>	<p>The valid <i>auth_args</i> are -D, -H, -l, -p, -r, and -u; they are all optional. If no <i>auth_args</i> are specified, certain defaults will be assumed and the user may be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either -D or - -domain.</p> <p>-D - -domain <i>domain</i></p> <p>Specifies the default domain that you want to manage. The syntax of <i>domain=type:/host_name/domain_name</i>, where <i>type</i> is <code>nis</code>, <code>nisplus</code>, <code>dns</code>, <code>ldap</code>, or <code>file</code>; <i>host_name</i> is the name of the machine that serves the domain; and <i>domain_name</i> is the name of the domain you want to manage. (Note: Do not use <code>nis+</code> for <code>nisplus</code>.)</p>								

smnettmpl(1M)

If you do not specify this option, the SMC assumes the `file` default domain on whatever server you choose to manage, meaning that changes are local to the server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.

`-H | - --hostname host_name:port`

Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the SMC connects to the local host on port 898.

`-l | - --rolepassword role_password`

Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

`-p | - --password password`

Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

`-r | - --rolename role_name`

Specifies a role name for authentication. If you do not specify this option, no role is assumed.

`-u | - --username user_name`

Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.

`- -`

This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the `- -` option.

subcommand_args

Note: Descriptions and other arg options that contain white spaces must be enclosed in double quotes.

`-a allowed_privilege`

Specifies the allowed privilege. Values can be a privilege name or number. Multiple privileges must be separated by a comma.

`-c clearance`

Specifies the clearance. Values can be a hex value or string (such as `admin_low`).

`-f forced_privilege`

Specifies the forced privilege. Values can be a privilege name or number. Multiple privileges must be separated by a comma.

`-h`

Displays the command's usage statement.

`-i ipLabelType`

Specifies the IP label type. Valid values are `none`, `ripso`, or `cipso`.

<code>-l label</code>	Specifies the label in CMW label format. Values can be a hex value or string (such as <code>[admin_low]</code>).
<code>-n templatename</code>	Specifies the template name of the template.
<code>-t hosttype=value</code>	Specifies the hosttype of the new host. Valid values are unlabeled, sun_tsol, cipso, ripso, and tsix.
<code>-x DOI=doi_value</code>	Specifies the DOI value.
<code>-x max=maximum_label</code>	Specifies the maximum label. Values can be a hex value or string (such as <code>admin_low</code>).
<code>-x min=minimum_label</code>	Specifies the minimum label. Values can be a hex value or string (such as <code>admin_low</code>).
<code>-x ripsoRPAF=RipsoReturnPAF</code>	Specifies the ripso return PAF. Valid values are GENSER, SIOP-ESI, SCI, NSA, or DOE.
<code>-x ripsoSC=RipsoSendClass</code>	Specifies the ripso send class. Valid values are Top Secret, Secret, Confidential, or Unclassified.
<code>-x ripsoSPAF=RipsoSendPAF</code>	Specifies the ripso send PAF. Valid values are GENSER, SIOP-ESI, SCI, NSA, or DOE.

- One of the following sets of arguments must be specified for subcommand add:

```

-n template name (
    -t hosttype=sun_tsol -x min=minimum_label -x max=maximum_label -a
    allowed_privilege ( -i none | -i ripso -x ripsoSC=RipsoSendClass -x
    ripsoSPAF=RipsoSendPAF -x ripsoRPAF=RipsoReturnPAF | -i cipso ) -x
    DOI=doi_value |
    -t hosttype=unlabeled -x min=minimum_label -x max=maximum_label -l
    label -c clearance -f forced_privilege [ ( -i none | -i ripso -x
    ripsoSC=RipsoSendClass -x ripsoSPAF=RipsoSendPAF -x
    ripsoRPAF=RipsoReturnPAF | -i cipso ) ] -x DOI=doi_value |
    -t hosttype=ripso -x min=minimum_label -x max=maximum_label -l label -c
    clearance -f forced_privilege -x ripsoSC=RipsoSendClass -x
    ripsoSPAF=RipsoSendPAF -x ripsoRPAF=RipsoReturnPAF -x DOI=doi_value
    |
    -t hosttype=cipso -x min=minimum_label -x max=maximum_label -c
    clearance -f forced_privilege -x DOI=doi_value |
    -t hosttype=tsix -x min=minimum_label -x max=maximum_label -a
    allowed_privilege ( -i none | -i ripso -x ripsoSC=RipsoSendClass -x
    ripsoSPAF=RipsoSendPAF -x ripsoRPAF=RipsoReturnPAF | -i cipso ) -x
    DOI=doi_value |
    -h
)

```

- One of the following sets of arguments must be specified for subcommand `modify`:

```
-n template name (
    -t hosttype=sun_tsol -x min=minimum_label -x max=maximum_label -a
    allowed_privilege ( -i none | -i ripso -x ripsoSC=RipsoSendClass -x
    ripsoSPAF=RipsoSendPAF -x ripsoRPAF=RipsoReturnPAF | -i cipso ) -x
    DOI=doi_value |
    -t hosttype=unlabeled -x min=minimum_label -x max=maximum_label -l
    label -c clearance -f forced_privilege [ ( -i none | -i ripso -x
    ripsoSC=RipsoSendClass -x ripsoSPAF=RipsoSendPAF -x
    ripsoRPAF=RipsoReturnPAF | -i cipso ) ] -x DOI=doi_value |
    -t hosttype=ripso -x min=minimum_label -x max=maximum_label -l label -c
    clearance -f forced_privilege -x ripsoSC=RipsoSendClass -x
    ripsoSPAF=RipsoSendPAF -x ripsoRPAF=RipsoReturnPAF -x DOI=doi_value
    |
    -t hosttype=cipso -x min=minimum_label -x max=maximum_label -c
    clearance -f forced_privilege -x DOI=doi_value |
    -t hosttype=tsix -x min=minimum_label -x max=maximum_label -a
    allowed_privilege ( -i none | -i ripso -x ripsoSC=RipsoSendClass -x
    ripsoSPAF=RipsoSendPAF -x ripsoRPAF=RipsoReturnPAF | -i cipso ) -x
    DOI=doi_value |
    -h
)
```

Note: If the host type is changed, all options for the new host type must be specified.

- One of the following sets of arguments must be specified for subcommand `delete`:

```
-n template name |
-h
```

- The following argument may be specified for subcommand `list`:

```
-h
```

EXAMPLES

EXAMPLE 1 Adding a new entry to the network template database

The admin role connects to port 898 (which happens to be the default) of the aviary server on the nis:/birds/aves.Sun.COM domain, and creates the `tsol` entry in the `tnrhtp` database. The new template is assigned a host type of `unlabeled`, minimum label of `confidential`, maximum label of `top secret`, label of `secret`, clearance of `top secret able baker`, forced privilege of `all`, IP label type of `cipso`, and domain of interpretation of `1`. The administrator is prompted for the admin password.

EXAMPLE 1 Adding a new entry to the network template database (Continued)

```
$ /usr/sadm/bin/smnettmpl add -D nis:/birds/aves.Sun.COM -H aviary:898 -- \
-n tsol -t hosttype=unlabeled -x min=confidential -x max="top secret" \
-l secret -c "ts a b" -f all -i cipso -x DOI=1
```

EXIT STATUS

The following exit values are returned:

- 0 Successful completion.
- 1 Invalid command syntax. A usage message displays.
- 2 An error occurred while executing the command. An error message displays.

FILES

The following files are used by the smnettmpl command:

/etc/security/tsol/tnrhttp Trusted network remote-host templates. See tnrhttp(4).

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmgapp

Trusted Solaris 8
4/01 Reference
Manual

smc(1M), tnrhttp(4)

attributes(5)

smnetwork(1M)

NAME	smnetwork – Manage entries in the networks database								
SYNOPSIS	<code>/usr/sadm/bin/smnetwork subcommand [auth_args] - - [subcommand_args]</code>								
DESCRIPTION	<p>The smnetwork command adds, modifies, deletes, and lists entries in the networks and netmasks databases.</p> <p>smnetwork subcommands are:</p> <table> <tr> <td>add</td><td>Adds a new entry to the networks database. To add an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.</td></tr> <tr> <td>delete</td><td>Deletes an entry from the networks database. To delete an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.</td></tr> <tr> <td>list</td><td>Lists all entries in the networks database. To list an entry, the administrator must have the <code>solaris.network.host.read</code> and <code>solaris.network.security.read</code> authorizations.</td></tr> <tr> <td>modify</td><td>Modifies an entry in the networks database. To modify an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.</td></tr> </table>	add	Adds a new entry to the networks database. To add an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.	delete	Deletes an entry from the networks database. To delete an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.	list	Lists all entries in the networks database. To list an entry, the administrator must have the <code>solaris.network.host.read</code> and <code>solaris.network.security.read</code> authorizations.	modify	Modifies an entry in the networks database. To modify an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.
add	Adds a new entry to the networks database. To add an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.								
delete	Deletes an entry from the networks database. To delete an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.								
list	Lists all entries in the networks database. To list an entry, the administrator must have the <code>solaris.network.host.read</code> and <code>solaris.network.security.read</code> authorizations.								
modify	Modifies an entry in the networks database. To modify an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.								
OPTIONS	<p>The smnetwork authentication arguments, <i>auth_args</i>, are derived from the smc(1M) arg set and are the same regardless of which subcommand you use. The smnetwork command requires the SMC to be initialized for the command to succeed (see smc(1M)). After rebooting the SMC server, the first smc connection may time out, so you may need to retry the command.</p> <p>The subcommand-specific options, <i>subcommand_args</i>, must be preceded by the - - option.</p>								
<i>auth_args</i>	<p>The valid <i>auth_args</i> are -D, -H, -l, -p, -r, and -u; they are all optional. If no <i>auth_args</i> are specified, certain defaults will be assumed and the user may be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either -D or - --domain.</p> <p>-D - --domain <i>domain</i></p> <p>Specifies the default domain that you want to manage. The syntax of <i>domain=type:/host_name/domain_name</i>, where <i>type</i> is <code>nis</code>, <code>nisplus</code>, <code>dns</code>, <code>ldap</code>, or <code>file</code>; <i>host_name</i> is the name of the machine that serves the domain; and <i>domain_name</i> is the name of the domain you want to manage. (Note: Do not use <code>nis+</code> for <code>nisplus</code>.)</p> <p>If you do not specify this option, the SMC assumes the <code>file</code> default domain on whatever server you choose to manage, meaning that changes are local to the</p>								

server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.

-H | **-** *-hostname host_name:port*

Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the SMC connects to the local host on port 898.

-l | **-** *-rolepassword role_password*

Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-p | **-** *-password password*

Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-r | **-** *-rolename role_name*

Specifies a role name for authentication. If you do not specify this option, no role is assumed.

-u | **-** *-username user_name*

Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.

- **-**

This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the **- -** option.

subcommand_args

Note: Descriptions and other arg options that contain white spaces must be enclosed in double quotes.

-a *aliases*

Specifies the aliases by which the network may be identified.

-d *description*

Specifies the description of the network.

-i *ipaddress*

Specifies the IP address of the network.

-m *netmask*

Specifies the network mask of the network.

-n *network_name*

Specifies the name of the network.

■ One of the following sets of arguments must be specified for subcommand **add**:

-n network_name -i ipaddress [-m netmask] [-a aliases] [-d description] | -h

■ One of the following sets of arguments must be specified for subcommand **modify**:

-n network_name -i ipaddress [-m netmask] [-a aliases] [-d description] |

smnetwork(1M)

	<ul style="list-style-type: none">-h■ One of the following sets of arguments must be specified for subcommand delete:<ul style="list-style-type: none">-n <i>network_name</i> -h■ The following argument may be specified for subcommand list:<ul style="list-style-type: none">-h						
EXAMPLES	<p>EXAMPLE 1 Adding a new entry to the networks database</p> <p>The admin role connects to port 898 (which happens to be the default) of the aviary server on the <code>nis:/birds/aves.Sun.COM</code> domain, and creates a new network, <code>aves</code>, with an IP address of <code>129.150.0.0</code>, netmask of <code>255.255.255.0</code>, and description of <code>Aviation network</code>. The administrator is prompted for the admin password.</p> <pre>\$ /usr/sadm/bin/smnetwork add -D nis:/birds/aves.Sun.COM -H aviary:898 -- \ -n aves -i 129.150 -m 255.255.255.0 -d "Aviation network"</pre> <p>EXAMPLE 2 Modifying an entry in the networks database</p> <p>The admin role modifies the network, <code>aves</code>, changing its IP address to <code>129.160.0.0</code>, and description to <code>New aviation network</code>. Since no authorization arguments were specified, the administrator connects to port 898 of the local host on the local server with the <code>file</code> domain type, which are the defaults. The administrator is prompted for the admin password.</p> <pre>\$ /usr/sadm/bin/smnetwork modify -- -n aves -i 129.160 \ -d "New aviation network"</pre> <p>EXAMPLE 3 Deleting an entry in the networks database</p> <p>The admin role connects to the <code>nis:/birds/aves.Sun.COM</code> domain and deletes the <code>aves</code> entry. Since the host and port were not specified, the local host and port 898 are used by default. The administrator is prompted for the admin password.</p> <pre>\$ /usr/sadm/bin/smnetwork delete -D nis:/birds/aves.Sun.COM -- -n aves</pre>						
EXIT STATUS	<p>The following exit values are returned:</p> <table><tr><td>0</td><td>Successful completion.</td></tr><tr><td>1</td><td>Invalid command syntax. A usage message displays.</td></tr><tr><td>2</td><td>An error occurred while executing the command. An error message displays.</td></tr></table>	0	Successful completion.	1	Invalid command syntax. A usage message displays.	2	An error occurred while executing the command. An error message displays.
0	Successful completion.						
1	Invalid command syntax. A usage message displays.						
2	An error occurred while executing the command. An error message displays.						
FILES	<p>The following files are used by the <code>smnetwork</code> command:</p>						

smnetwork(1M)

- /etc/netmasks Network mask database. See netmasks(4).
- /etc/networks Network name database. See networks(4)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmgapp

smc(1M)
netmasks(4), networks(4), attributes(5)

smprofile(1M)

NAME	smprofile – manage profiles in the prof_attr and exec_attr databases								
SYNOPSIS	<code>/usr/sadm/bin/smprofile subcommand [auth_args] - - [subcommand_args]</code>								
DESCRIPTION	The smprofile command manages one or more profiles in the prof_attr(4) or exec_attr(4) databases in the local /etc files name service or a NIS or NIS+ name service.								
<i>subcommands</i>	<p>smprofile subcommands are:</p> <table> <tr> <td>add</td><td>Adds a new profile (right) to the prof_attr(4) database. To add a profile, the administrator must have the solaris.profmgr.write authorizations.</td></tr> <tr> <td>delete</td><td>Deletes a profile from the prof_attr(4) database, deletes all associated entries from the exec_attr(4) database, and deletes the assigned profile from the user_attr(4) database. To delete a profile, the administrator must have the solaris.profmgr.execattr.write and solaris.profmgr.write authorizations.</td></tr> <tr> <td>list</td><td>Lists one or more profiles from the prof_attr(4) or exec_attr(4) databases. To list a profile, the administrator must have the solaris.profmgr.read authorization.</td></tr> <tr> <td>modify</td><td>Modifies a profile in the prof_attr(4) database. To modify a profile, the administrator must have the solaris.profmgr.write authorizations.</td></tr> </table>	add	Adds a new profile (right) to the prof_attr(4) database. To add a profile, the administrator must have the solaris.profmgr.write authorizations.	delete	Deletes a profile from the prof_attr(4) database, deletes all associated entries from the exec_attr(4) database, and deletes the assigned profile from the user_attr(4) database. To delete a profile, the administrator must have the solaris.profmgr.execattr.write and solaris.profmgr.write authorizations.	list	Lists one or more profiles from the prof_attr(4) or exec_attr(4) databases. To list a profile, the administrator must have the solaris.profmgr.read authorization.	modify	Modifies a profile in the prof_attr(4) database. To modify a profile, the administrator must have the solaris.profmgr.write authorizations.
add	Adds a new profile (right) to the prof_attr(4) database. To add a profile, the administrator must have the solaris.profmgr.write authorizations.								
delete	Deletes a profile from the prof_attr(4) database, deletes all associated entries from the exec_attr(4) database, and deletes the assigned profile from the user_attr(4) database. To delete a profile, the administrator must have the solaris.profmgr.execattr.write and solaris.profmgr.write authorizations.								
list	Lists one or more profiles from the prof_attr(4) or exec_attr(4) databases. To list a profile, the administrator must have the solaris.profmgr.read authorization.								
modify	Modifies a profile in the prof_attr(4) database. To modify a profile, the administrator must have the solaris.profmgr.write authorizations.								
OPTIONS	<p>The smprofile authentication arguments, <i>auth_args</i>, are derived from the smc(1M) arg set and are the same regardless of which subcommand you use. The smprofile command requires the SMC to be initialized for the command to succeed (see smc(1M)). After rebooting the SMC server, the first smc connection may time out, so you may need to retry the command.</p> <p>The subcommand-specific options, <i>subcommand_args</i>, must come after the <i>auth_args</i> and must be separated from them by the - - option.</p>								
<i>auth_args</i>	<p>The valid <i>auth_args</i> are -D, -H, -l, -p, -r, and -u; they are all optional. If no <i>auth_args</i> are specified, certain defaults will be assumed and the user may be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either -D or - --domain with the <i>domain</i> argument.</p> <p>-D - --domain <i>domain</i></p> <p>Specifies the default domain that you want to manage. The syntax of <i>domain</i> is <i>type:/host_name/domain_name</i>, where <i>type</i> is nis, nisplus, dns, ldap, or file; <i>host_name</i> is the name of the machine that serves the domain; and <i>domain_name</i> is the name of the domain you want to manage. (Note: Do not use nis+ for nisplus.)</p>								

If you do not specify this option, the SMC assumes the `file` default domain on whatever server you choose to manage, meaning that changes are local to the server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.

-H | **-** `-hostname host_name:port`

Specifies the `host_name` and `port` to which you want to connect. If you do not specify a `port`, the system connects to the default port, 898. If you do not specify `host_name:port`, the SMC connects to the local host on port 898. You may still have to choose a toolbox to load into the console. To override this behavior, use the `smc(1M) -B` option, or set your console preferences to load a “home toolbox” by default.

-l | **-** `-rolepassword role_password`

Specifies the password for the `role_name`. If you specify a `role_name` but do not specify a `role_password`, the system prompts you to supply a `role_password`. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-p | **-** `-password password`

Specifies the password for the `user_name`. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-r | **-** `-rolename role_name`

Specifies a role name for authentication. If you do not specify this option, no role is assumed.

-u | **-** `-username user_name`

Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.

- **-**

This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the **- -** option.

subcommand_args

Note: Descriptions and other arg options that contain white spaces must be enclosed in double quotes.

■ For subcommand `add`:

-a *addauth1* **-a** *addauth2* . . . (Optional) Specifies the authorization name(s) to add to the new profile.

-d *description* Specifies the description of the new profile.

-h (Optional) Displays the command’s usage statement.

smprofile(1M)

<code>-m <i>html_help</i></code>	Specifies the HTML help file name for the new profile. The help file name must be put in the <code>/usr/lib/help/profiles/locale/C</code> directory.
<code>-n <i>name</i></code>	Specifies the name of the new profile.
<code>-p <i>addprof1</i> -p <i>addprof2</i> ...</code>	(Optional) Specifies the supplementary profile name(s) to add to the new profile.
■ For subcommand <code>delete</code> :	
<code>-h</code>	(Optional) Displays the command's usage statement.
<code>-n <i>name</i></code>	Specifies the name of the profile you want to delete.
■ For subcommand <code>list</code> :	
<code>-h</code>	(Optional) Displays the command's usage statement.
<code>-l</code>	(Optional) Displays the detailed output for each profile in a block of <i>key:value</i> pairs, followed by a blank line that delimits each profile block. Each <i>key:value</i> pair is displayed on a separate line. All the attributes associated with a profile from the <code>prof_attr</code> and <code>exec_attr</code> databases are displayed. If you do not specify this option, only the specified profile name(s) and associated profile description(s) are displayed.
<code>-n <i>name1</i> -n <i>name2</i> ...</code>	(Optional) Specifies the profile(s) that you want to display. If you do not specify a profile name, all profiles are displayed.
<code>-p <i>policy</i></code>	(Optional) Specifies the policy (<code>tsol</code> or <code>suser</code>). If this option is not specified, the default is <code>suser</code> . This option takes effect only when the <code>-l</code> option is specified. When <code>tsol</code> policy is specified, <code>tsol</code> entries in the <code>exec_attr</code> database are listed. When <code>suser</code> policy is specified, <code>suser</code> entries in the <code>exec_attr</code> database are listed.
■ For subcommand <code>modify</code> :	
<code>-a <i>addauth1</i> -a <i>addauth2</i> ...</code>	(Optional) Specifies the authorization name(s) to add to the profile.
<code>-d <i>description</i></code>	(Optional) Specifies the new description of the profile.
<code>-h</code>	(Optional) Displays the command's usage statement.
<code>-m <i>html_help</i></code>	(Optional) Specifies the new HTML help file name of the profile. If you change this name, you must

accordingly rename the help file name entered in the /usr/lib/help/profiles/locale/C directory.

- n *name* Specifies the name of the profile you want to modify.
- p *addprof1* -p *addprof2* ... (Optional) Specifies the supplementary profile name(s) to add to the profile.
- q *delprof1* -q *delprof2* ... (Optional) Specifies the supplementary profile name(s) to delete from the profile.
- r *delauth1* -r *delauth2* ... (Optional) Specifies the authorization name(s) to delete from the profile.

EXAMPLES

EXAMPLE 1 Adding a new profile

The admin role connects to port 898 (which happens to be the default) of the aviary server on the nisplus:/eagle/accipitridae.sun.com domain, and adds the User Manager profile. The new profile description is Manage users and groups, and the authorizations assigned are solaris.admin.usermgr.write and solaris.admin.usermgr.read. The supplementary profile assigned is Operator. The help file name is RtUserMgmt.html. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smpprofile add -D nisplus:/eagle/accipitridae.sun.com \
-H aviary:898 -- -n "User Manager" -d "Manage users and groups" \
-a solaris.admin.usermgr.write -a solaris.admin.usermgr.read \
-p Operator -m RtUserMgmt.html
```

EXAMPLE 2 Deleting a profile

The admin role deletes the All Actions profile from the local file system. Since no authorization arguments were specified, the administrator connects to port 898 of the local host on the local server with the file domain type, which are the defaults. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smpprofile delete -- -n "All Actions"
```

EXAMPLE 3 Listing all profiles

The admin role connects to the nisplus:/eagle/accipitridae.sun.com domain and lists all profiles and their associated profile descriptions on the local file system. Since the host and port were not specified, the local host and port 898 are used by default. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smpprofile list \
-D nisplus:/eagle/accipitridae.sun.com --
```

smprofile(1M)

EXAMPLE 4 Listing the boot profile

The admin role lists the boot profile on the local file system. Since no authorization arguments were specified, the administrator connects to port 898 of the local host on the local server with the file domain type, which are the defaults. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smpfile list -- -n boot -l -p tsol
```

EXAMPLE 5 Modifying a profile

The admin role connects to the aviary server on the nisplus:/eagle/accipitridae.sun.com domain and modifies the User Manager profile. The new profile description is Manage world, the new authorization assignment is solaris.admin.usermgr.* authorizations, and the new supplementary profile assignment is All. (The -a option argument must be enclosed in double quotes when the wildcard character (*) is used.). Since the port was not specified, port 898 is used by default. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smpfile modify \
-D nisplus:/eagle/accipitridae.sun.com -H aviary -- \
-n "User Manager" -d "Manage world" -a "solaris.admin.usermgr.*" \
-p All
```

ENVIRONMENT VARIABLES

See environ(5) for a description of the JAVA_HOME environment variable, which affects the execution of the smprofile command. If this environment variable is not specified, the /usr/java location is used. See smc(1M).

EXIT STATUS

The following exit values are returned:

- 0 Successful completion.
- 1 Invalid command syntax. A usage message displays.
- 2 An error occurred while executing the command. An error message displays.

FILES

The following files are used by the smprofile command:

/etc/security/exec_attr	Execution profiles database. See exec_attr(4).
/etc/security/prof_attr	Profile description database. See prof_attr(4).
/etc/user_attr	Extended user attribute database. See user_attr(4).

ATTRIBUTES

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

smprofile(1M)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmga

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**

To add, modify, or delete a profile, the administrator must have the `solaris.profmgr.write` authorization. To delete a profile, the administrator must also have the `solaris.profmgr.execattr.write` authorization. To list a profile, the administrator must have the `solaris.profmgr.read` authorization.

**Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual**

`smc(1M)`, `smrole(1M)`, `exec_attr(4)`, `prof_attr(4)`, `user_attr(4)`
`attributes(5)`, `environ(5)`

smrole(1M)

NAME	smrole – manage roles and users in role accounts								
SYNOPSIS	<code>/usr/sadm/bin/smrole subcommand [auth_args] - - [subcommand_args]</code>								
DESCRIPTION	<p>The <code>smrole</code> command manages roles and adds or deletes users in role accounts. To set audit classes, the administrator must have the <code>solaris.admin.usermgr.audit</code> authorization.</p> <p><i>subcommands</i></p> <p><code>smrole subcommands</code> are:</p> <table> <tr> <td><code>add</code></td><td>Adds a new role entry. To add an entry, the administrator must have the <code>solaris.role.write</code> authorization.</td></tr> <tr> <td><code>delete</code></td><td>Deletes one or more roles. To delete an entry, the administrator must have the <code>solaris.role.write</code> authorization.</td></tr> <tr> <td><code>list</code></td><td>Lists one or more roles. If you do not specify a role name, all roles are listed. To list an entry, the administrator must have the <code>solaris.admin.usermgr.read</code> authorization.</td></tr> <tr> <td><code>modify</code></td><td>Adds or deletes users from a role account. To modify an entry, the administrator must have the <code>solaris.role.write</code> authorization.</td></tr> </table>	<code>add</code>	Adds a new role entry. To add an entry, the administrator must have the <code>solaris.role.write</code> authorization.	<code>delete</code>	Deletes one or more roles. To delete an entry, the administrator must have the <code>solaris.role.write</code> authorization.	<code>list</code>	Lists one or more roles. If you do not specify a role name, all roles are listed. To list an entry, the administrator must have the <code>solaris.admin.usermgr.read</code> authorization.	<code>modify</code>	Adds or deletes users from a role account. To modify an entry, the administrator must have the <code>solaris.role.write</code> authorization.
<code>add</code>	Adds a new role entry. To add an entry, the administrator must have the <code>solaris.role.write</code> authorization.								
<code>delete</code>	Deletes one or more roles. To delete an entry, the administrator must have the <code>solaris.role.write</code> authorization.								
<code>list</code>	Lists one or more roles. If you do not specify a role name, all roles are listed. To list an entry, the administrator must have the <code>solaris.admin.usermgr.read</code> authorization.								
<code>modify</code>	Adds or deletes users from a role account. To modify an entry, the administrator must have the <code>solaris.role.write</code> authorization.								
OPTIONS	<p>The <code>smrole</code> authentication arguments, <i>auth_args</i>, are derived from the <code>smc(1M)</code> arg set and are the same regardless of which subcommand you use. The <code>smrole</code> command requires the SMC to be initialized for the command to succeed (see <code>smc(1M)</code>). After rebooting the SMC server, the first <code>smc</code> connection may time out, so you may need to retry the command.</p> <p>The subcommand-specific options, <i>subcommand_args</i>, must come after the <i>auth_args</i> and must be separated from them by the <code>- -</code> option.</p>								
<i>auth_args</i>	<p>The valid <i>auth_args</i> are <code>-D</code>, <code>-H</code>, <code>-l</code>, <code>-p</code>, <code>-r</code>, and <code>-u</code>; they are all optional. If no <i>auth_args</i> are specified, certain defaults will be assumed and the user may be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either <code>-D</code> or <code>- -domain</code> with the <i>domain</i> argument.</p> <p><code>-D - -domain domain</code></p> <p>Specifies the default domain that you want to manage. The syntax of <i>domain</i> is <i>type</i>:<i>/host_name/domain_name</i>, where <i>type</i> is <code>nis</code>, <code>nisplus</code>, <code>dns</code>, <code>ldap</code>, or <code>file</code>; <i>host_name</i> is the name of the machine that serves the domain; and <i>domain_name</i> is the name of the domain you want to manage. (Note: Do not use <code>nis+</code> for <code>nisplus</code>.)</p> <p>If you do not specify this option, the SMC assumes the <code>file</code> default domain on whatever server you choose to manage, meaning that changes are local to the server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.</p>								

-H | **-** **-hostname** *host_name:port*

Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the SMC connects to the local host on port 898. You may still have to choose a toolbox to load into the console. To override this behavior, use the smc(1M) -B option, or set your console preferences to load a “home toolbox” by default.

-l | **-** **-rolepassword** *role_password*

Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-p | **-** **-password** *password*

Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-r | **-** **-rolename** *role_name*

Specifies a role name for authentication. If you do not specify this option, no role is assumed.

-u | **-** **-username** *user_name*

Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.

- **-**

This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the **-** **-** option.

subcommand_args

Note: Descriptions and other arg options that contain white spaces must be enclosed in double quotes.

■ For subcommand **add**:

-a *adduser1 -a adduser2 . . .* (Optional) Specifies the user name(s) to add to the new role.

-c *comment* (Optional) Includes a short description of the role. Consists of a string of up to 256 printable characters, excluding the colon (:).

-d *dir* (Optional) Specifies the home directory of the new role, limited to 1024 characters.

-F *full_name* (Optional) Specifies the full, descriptive name of the role. The *full_name* must be unique within a domain, and can contain alphanumeric characters and spaces. If you use spaces, you must enclose the *full_name* in double quotes.

smrole(1M)

<code>-G group1 -G group2 . . .</code>	(Optional) Specifies the new role's supplementary group membership in the system group database with the character string names of one or more existing groups. <i>Note:</i> You cannot assign a primary group to a role. A role's primary group is always <code>sysadmin</code> (group 14).
<code>-h</code>	(Optional) Displays the command's usage statement.
<code>-n rolename</code>	Specifies the name of the role you want to create.
<code>-p addprof1 -p addprof2 . . .</code>	(Optional) Specifies the profile(s) to add to the role. To assign a profile to a role, the administrator must have the <code>solaris.profmgr.assign</code> or <code>solaris.profmgr.delegate</code> authorization.
<code>-P password</code>	(Optional) Specifies the role's password. The <i>password</i> can contain up to eight characters. If you do not specify a password, the system prompts you for one. <i>Note:</i> When you specify a password using the <code>-P</code> option, you type the password in plain text. Specifying a password using this method introduces a security gap while the command is running. However, if you do not specify a password (and the system prompts you for one), the echo is turned off when you type in the password. To set the password, the administrator must have the <code>solaris.admin.usermgr.pswd</code> authorization.
<code>-s shell</code>	(Optional) Specifies the full pathname of the program used as the role's shell on login. Valid entries are <code>/bin/pfcsh</code> (C shell), <code>/bin/pfksh</code> (Korn shell), and <code>/bin/pfsh</code> (Bourne shell), the default.
<code>-u uid</code>	(Optional) Specifies the ID of the role you want to add. If you do not specify this option, the system assigns the next available unique ID greater than 100.
<code>-x autohome=Y N</code>	(Optional) Sets the role's home directory. The home directory path in the password entry is set to <code>/home/login name</code> .
<code>-x clear=clearanceval</code>	(Optional) Specifies the role's clearance. <i>clearanceval</i> can be a string value or a hex value. If this option is not specified, the default is <code>admin_high</code> . To set the

clearance, the administrator must have the `solaris.admin.usermgr.labels` authorization.

- x `label=labelval` (Optional) Specifies the role's minimum label. *labelval* can be a string label or a hex label. If this option is not specified, the default is `admin_low`. To set the minimum label, the administrator must have the `solaris.admin.usermgr.labels` authorization.
- x `labelview=HIDE|SHOW` (Optional) Specifies the second part of the `labelview` key value pair. If `SHOW` is specified, `labelview=*showsl` will be recorded. If `HIDE` is specified, `labelview=*hidesl` will be recorded. * can be "internal," "external," or "". If this option is not specified, the default is `SHOW`.
- x `perm=home_perm` (Optional) Sets the permissions on the role's home directory. `perm` is interpreted as an octal number. If this option is not specified, the default is `0775`.
- x `pwupdate=AUTO|MANUAL` (Optional) Specifies how the password is changed by the user. If the `AUTO` option is specified, `passwd=automatic` will be recorded in `user_attr`. If `MANUAL` is specified, `passwd>manual` will be recorded in `user_attr`. If this option is not specified, the default is `MANUAL`.
- x `serv=homedir_server` (Optional) If `-D` is `nis`, `nisplus`, or `ldap`, use this option to specify the name of the server where the user's home directory resides. Users created in a local scope must have their home directory server created on their local machines.
- x `view=INTERNAL|EXTERNAL|DEFAULT` (Optional) Specifies the label view type for the `labelview` in `user_attr`. If `INTERNAL` is specified, `labelview=internal` will be recorded; if `EXTERNAL` is specified, `labelview=external` will be recorded; if `DEFAULT` is specified, nothing will be recorded to `user_attr`. If this option is not specified, the default is `INTERNAL`.
- For subcommand `delete`:
 - h (Optional) Displays the command's usage statement.
 - n *rolename1* -n *rolename2* ... Specifies the name of the role(s) you want to delete.
- For subcommand `list`:
 - h (Optional) Displays the command's usage statement.

smrole(1M)

-l	(Optional) Displays the output for each user in a block of <i>key:value</i> pairs (for example, <i>user name:root</i>), followed by a blank line that delimits each user block. Each <i>key:value</i> pair is displayed on a separate line. The keys are: <i>autohome</i> , <i>setup</i> , <i>comment</i> , <i>home directory</i> , <i>login shell</i> , <i>primary group</i> , <i>secondary groups</i> , <i>server</i> , <i>user ID (UID)</i> , and <i>user name</i> .
-n <i>role1</i> -n <i>role2</i> ...	(Optional) Specifies the role(s) that you want to list. If you do not specify a role name, all roles are listed.
■ For subcommand <i>modify</i> :	
-a <i>adduser1</i> -a <i>adduser2</i> ...	(Optional) Specifies the user name(s) to add to the role.
-c <i>comment</i>	(Optional) Includes a short description of the role. Consists of a string of up to 256 printable characters, excluding the colon (:).
-d <i>dir</i>	(Optional) Specifies the home directory of the new role, limited to 1024 characters.
-F <i>full_name</i>	(Optional) Specifies the full, descriptive name of the role. The <i>full_name</i> must be unique within a domain, and can contain alphanumeric characters and spaces. If you use spaces, you must enclose the <i>full_name</i> in double quotes.
-G <i>group1</i> -G <i>group2</i> ...	(Optional) Specifies the new role's secondary group membership in the system group database with the character string names of one or more existing groups. <i>Note:</i> You cannot assign a primary group to a role. A role's primary group is always <i>sysadmin</i> (group 14).
-h	(Optional) Displays the command's usage statement.
-n <i>rolename</i>	Specifies the name of the role you want to modify.
-N <i>new_rolename</i>	(Optional) Specifies the new name of the role.
-p <i>addprof1</i> -p <i>addprof2</i> ...	(Optional) Specifies the profile(s) to add to the role. To assign a profile to a role, the administrator must have the <i>solaris.profmgr.assign</i> or <i>solaris.profmgr.delegate</i> authorization.
-P <i>password</i>	(Optional) Specifies the role's password. The <i>password</i> can contain up to eight characters. <i>Note:</i> When you specify a password, you type the password in plain text. Specifying a password using this method introduces a security gap while

	the command is running. To set the password, the administrator must have the <code>solaris.admin.usermgr.pswd</code> authorization.
<code>-q delprof1 -q delprof2 . . .</code>	(Optional) Specifies the profile(s) to delete from the role.
<code>-r deluser1 -r deluser2 . . .</code>	(Optional) Specifies the user name(s) to delete from the role.
<code>-s shell</code>	(Optional) Specifies the full pathname of the program used as the role's shell on login. Valid entries are <code>/bin/pfcsh</code> (C shell), <code>/bin/pfksh</code> (Korn shell), and <code>/bin/pfsh</code> (Bourne shell), the default.
<code>-x autohome=Y N</code>	(Optional) Sets the role's home directory. The home directory path in the password entry is set to <code>/home/login_name</code> .
<code>-x clear=clearanceval</code>	(Optional) Specifies the role's clearance. <i>clearanceval</i> can be a string value or a hex value. If this option is not specified, the default is <code>admin_high</code> . To set the clearance, the administrator must have the <code>solaris.admin.usermgr.labels</code> authorization.
<code>-x label=labelval</code>	(Optional) Specifies the role's minimum label. <i>labelval</i> can be a string label or a hex label. If this option is not specified, the default is <code>admin_low</code> . To set the minimum label, the administrator must have the <code>solaris.admin.usermgr.labels</code> authorization.
<code>-x labelview=HIDE SHOW</code>	(Optional) Specifies the second part of the <code>labelview</code> key value pair. If <code>SHOW</code> is specified, <code>labelview=*showsl</code> will be recorded. If <code>HIDE</code> is specified, <code>labelview=*hidesl</code> will be recorded. * can be "internal," "external," or "". If this option is not specified, the default is <code>SHOW</code> .
<code>-x perm=home_perm</code>	(Optional) Sets the permissions on the role's home directory. <code>perm</code> is interpreted as an octal number. If this option is not specified, the default is <code>0775</code> .
<code>-x pwupdate=AUTO MANUAL</code>	(Optional) Specifies how the password is changed by the user. If the <code>AUTO</code> option is specified, <code>passwd=automatic</code> will be recorded in <code>user_attr</code> . If <code>MANUAL</code> is specified, <code>passwd>manual</code> will be recorded in <code>user_attr</code> . If this option is not specified, the default is <code>MANUAL</code> .

```
-x view=INTERNAL|EXTERNAL|DEFAULT
```

(Optional) Specifies the label view type for the `labelview` in `user_attr`. If `INTERNAL` is specified, `labelview=internal` will be recorded; if `EXTERNAL` is specified, `labelview=external` will be recorded; if `DEFAULT` is specified, nothing will be recorded to `user_attr`. If this option is not specified, the default is `INTERNAL`.

EXAMPLES

EXAMPLE 1 Adding a role account

The admin role connects to port 898 (which happens to be the default) of the aviary server on the `nisplus:/eagle/accipitridae.sun.com` domain, and adds the singer role account with a full name of Kaori Mochida and a password of `$jPoP213` on the local file system, and assigns users `kmochida` and `migarashi` to the role. This role has Name Service Security and Audit Review profiles, a clearance of Top Secret Able Baker, a minimum label of confidential, the `labelview` is set to show the label, the password update is set to provide a list from which the user must choose for the new password, and the label view is `EXTERNAL`. The system assigns the next available unique UID greater than 100. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smrole add -D nisplus:/eagle/accipitridae.sun.com \
-H aviary:898 -- -n singer -F "Kaori Mochida" -P $jPoP213
-a kmochida -a migarashi -p "Name Service Security" \
-p "Audit Review" -x clear="TS A B" -x label=confidential \
-x labelview=SHOW -x pwupdate=AUTO -x view=EXTERNAL
```

EXAMPLE 2 Deleting role accounts

The admin role deletes the singer and musician role accounts from the local file system. Since no authorization arguments were specified, the administrator connects to port 898 of the local host on the local server with the `file` domain type, which are the defaults. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smrole delete -- -n singer -n musician
```

EXAMPLE 3 Listing role accounts

The admin role connects to the `nisplus:/eagle/accipitridae.sun.com` domain and lists all role accounts on the local file system in summary form. Since the host and port were not specified, the local host and port 898 are used by default. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smrole list \
-D nisplus:/eagle/accipitridae.sun.com --
```

EXAMPLE 4 Modifying a role account

The admin role connects to port 898 of the aviary server and modifies the singer role account so the role defaults to the Korn shell, includes the `iito` account, does not

EXAMPLE 4 Modifying a role account *(Continued)*

include the migarashi account, changes the clearance to Secret Able, and the label view is changed to INTERNAL. Since the domain was not specified, the file domain type and local server are used by default. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smrole modify -H aviary:898 -- -n singer \
-s /bin/pfksh -a iito -r migarashi -x clear "S A" -x view=internal
```

ENVIRONMENT VARIABLES

See environ(5) for a description of the JAVA_HOME environment variable, which affects the execution of the smrole command. If this environment variable is not specified, the /usr/java location is used. See smc(1M).

EXIT STATUS

The following exit values are returned:

- 0 Successful completion.
- 1 Invalid command syntax. A usage message displays.
- 2 An error occurred while executing the command. An error message displays.

FILES

The following files are used by the smrole command:

/etc/aliases	Mail aliases. See aliases(4).
/etc/auto_home	Automatic mount points. See automount(1M).
/etc/group	Group file. See group(4).
/etc/passwd	Password file. See passwd(4).
/etc/security/policy.conf	Configuration file for security policy. See policy.conf(4).
/etc/shadow	Shadow password file. See shadow(4).
/etc/user_attr	Extended user attribute database. See user_attr(4).

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmga

SUMMARY OF TRUSTED SOLARIS CHANGES

To set audit classes, the administrator must have the solaris.admin.usermgr.audit authorization. To add, modify, or delete an entry, the administrator must have the solaris.role.write authorization. To list an

smrole(1M)

entry, the administrator must have the `solaris.admin.usermgr.read` authorization. To set the clearance or minimum label, the administrator must have the `solaris.admin.usermgr.labels` authorization. To set the password, the administrator must have the `solaris.admin.usermgr.pswd` authorization. To assign a profile to a user, the administrator must have the `solaris.profmgr.assign` or `solaris.profmgr.delegate` authorization.

Additional `-x` security options may be specified for the `add` and `modify` subcommands.

**Trusted Solaris 8
4/01 Reference
Manual**
**SunOS 5.8
Reference Manual**

`automount(1M)`, `smc(1M)`, `smprofile(1M)`, `policy.conf(4)`, `shadow(4)`,
`user_attr(4)`

`aliases(4)`, `group(4)`, `passwd(4)`, `attributes(5)`, `environ(5)`

NAME	smuser – manage user entries								
SYNOPSIS	/usr/sadm/bin/smuser <i>subcommand</i> [<i>auth_args</i>] - - [<i>subcommand_args</i>]								
DESCRIPTION	The smuser command manages one or more user entries in the local /etc filesystem or a NIS or NIS+ target name service. To set audit classes, the administrator must have the <code>solaris.admin.usermgr.audit</code> authorization.								
<i>subcommands</i>	<p>smuser <i>subcommands</i> are:</p> <table> <tr> <td>add</td><td>Adds a new user entry to the appropriate files. You can use a template and input file instead of supplying the additional command line options. If you use a template and command line options, the command line options take precedence and override any conflicting template values. To add an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.</td></tr> <tr> <td>delete</td><td>Deletes one or more user entries from the appropriate files. <i>Note:</i> You cannot delete the system accounts with IDs less than 100, or 60001, 60002, or 65534. To delete an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.</td></tr> <tr> <td>list</td><td>Lists one more user entries from the appropriate files. To list entries, the administrator must have the <code>solaris.admin.usermgr.read</code> authorization.</td></tr> <tr> <td>modify</td><td>Modifies a user entry in the appropriate files. To modify an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.</td></tr> </table>	add	Adds a new user entry to the appropriate files. You can use a template and input file instead of supplying the additional command line options. If you use a template and command line options, the command line options take precedence and override any conflicting template values. To add an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.	delete	Deletes one or more user entries from the appropriate files. <i>Note:</i> You cannot delete the system accounts with IDs less than 100, or 60001, 60002, or 65534. To delete an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.	list	Lists one more user entries from the appropriate files. To list entries, the administrator must have the <code>solaris.admin.usermgr.read</code> authorization.	modify	Modifies a user entry in the appropriate files. To modify an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.
add	Adds a new user entry to the appropriate files. You can use a template and input file instead of supplying the additional command line options. If you use a template and command line options, the command line options take precedence and override any conflicting template values. To add an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.								
delete	Deletes one or more user entries from the appropriate files. <i>Note:</i> You cannot delete the system accounts with IDs less than 100, or 60001, 60002, or 65534. To delete an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.								
list	Lists one more user entries from the appropriate files. To list entries, the administrator must have the <code>solaris.admin.usermgr.read</code> authorization.								
modify	Modifies a user entry in the appropriate files. To modify an entry, the administrator must have the <code>solaris.admin.usermgr.write</code> authorization.								
OPTIONS	<p>The smuser authentication arguments, <i>auth_args</i>, are derived from the smc(1M) arg set and are the same regardless of which subcommand you use. The smuser command requires the SMC to be initialized for the command to succeed (see smc(1M)). After rebooting the SMC server, the first smc connection may time out, so you may need to retry the command.</p> <p>The subcommand-specific options, <i>subcommand_args</i>, must come after the <i>auth_args</i> and must be separated from them by the - - option.</p>								
<i>auth_args</i>	<p>The valid <i>auth_args</i> are -D, -H, -l, -p, -r, and -u; they are all optional. If no <i>auth_args</i> are specified, certain defaults will be assumed and the user may be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either -D or - --domain with the <i>domain</i> argument.</p> <p>-D - --domain <i>domain</i></p> <p>Specifies the default domain that you want to manage. The syntax of <i>domain</i> is <i>type</i>:/<i>host_name</i>/<i>domain_name</i>, where <i>type</i> is nis, nisplus, dns, ldap, or file; <i>host_name</i> is the name of the machine that serves the domain; and <i>domain_name</i> is the name of the domain you want to manage. (<i>Note:</i> Do not use nis+ for nisplus.)</p>								

smuser(1M)

If you do not specify this option, the SMC assumes the `file` default domain on whatever server you choose to manage, meaning that changes are local to the server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.

`-H | - --hostname host_name:port`

Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the SMC connects to the local host on port 898. You may still have to choose a toolbox to load into the console. To override this behavior, use the `smc(1M) -B` option, or set your console preferences to load a “home toolbox” by default.

`-l | - --rolepassword role_password`

Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

`-p | - --password password`

Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

`-r | - --rolename role_name`

Specifies a role name for authentication. If you do not specify this option, no role is assumed.

`-u | - --username user_name`

Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.

`- -`

This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the `- -` option.

subcommand_args

Note: Descriptions and other arg options that contain white spaces must be enclosed in double quotes.

■ For subcommand `add`:

`-c comment`

(Optional) Includes a short description of the login, which is typically the user’s name. Consists of a string of up to 256 printable characters, excluding the colon (:).

`-d dir`

(Optional) Specifies the home directory of the new user, limited to 1024 characters.

`-e ddmmyyyy`

(Optional) Specifies the expiration date for a login. After this date, no user can access this login. This

	option is useful for creating temporary logins. Specify a null value (" ") to indicate that the login is always valid.
<code>-f <i>inactive</i></code>	(Optional) Specifies the maximum number of days allowed between uses of a login ID before that ID is declared invalid. Normal values are positive integers. Enter zero to indicate that the login account is always active.
<code>-F <i>full_name</i></code>	(Optional) Specifies the full, descriptive name of the user. The <i>full_name</i> must be unique within a domain and can contain alphanumeric characters and spaces. If you use spaces, you must enclose the <i>full_name</i> in double quotes.
<code>-g <i>group</i></code>	(Optional) Specifies the new user's primary group membership in the system group database with an existing group's integer ID.
<code>-G <i>group1</i> -G <i>group2</i> . . .</code>	(Optional) Specifies the new user's supplementary group membership in the system group database with the character string names of one or more existing groups. Duplicates of groups specified with the <code>-g</code> and <code>-G</code> options are ignored.
<code>-h</code>	(Optional) Displays the command's usage statement.
<code>-n <i>login</i></code>	Specifies the new user's login name. The login name must be unique within a domain, contain 2–32 alphanumeric characters, begin with a letter, and contain at least one lowercase letter.
<code>-P <i>password</i></code>	(Optional) Specifies up to an eight-character password assigned to the user account. Note: When you specify a password, you type the password in plain text. Specifying a password using this method introduces a security gap while the command is running. If this option is not specified, the user is prompted to input a password upon first login to the new account. This only works on the local file system. To set the password, the administrator must have the <code>solaris.admin.usermgr.pswd</code> authorization.
<code>-s <i>shell</i></code>	(Optional) Specifies the full pathname (limited to 1024 characters) of the program used as the user's shell on login. Valid entries are a user-defined shell, <code>/bin/csh</code> (C shell), <code>bin/ksh</code> (Korn shell), and the default, <code>/bin/sh</code> (Bourne shell).

smuser(1M)

<code>-t <i>template</i></code>	(Optional) Specifies a template, created using the User Manager tool, that contains a set of pre-defined user attributes. You may have entered a name service server in the template. However, when a user is actually added with this template, if a name service is unavailable, the user's local server will be used for both the Home Directory Server and Mail Server.
<code>-u <i>uid</i></code>	(Optional) Specifies the user ID of the user you want to add. If you do not specify this option, the system assigns the next available unique user ID greater than 100.
<code>-x autohome=Y N</code>	(Optional) Sets the home directory to automount if set to Y. The user's home directory path in the password entry is set to <code>/home/login name</code> .
<code>-x clear=<i>clearanceval</i></code>	(Optional) Specifies the user's clearance. <i>clearanceval</i> can be a string value or a hex value. If this option is not specified, the default is the user's system default clearance. To set the clearance, the administrator must have the <code>solaris.admin.usermgr.labels</code> authorization.
<code>-x idlecmd=LOGOUT LOCK</code>	(Optional) Specifies the command to execute if the system has been idled. If LOGOUT is specified, <code>idlecmd=logout</code> will be recorded in <code>user_attr</code> . If LOCK is specified, <code>idlecmd=lock</code> will be recorded in <code>user_attr</code> . If this option is not specified, the default is the IDLECMD in the <code>/etc/security/policy.conf</code> file.
<code>-x idletime=<i>minutes</i></code>	(Optional) Specifies the number of minutes before the specified idle command gets executed. Any integer value between 1 and 120 is valid. This value is recorded into <code>user_attr</code> as <code>idletime=val</code> . If this option is not specified, the default is the IDLETIME in the <code>/etc/security/policy.conf</code> file.
<code>-x label=<i>labelval</i></code>	(Optional) Specifies the user's minimum label. <i>labelval</i> can be a string label or a hex label. If this option is not specified, the default is the user's system default minimum label. To set the minimum label, the administrator must have the <code>solaris.admin.usermgr.labels</code> authorization.
<code>-x labelview=HIDE SHOW</code>	(Optional) Specifies the second part of the <code>labelview</code> key value pair. If SHOW is specified, <code>labelview=*showsl</code> will be recorded. If HIDE is specified, <code>labelview=*hidesl</code> will be recorded. * can

	be "internal," "external," or "". If this option is not specified, the default is the LABELVIEW in the /etc/security/policy.conf file.
-x lock=Y N	(Optional) Specifies if an account is locked after a specified number of failed logins. This value is recorded in user_attr as lock_after_retries. If this option is not specified, the default is the LOCK_AFTER_RETRIES in the /etc/security/policy.conf file.
-x mail=mail_server	(Optional) Specifies the host name of the user's mail server, and creates a mail file on the server. Users created in a local scope must have a mail server created on their local machines.
-x perm=home_perm	(Optional) Sets the permissions on the user's home directory. perm is interpreted as an octal number. If this option is not specified, the default is 0775.
-x pwmax=days	(Optional) Specifies the maximum number of days that the user's password is valid.
-x pwmin=days	(Optional) Specifies the minimum number of days between user password changes.
-x pwupdate=AUTO MANUAL	(Optional) Specifies how the password is changed by the user. If the AUTO option is specified, passwd=automatic will be recorded in user_attr. If MANUAL is specified, passwd=manual will be recorded in user_attr. If this option is not specified, the default is the PASSWORD in the /etc/security/policy.conf file.
-x pwwarn=days	(Optional) Specifies the number of days relative to pwmax that the user is warned about password expiration prior to the password expiring.
-x serv=homedir_server	(Optional) Specifies the name of the server where the user's home directory resides. Users created in a local scope must have their home directory server created on their local machines.
-x view=INTERNAL EXTERNAL DEFAULT	(Optional) Specifies the label view type for the labelview in user_attr. If INTERNAL is specified, labelview=internal will be recorded; if EXTERNAL is specified, labelview=external will be recorded; if DEFAULT is specified, nothing will be recorded to user_attr. If this option is not specified, nothing will get recorded to user_attr by default.
■ For subcommand delete:	
-h	(Optional) Displays the command's usage statement.

smuser(1M)

- n *login1* Specifies the login name of the user you want to delete.
- n *login2* . . . (Optional) Specifies the additional login name(s) of the user(s) you want to delete.
- For subcommand *list*:
 - h (Optional) Displays the command's usage statement.
 - l Displays the output for each user in a block of *key:value* pairs (for example, *user name:root*) followed by a blank line to delimit each user block. Each *key:value* pair is displayed on a separate line. The keys are: *autohome setup*, *comment*, *days to warn*, *full name*, *home directory*, *home directory permissions*, *login shell*, *mail server*, *max days change*, *max days inactive*, *min days change*, *password expires*, *password type*, *primary group*, *rights*, *roles*, *secondary groups*, *server*, *user ID (UID)*, and *user name*.
 - n *login1* Specifies the login name of the user you want to list.
 - n *login2* . . . (Optional) Specifies the additional login name(s) of the user(s) you want to list.
- For subcommand *modify*:
 - a *addrole1* -a *addrole2* . . . (Optional) Specifies the role(s) to add to the user account. To assign a role to a user, the administrator must have the *solaris.role.assign* or *solaris.role.delegate* authorization.
 - c *comment* (Optional) Describes the changes you made to the user account. Consists of a string of up to 256 printable characters, excluding the colon (:).
 - d *description* (Optional) Specifies the user's home directory, limited to 1024 characters.
 - e *ddmmyyyy* (Optional) Specifies the expiration date for a login in a format appropriate to the locale. After this date, no user can access this login. This option is useful for creating temporary logins. Specify a null value (" ") to indicate that the login is always valid.
 - f *inactive* (Optional) Specifies the maximum number of days allowed between uses of a login ID before the ID is declared invalid. Normal values are positive integers. Specify zero to indicate that the login account is always active.
 - F *full_name* (Optional) Specifies the full, descriptive name of the user. The *full_name* must be unique within a domain

and can contain alphanumeric characters and spaces. If you use spaces, you must enclose the *full_name* in double quotes.

<code>-g group</code>	(Optional) Specifies the new user's primary group membership in the system group database with an existing group's integer ID.
<code>-G group1 -G group2 . . .</code>	(Optional) Specifies the new user's supplementary group membership in the system group database with the character string names of one or more existing groups. Duplicates of groups specified with the <code>-g</code> and <code>-G</code> options are ignored.
<code>-h</code>	(Optional) Displays the command's usage statement.
<code>-n name</code>	Specifies the user's current login name.
<code>-N new_name</code>	(Optional) Specifies the user's new login name. The login name must be unique within a domain, contain 2–32 alphanumeric characters, begin with a letter, and contain at least one lowercase letter.
<code>-p addprof1 -p addprof2 . . .</code>	(Optional) Specifies the profile(s) to add to the user account. To assign a profile to a user, the administrator must have the <code>solaris.profmgr.assign</code> or <code>solaris.profmgr.delegate</code> authorization.
<code>-P password</code>	(Optional) Specifies up to an eight-character password assigned to the user account. When you specify a password, you type the password in plain text. Specifying a password using this method introduces a security gap while the command is running.
<code>-q delprof1 -q delprof2 . . .</code>	(Optional) Specifies the profile(s) to delete from the user account.
<code>-r delrole1 -r delrole2 . . .</code>	(Optional) Specifies the role(s) to delete from the user account.
<code>-s shell</code>	(Optional) Specifies the full pathname (limited to 1024 characters) of the program used as the user's shell on login. Valid entries are a user-defined shell, <code>/bin/csh</code> (C shell), <code>bin/ksh</code> (Korn shell), and the default, <code>/bin/sh</code> (Bourne shell).l)
<code>-x autohome=Y N</code>	(Optional) Sets up the home directory to automount if set to <code>Y</code> . The user's home directory path in the password entry is set to <code>/home/login_name</code> .

smuser(1M)

-x clear= <i>clearanceval</i>	(Optional) Specifies the user's clearance. <i>clearanceval</i> can be a string value or a hex value. If this option is not specified, the default is the user's system default clearance. To set the clearance, the administrator must have the <code>solaris.admin.usermgr.labels</code> authorization.
-x idlecmd=LOGOUT LOCK	(Optional) Specifies the command to execute if the system has been idled. If LOGOUT is specified, <code>idlecmd=logout</code> will be recorded in <code>user_attr</code> . If LOCK is specified, <code>idlecmd=lock</code> will be recorded in <code>user_attr</code> . If this option is not specified, the default is the IDLECMD in the <code>/etc/security/policy.conf</code> file.
-x idletime= <i>minutes</i>	(Optional) Specifies the number of minutes before the specified idle command gets executed. Any integer value between 1 and 120 is valid. This value is recorded into <code>user_attr</code> as <code>idletime=val</code> . If this option is not specified, the default is the IDLETIME in the <code>/etc/security/policy.conf</code> file.
-x label= <i>labelval</i>	(Optional) Specifies the user's minimum label. <i>labelval</i> can be a string label or a hex label. If this option is not specified, the default is the user's system default minimum label. To set the minimum label, the administrator must have the <code>solaris.admin.usermgr.labels</code> authorization.
-x labelview=HIDE SHOW	(Optional) Specifies the second part of the labelview key value pair. If SHOW is specified, <code>labelview=*showsl</code> will be recorded. If HIDE is specified, <code>labelview=*hidesl</code> will be recorded. * can be "internal," "external," or "". If this option is not specified, the default is in the LABELVIEW in the <code>/etc/security/policy.conf</code> file.
-x lock=Y N	(Optional) Specifies if an account is locked after a specified number of failed logins. This value is recorded in <code>user_attr</code> as <code>lock_after_retries</code> . If this option is not specified, the default is the LOCK_AFTER_RETRIES in the <code>/etc/security/policy.conf</code> file.
-x pwmax= <i>days</i>	(Optional) Specifies the maximum number of days that the user's password is valid.

- x pwmin=*days* (Optional) Specifies the minimum number of days between password changes.
- x pwupdate=AUTO|MANUAL (Optional) Specifies how the password is changed by the user. If the AUTO option is specified, passwd=automatic will be recorded in user_attr. If MANUAL is specified, passwd>manual will be recorded in user_attr. If this option is not specified, the default is the PASSWORD in the /etc/security/policy.conf file.
- x pwwarn=*days* (Optional) Specifies the number of days relative to pwmax that the user is warned about password expiration before the password expires.
- x view=INTERNAL|EXTERNAL|DEFAULT (Optional) Specifies the label view type for the labelview in user_attr. If INTERNAL is specified, labelview=internal will be recorded; if EXTERNAL is specified, labelview=external will be recorded; if DEFAULT is specified, nothing will be recorded to user_attr. If this option is not specified, nothing will get recorded to user_attr by default.

EXAMPLES**EXAMPLE 1** Adding a new user account

The admin role connects to port 898 (which happens to be the default) of the aviary server on the nis:/birds/aves.Sun.COM domain, and adds a new user entry, kmochida. Several options are given: the password is set to \$jPoP213 (note: this is insecure), the comment is set to Kaori's account, the full name is set to Kaori Mochida, the password update is set to provide a list from which the user must choose for the new password, the minimum label is set to Confidential, the clearance is set to Top Secret Able Baker, the view is internal, the labelview is set to show the labels, the idle command will be executed in 30 minutes, the system will lock when idled too long, and the account will lock after the maximum number of failed logins is reached. The system will assign the next available user ID greater than 100 to this account. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smuser add -D nis:/birds/aves.Sun.COM \
-H aviary:898 -- -n kmochida -P $jPoP213 -c "Kaori's account" \
-F "Kaori Mochida" -x pwupdate=AUTO -x label=confidential \
-x clear="TS A B" -x view=INTERNAL -x labelview=SHOW \
-x idletime=30 -x idlcmd=LOCK -x lock=Y
```

EXAMPLE 2 Deleting a user account

The admin role deletes the kmochida user entry in the local file system. Since no authorization arguments were specified, the administrator connects to port 898 of the local host on the local server with the file domain type, which are the defaults. The administrator is prompted for the admin password.

smuser(1M)

EXAMPLE 2 Deleting a user account (Continued)

```
$ /usr/sadm/bin/smuser delete -- -n kmochida
```

EXAMPLE 3 Listing all user accounts

The admin role connects to the `nis:/birds/aves.Sun.COM` domain and lists all user accounts on the local file system. Since the host and port were not specified, the local host and port 898 are used by default. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smuser list -D nis:/birds/aves.Sun.COM --
```

EXAMPLE 4 Modifying a user account

The admin role connects to port 898 of the `aviary` server and modifies the `kmochida` user entry by changing the default shell to the Korn shell, the supplementary group to `qa_group`, the password update method to `manual`, the clearance to `Secret Able`, and having the account not lock after the maximum number of failed logins is reached. Since the domain was not specified, the file domain type and local server are used by default. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smuser modify -H aviary:898 -- -n kmochida \
-s /bin/ksh -G qa_group -x pwupdate=MANUAL -x clear="S A" \
-x lock=N
```

ENVIRONMENT VARIABLES

See `environ(5)` for a description of the `JAVA_HOME` environment variable, which affects the execution of the `smuser` command. If this environment variable is not specified, the `/usr/java` location is used. See `smc(1M)`.

EXIT STATUS

The following exit values are returned:

- | | |
|---|---|
| 0 | Successful completion. |
| 1 | Invalid command syntax. A usage message displays. |
| 2 | An error occurred while executing the command. An error message displays. |

FILES

The following files are used by the `smuser` command:

<code>/etc/aliases</code>	Mail aliases. See <code>aliases(4)</code> .
<code>/etc/auto_home</code>	Automatic mount points. See <code>automount(1M)</code> .
<code>/etc/group</code>	Group file. See <code>group(4)</code> .
<code>/etc/passwd</code>	Password file. See <code>passwd(4)</code> .

	/etc/security/policy.conf	Configuration file for security policy. See policy.conf(4).				
	/etc/shadow	Shadow password file. See shadow(4).				
	/etc/user_attr	Extended user attribute database. See user_attr(4).				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:					
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWmga</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWmga
ATTRIBUTE TYPE	ATTRIBUTE VALUE					
Availability	SUNWmga					
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>For the list subcommand, the user must have the solaris.admin.usermgr.read authorization. To set audit classes, the user must have the solaris.admin.usermgr.audit authorization. To set the clearance or minimum label, the user must have the solaris.admin.usermgr.labels authorization. To set the password, the user must have the solaris.admin.usermgr.pswd authorization. To make most other changes, the user must have the solaris.admin.usermgr.write authorization. To assign a role to a user, the user must have the solaris.role.assign or solaris.role.delegate authorization. To assign a profile to a user, the user must have the solaris.profmgr.assign or solaris.profmgr.delegate authorization.</p> <p>Additional -x security options may be specified for the add and modify subcommands.</p>					
Trusted Solaris 8 4/01 Reference Manual	automount(1M), smc(1M), smprofile(1M), smrole(1M), policy.conf(4), shadow(4), user_attr(4)					
SunOS 5.8 Reference Manual	aliases(4), group(4), passwd(4), attributes(5), environ(5)					

snoop(1M)

NAME	snoop – capture and inspect network packets
SYNOPSIS	snoop [-aqrCDNPSvV] [-t [r a d]] [-c <i>maxcount</i>] [-d <i>device</i>] [-i <i>filename</i>] [-n <i>filename</i>] [-o <i>filename</i>] [-p <i>first</i> [, <i>last</i>]] [-s <i>snaplen</i>] [-x <i>offset</i> [, <i>length</i>]] [<i>expression</i>]
DESCRIPTION	<p>snoop captures packets from the network and displays their contents. snoop uses both the network packet filter and streams buffer modules to provide efficient capture of packets from the network. Captured packets can be displayed as they are received, or saved to a file for later inspection.</p> <p>snoop can display packets in a single-line summary form or in verbose multi-line forms. In summary form, only the data pertaining to the highest level protocol is displayed. For example, an NFS packet will have only NFS information displayed. The underlying RPC, UDP, IP, and ethernet frame information is suppressed but can be displayed if either of the verbose options are chosen.</p> <p>snoop requires an interactive interface.</p> <p>TSIX token mapping requests and responses have a special Session ID value that prevents them from being received by a process that does not have the Trusted Path process attribute. The snoop command must have the Trusted Path process attribute to capture TSIX token mapping packets.</p>
OPTIONS	<ul style="list-style-type: none"> -a Listen to packets on /dev/audio (warning: can be noisy). -C List the code generated from the filter expression for either the kernel packet filter, or snoop's own filter. -D Display number of packets dropped during capture on the summary line. -N Create an IP address-to-name file from a capture file. This must be set together with the -i option that names a capture file. The address-to-name file has the same name as the capture file with .names appended. This file records the IP address to hostname mapping at the capture site and increases the portability of the capture file. Generate a .names file if the capture file is to be analyzed elsewhere. Packets are not displayed when this flag is used. -P Capture packets in non-promiscuous mode. Only broadcast, multicast, or packets addressed to the host machine will be seen. -q When capturing network packets into a file, do not display the packet count. This can improve packet capturing performance. -r Do not resolve the IP address to the symbolic name. This prevents snoop from generating network traffic while capturing and displaying packets. However, if the -n option is used, and an address is found in the mapping file, its corresponding name will be used. -S Display size of the entire ethernet frame in bytes on the summary line.

- v Verbose mode. Print packet headers in lots of detail. This display consumes many lines per packet and should be used only on selected packets.
- V Verbose summary mode. This is halfway between summary mode and verbose mode in degree of verbosity. Instead of displaying just the summary line for the highest level protocol in a packet, it displays a summary line for each protocol layer in the packet. For instance, for an NFS packet it will display a line each for the ETHER, IP, UDP, RPC and NFS layers. Verbose summary mode output may be easily piped through `grep` to extract packets of interest. For example to view only RPC summary lines:


```
example# snoop -i rpc.cap -V | grep RPC
```
- t [r | a | d] Time-stamp presentation. Time-stamps are accurate to within 4 microseconds. The default is for times to be presented in d (delta) format (the time since receiving the previous packet). Option a (absolute) gives wall-clock time. Option r (relative) gives time relative to the first packet displayed. This can be used with the -p option to display time relative to any selected packet.
- c *maxcount* Quit after capturing *maxcount* packets. Otherwise keep capturing until there is no disk left or until interrupted with Control-C.
- d *device* Receive packets from the network using the interface specified by *device*. Usually `le0` or `ie0`. The program `netstat(1M)`, when invoked with the -i flag, lists all the interfaces that a machine has. Normally, `snoop` will automatically choose the first non-loopback interface it finds.
- i *filename* Display packets previously captured in *filename*. Without this option, `snoop` reads packets from the network interface. If a *filename.names* file is present, it is automatically loaded into the `snoop` IP address-to-name mapping table (See -N flag).
- n *filename* Use *filename* as an IP address-to-name mapping table. This file must have the same format as the `/etc/hosts` file (IP address followed by the hostname).
- o *filename* Save captured packets in *filename* as they are captured. During packet capture, a count of the number of packets saved in the file is displayed. If you wish just to count packets without saving to a file, name the file `/dev/null`.
- p *first* [, *last*] Select one or more packets to be displayed from a capture file. The *first* packet in the file is packet number 1.
- s *snaplen* Truncate each packet after *snaplen* bytes. Usually the whole packet is captured. This option is useful if only certain packet header information is required. The packet truncation is done within the kernel giving better utilization of the streams packet buffer. This means less chance of dropped packets due to buffer overflow

snoop(1M)

OPERANDS

`-x offset [, length]`

expression

during periods of high traffic. It also saves disk space when capturing large traces to a capture file. To capture only IP headers (no options) use a *snaplen* of 34. For UDP use 42, and for TCP use 54. You can capture RPC headers with a *snaplen* of 80 bytes. NFS headers can be captured in 120 bytes.

Display packet data in hexadecimal and ASCII format. The *offset* and *length* values select a portion of the packet to be displayed. To display the whole packet, use an *offset* of 0. If a *length* value is not provided, the rest of the packet is displayed.

Select packets either from the network or from a capture file. Only packets for which the expression is true will be selected. If no expression is provided it is assumed to be true.

Given a filter expression, *snoop* generates code for either the kernel packet filter or for its own internal filter. If capturing packets with the network interface, code for the kernel packet filter is generated. This filter is implemented as a streams module, upstream of the buffer module. The buffer module accumulates packets until it becomes full and passes the packets on to *snoop*. The kernel packet filter is very efficient, since it rejects unwanted packets in the kernel before they reach the packet buffer or *snoop*. The kernel packet filter has some limitations in its implementation; it is possible to construct filter expressions that it cannot handle. In this event, *snoop* tries to split the filter and do as much filtering in the kernel as possible. The remaining filtering is done by the packet filter for *snoop*. The `-C` flag can be used to view generated code for either the packet filter for the kernel or the packet filter for *snoop*. If packets are read from a capture file using the `-i` option, only the packet filter for *snoop* is used.

A filter *expression* consists of a series of one or more boolean primitives that may be combined with boolean operators (AND, OR, and NOT). Normal precedence rules for boolean operators apply. Order of evaluation of these operators may be controlled with parentheses. Since parentheses and other filter expression characters are known to the shell, it is often necessary to enclose the filter expression in quotes. Refer to Example 4 for information about setting up more efficient filters.

The primitives are:

`host hostname`

True if the source or destination address is that of *hostname*. The *hostname* argument may be a literal address. The keyword *host* may be omitted if the name does not conflict with the name of another expression primitive. For example, "pinky" selects packets transmitted to or received from the host pinky,

whereas "pinky and dinky" selects packets exchanged between hosts pinky AND dinky.

The type of address used depends on the primitive which precedes the host primitive. The possible qualifiers are "inet", "inet6", "ether", or none. These three primitives are discussed below. Having none of the primitives present is equivalent to "inet host hostname or inet6 host hostname". In other words, snoop tries to filter on all IP addresses associate with hostname.

inet or inet6

A qualifier that modifies the host primitive that follows. If it is *inet*, then snoop tries to filter on all IPv4 addresses returned from a name lookup. If it is *inet6*, snoop tries to filter on all IPv6 addresses returned from a name lookup.

ipaddr or etheraddr

Literal addresses, both IP dotted and ethernet colon are recognized. For example,

- "129.144.40.13" matches all packets with that IP ;
- "2::9255:a00:20ff:fe73:6e35" matches all packets with that IPv6 address as source or destination;
- "8:0:20:f:b1:51" matches all packets with the ethernet address as source or destination.

An ethernet address beginning with a letter is interpreted as a hostname. To avoid this, prepend a zero when specifying the address. For example, if the ethernet address is "aa:0:45:23:52:44", then specify it by add a leading zero to make it "0aa:0:45:23:52:44".

from or src

A qualifier that modifies the following host, net, *ipaddr*, *etheraddr*, port or rpc primitive to match just the source address, port, or RPC reply.

to or dst

A qualifier that modifies the following host, net, *ipaddr*, *etheraddr*, port or rpc primitive to match just the destination address, port, or RPC call.

ether

A qualifier that modifies the following host primitive to resolve a name to an ethernet address. Normally, IP address matching is performed.

ethertype number

True if the ethernet type field has value *number*. Equivalent to "ether[12:2] = *number*".

snoop(1M)

ip, ip6, arp, rarp
True if the packet is of the appropriate ethertype.

broadcast
True if the packet is a broadcast packet. Equivalent to
"ether[2:4] = 0xffffffff".

multicast
True if the packet is a multicast packet. Equivalent to
"ether[0] & 1 = 1".

apple
True if the packet is an Apple Ethertalk packet. Equivalent to
"ethertype 0x809b or ethertype 0x803f".

decnet
True if the packet is a DECNET packet.

greater length
True if the packet is longer than *length*.

less length
True if the packet is shorter than *length*.

udp, tcp, icmp, icmp6, ah, esp
True if the IP or IPv6 protocol is of the appropriate type.

net net
True if either the IP source or destination address has a network number of *net*. The *from* or *to* qualifier may be used to select packets for which the network number occurs only in the source or destination address.

port port
True if either the source or destination port is *port*. The *port* may be either a port number or name from */etc/services*. The *tcp* or *udp* primitives may be used to select TCP or UDP ports only. The *from* or *to* qualifier may be used to select packets for which the *port* occurs only as the source or destination.

rpc prog
[, *vers* [, *proc*]] True if the packet is an RPC call or reply packet for the protocol identified by *prog*. The *prog* may be either the name of an RPC protocol from */etc/rpc* or a program number. The *vers* and *proc* may be used to further qualify the program *version* and *procedure* number, for example, "rpc nfs, 2, 0" selects all calls and replies for the NFS null procedure. The *to* or *from* qualifier may be used to select either call or reply packets only.

gateway host

True if the packet used *host* as a gateway, that is, the ethernet source or destination address was for *host* but not the IP address. Equivalent to "ether host *host* and not host *host*".

nofrag

True if the packet is unfragmented or is the first in a series of IP fragments. Equivalent to "ip[6:2] & 0x1fff = 0".

sectype type

True if the packet security type is *type*. The valid values for *type* are unlabeled, tsix, and tsol.

expr relop expr

True if the relation holds, where *relop* is one of >, <, >=, <=, =, !=, and *expr* is an arithmetic expression composed of numbers, packet field selectors, the *length* primitive, and arithmetic operators +, -, *, &, |, ^, and %. The arithmetic operators within *expr* are evaluated before the relational operator and normal precedence rules apply between the arithmetic operators, such as multiplication before addition. Parentheses may be used to control the order of evaluation. To use the value of a field in the packet use the following syntax:

base[*expr* [: *size*]]

where *expr* evaluates the value of an offset into the packet from a *base* offset which may be ether, ip, udp, tcp, or icmp. The *size* value specifies the size of the field. If not given, 1 is assumed. Other legal values are 2 and 4. For example,

ether[0] & 1 = 1

is equivalent to multicast.

ether[2:4] = 0xffffffff

is equivalent to broadcast.

ip[ip[0] & 0xf * 4 : 2] = 2049

is equivalent to udp[0:2] = 2049

ip[0] & 0xf > 5

selects IP packets with options.

ip[6:2] & 0x1fff = 0

eliminates IP fragments.

udp and ip[6:2]&0x1fff = 0 and udp[6:2] != 0

finds all packets with UDP checksums.

The *length* primitive may be used to obtain the length of the packet. For instance "length > 60" is equivalent to "greater 60", and "ether[length - 1]" obtains the value of the last byte in a packet.

snoop(1M)

and

Perform a logical AND operation between two boolean values. The AND operation is implied by the juxtaposition of two boolean expressions, for example "dinky pinky" is the same as "dinky AND pinky".

or or ,

Perform a logical OR operation between two boolean values. A comma may be used instead, for example, "dinky, pinky" is the same as "dinky OR pinky".

not or !

Perform a logical NOT operation on the following boolean value. This operator is evaluated before AND or OR.

slp

True if the packet is an SLP packet.

EXAMPLES

EXAMPLE 1 Sample output from the snoop command.

Capture all packets and display them as they are received:

```
example# snoop
```

Capture packets with host funky as either the source or destination and display them as they are received:

```
example# snoop funky
```

Capture packets between funky and pinky and save them to a file. Then inspect the packets using times (in seconds) relative to the first captured packet:

```
example# snoop -o cap funky pinky
example$ snoop -i cap -t r | more
```

Look at selected packets in another capture file:

```
example$ snoop -i pkts -p 99,108
 99  0.0027  boutique -> sunroof      NFS C GETATTR FH=8E6C
100  0.0046  sunroof -> boutique      NFS R GETATTR OK
101  0.0080  boutique -> sunroof      NFS C RENAME FH=8E6C MTra00192 to .nfs08
102  0.0102  marmot -> viper          NFS C LOOKUP FH=561E screen.r.13.i386
103  0.0072  viper -> marmot          NFS R LOOKUP No such file or directory
104  0.0085  bugbomb -> sunroof      RLOGIN C PORT=1023 h
105  0.0005  kandinsky -> sparky      RSTAT C Get Statistics
106  0.0004  beebledbrox -> sunroof   NFS C GETATTR FH=0307
107  0.0021  sparky -> kandinsky      RSTAT R
108  0.0073  office -> jeremiah       NFS C READ FH=2584 at 40960 for 8192
```

EXAMPLE 2 Sample TSOL packets

Now select only those TSOL packets in the above capture file:

```
example$ snoop -i pkts -p 99,108 sectype tsol
 99  0.0027  boutique -> sunroof      NFS C GETATTR FH=8E6C
```

EXAMPLE 2 Sample TSOL packets (Continued)

```

100  0.0046  sunroof -> boutique      NFS R GETATTR OK
101  0.0080  boutique -> sunroof      NFS C RENAME FH=8E6C MTra00192 to .nfs08

```

Packet 101 looks interesting. Take a look in more detail:

```

example$ snoop -i pkts -v -p 101

ETHER: ----- Ether Header -----
ETHER:
ETHER: Packet 101 arrived at 16:09:53.59
ETHER: Packet size = 210 bytes
ETHER: Destination = 8:0:20:1:3d:94, Sun
ETHER: Source      = 8:0:69:1:5f:e, Silicon Graphics
ETHER: Ethertype = 0800 (IP)
ETHER:
IP: ----- IP Header -----
IP:
IP: Version = 4, header length = 20 bytes
IP: Type of service = 00
IP:   ..0. .... = routine
IP:   ...0 .... = normal delay
IP:   .... 0... = normal throughput
IP:   .... .0.. = normal reliability
IP: Total length = 196 bytes
IP: Identification 19846
IP: Flags = 0X
IP:  .0.. .... = may fragment
IP:  ..0. .... = more fragments
IP: Fragment offset = 0 bytes
IP: Time to live = 255 seconds/hops
IP: Protocol = 17 (UDP)
IP: Header checksum = 18DC
IP: Source address = 129.144.40.222, boutique
IP: Destination address = 129.144.40.200, sunroof
IP:
UDP: ----- UDP Header -----
UDP:
UDP: Source port = 1023
UDP: Destination port = 2049 (Sun RPC)
UDP: Length = 176
UDP: Checksum = 0
UDP:
TSOL: ----- TSOL SECURITY ATTRIBUTES -----
TSOL:
TSOL: SM Type = 0x0002, Version = 0x3032
TSOL: Total Length = 200
TSOL: Attribute Type = 4 (Raw)
TSOL: Attributes Length = 192
TSOL: Domain = 0x00000000
TSOL: Generation = 0x00000000
TSOL: Attribute Mask = 0x0000856f
TSOL: Attribute List:
TSOL: Sensitivity Label = SECRET
TSOL: Session ID = 35
TSOL: Clearance = TOP SECRET

```

snoop(1M)

EXAMPLE 2 Sample TSOL packets *(Continued)*

```
TSOL:   Information Label = SECRET ALL EYES
TSOL:   Effective Privilege Mask = 0x20800041084020000200108020000000
TSOL:           89
TSOL:           99
TSOL:           file_dac_read
TSOL:           file_mac_read
TSOL:           ipc_owner
TSOL:           net_downgrade_sl
TSOL:           net_rawaccess
TSOL:           net_upgrade_sl
TSOL:           proc_owner
TSOL:           sys_trans_label
TSOL:           win_upgrade_sl
TSOL:   Process ID = 22540
TSOL:   Effective User ID = 27042
TSOL:   Effective Group ID = 100
TSOL:   Process Attributes Flags = 0x00000001
TSOL:           Trusted Path Flag = 1
TSOL:           Privilege Debug Flag = 0
TSOL:           Trusted Net Process Flag = 0
TSOL:           Label Translation Flags = 0x0
TSOL:           Label View Flags = 0x0
TSOL:
RPC:   ----- SUN RPC Header -----
RPC:
RPC:   Transaction id = 665905
RPC:   Type = 0 (Call)
RPC:   RPC version = 2
RPC:   Program = 100003 (NFS), version = 2, procedure = 1
RPC:   Credentials: Flavor = 1 (Unix), len = 32 bytes
RPC:       Time = 06-Mar-90 07:26:58
RPC:       Hostname = boutique
RPC:       Uid = 0, Gid = 1
RPC:       Groups = 1
RPC:   Verifier : Flavor = 0 (None), len = 0 bytes
RPC:
NFS:   ----- SUN NFS -----
NFS:
NFS:   Proc = 11 (Rename)
NFS:   File handle = 0000164300000000100080000305A1C47
NFS:               597A0000000800002046314AFC450000
NFS:   File name = MTra00192
NFS:   File handle = 0000164300000000100080000305A1C47
NFS:               597A0000000800002046314AFC450000
NFS:   File name = .nfs08
NFS:
```

EXAMPLE 3 Sample NFS packets

View just the NFS packets between sunroof and boutique:

```
example$ snoop -i pkts rpc nfs and sunroof and boutique
```

EXAMPLE 3 Sample NFS packets (Continued)

```

1  0.0000  boutique -> sunroof    NFS C GETATTR FH=8E6C
2  0.0046  sunroof -> boutique    NFS R GETATTR OK
3  0.0080  boutique -> sunroof    NFS C RENAME FH=8E6C MTra00192 to .nfs08

```

To save these packets to a new capture file:

```
example# snoop -i pkts -o pkts.nfs rpc nfs sunroof boutique
```

To view encapsulated packets, there will be an indicator of encapsulation:

```
example# snoop ip-in-ip
sunroof -> boutique ICMP Echo request      (1 encap)
```

If -V is used on an encapsulated packet:

```
example# snoop -V ip-in-ip
sunroof -> boutique ETHER Type=0800 (IP), size = 118 bytes
sunroof -> boutique IP  D=129.144.40.222 S=129.144.40.200 LEN=104, ID=27497
sunroof -> boutique IP  D=10.1.1.2 S=10.1.1.1 LEN=84, ID=27497
sunroof -> boutique ICMP Echo request

```

EXAMPLE 4 Setting Up A More Efficient Filter

To set up a more efficient filter, the following filters should be used toward the end of the expression, so that the first part of the expression can be set up in the kernel: greater, less, port, rpc, nofrag, and relop. The presence of OR makes it difficult to split the filtering when using these primitives that cannot be set in the kernel. Instead, use parenthesis to enforce the primitives that should be OR'd.

To capture packets between funky and pinky of type tcp or udp on port 80:

```
example# snoop funky and pinky and port 80 and tcp or udp
```

Since the primitive port cannot be handled by the kernel filter, and there is also an OR in the expression, a more efficient way to filter is to move the OR to the end of the expression and to use parenthesis to enforce the OR between tcp and udp:

```
example# snoop funky and pinky and (tcp or udp) and port 80
```

EXIT STATUS	0	Successful completion.
	1	An error occurred.
FILES	/dev/audio	Symbolic link to the system's primary audio device.
	/dev/null	The null file.
	/etc/hosts	Host name database.
	/etc/rpc	RPC program number database.
	/etc/services	Internet services and aliases.
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:	

snoop(1M)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

Except when using the `-i` option alone, this program should be run with an effective user ID of 0 to open the network device. The `file_dac_read` and `file_dac_write` privileges can override this restriction. This program must inherit the `sys_net_config` privilege. It must also inherit the `sys_devices` privilege for drivers which either use the `gld(7D)` facility, for example, an `elx1` driver, or make a `drv_priv(9F)` call.

The `sectype` primitive described under OPERANDS is new in the Trusted Solaris environment.

TSIX token mapping requests and responses have a special Session ID value that prevents them from being received by a process that does not have the Trusted Path process attribute. The `snoop` command must have the Trusted Path process attribute to capture TSIX token mapping packets.

Trusted Solaris 8 4/01 Reference Manual

`netstat(1M)`

`hosts(4)`, `rpc(4)`, `services(4)`, `attributes(5)`, `audio(7I)`, `bufmod(7M)`, `dlpi(7P)`, `gld(7D)`, `iee(7D)`, `le(7D)`, `pfmod(7M)`, `tun(7M)`, `drv_priv(9F)`

WARNINGS

The processing overhead is much higher for realtime packet interpretation. Consequently, the packet drop count may be higher. For more reliable capture, output raw packets to a file using the `-o` option and analyze the packets off-line.

Unfiltered packet capture imposes a heavy processing load on the host computer—particularly if the captured packets are interpreted realtime. This processing load further increases if verbose options are used. Since heavy use of `snoop` may deny computing resources to other processes, it should not be used on production servers. Heavy use of `snoop` should be restricted to a dedicated computer.

`snoop` does not reassemble IP fragments. Interpretation of higher level protocol halts at the end of the first IP fragment.

`snoop` may generate extra packets as a side-effect of its use. For example it may use a network name service (NIS or NIS+) to convert IP addresses to host names for display. Capturing into a file for later display can be used to postpone the address-to-name mapping until after the capture session is complete. Capturing into an NFS-mounted file may also generate extra packets.

Setting the `snaplen` (`-s` option) to small values may remove header information that is needed to interpret higher level protocols. The exact cutoff value depends on the network and protocols being used. For NFS Version 2 traffic using UDP on 10 Mb/s ethernet, do not set `snaplen` less than 150 bytes. For NFS Version 3 traffic using TCP on 100 Mb/s ethernet, `snaplen` should be 250 bytes or more.

snoop(1M)

snoop requires information from an RPC request to fully interpret an RPC reply. If an RPC reply in a capture file or packet range does not have a request preceding it, then only the RPC reply header will be displayed.

spray(1M)

NAME	spray – Spray packets								
SYNOPSIS	/usr/sbin/spray [-c <i>count</i>] [-d <i>delay</i>] [-l <i>length</i>] [-t <i>nettype</i>] <i>host</i>								
DESCRIPTION	<p>spray sends a one-way stream of packets to <i>host</i> using RPC, and reports how many were received, as well as the transfer rate. The <i>host</i> argument can be either a name or an Internet address. If the <i>host</i> is a broadcast address, this program needs to inherit the net_broadcast privilege to run properly.</p> <p>spray is not useful as a networking benchmark as it uses unreliable connectionless transports, (upd for example). spray can report a large number of packets dropped when the drops were caused by spray sending packets faster than they can be buffered locally (before the packets get to the network medium).</p>								
OPTIONS	<table><tr><td>-ccount</td><td>Specify how many packets to send. The default value of <i>count</i> is the number of packets required to make the total stream size 100000 bytes.</td></tr><tr><td>-ddelay</td><td>Specify how many microseconds to pause between sending each packet. The default is 0.</td></tr><tr><td>-llength</td><td>The <i>length</i> parameter is the numbers of bytes in the Ethernet packet that holds the RPC call message. Since the data is encoded using XDR, and XDR only deals with 32 bit quantities, not all values of <i>length</i> are possible, and spray rounds up to the nearest possible value. When <i>length</i> is greater than 1514, then the RPC call can no longer be encapsulated in one Ethernet packet, so the <i>length</i> field no longer has a simple correspondence to Ethernet packet size. The default value of <i>length</i> is 86 bytes (the size of the RPC and UDP headers).</td></tr><tr><td>-tnettype</td><td>Specify class of transports. Defaults to netpath. See rpc(3NSL) for a description of supported classes.</td></tr></table>	-ccount	Specify how many packets to send. The default value of <i>count</i> is the number of packets required to make the total stream size 100000 bytes.	-ddelay	Specify how many microseconds to pause between sending each packet. The default is 0.	-llength	The <i>length</i> parameter is the numbers of bytes in the Ethernet packet that holds the RPC call message. Since the data is encoded using XDR, and XDR only deals with 32 bit quantities, not all values of <i>length</i> are possible, and spray rounds up to the nearest possible value. When <i>length</i> is greater than 1514, then the RPC call can no longer be encapsulated in one Ethernet packet, so the <i>length</i> field no longer has a simple correspondence to Ethernet packet size. The default value of <i>length</i> is 86 bytes (the size of the RPC and UDP headers).	-tnettype	Specify class of transports. Defaults to netpath. See rpc(3NSL) for a description of supported classes.
-ccount	Specify how many packets to send. The default value of <i>count</i> is the number of packets required to make the total stream size 100000 bytes.								
-ddelay	Specify how many microseconds to pause between sending each packet. The default is 0.								
-llength	The <i>length</i> parameter is the numbers of bytes in the Ethernet packet that holds the RPC call message. Since the data is encoded using XDR, and XDR only deals with 32 bit quantities, not all values of <i>length</i> are possible, and spray rounds up to the nearest possible value. When <i>length</i> is greater than 1514, then the RPC call can no longer be encapsulated in one Ethernet packet, so the <i>length</i> field no longer has a simple correspondence to Ethernet packet size. The default value of <i>length</i> is 86 bytes (the size of the RPC and UDP headers).								
-tnettype	Specify class of transports. Defaults to netpath. See rpc(3NSL) for a description of supported classes.								
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWcsu</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu				
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWcsu								
SUMMARY OF TRUSTED SOLARIS CHANGES Trusted Solaris 8 4/01 Reference Manual	<p>If the <i>host</i> is a broadcast address, this program needs to inherit the net_broadcast privilege to run properly.</p> <p>rpc(3NSL)</p> <p>attributes(5)</p>								

NAME	statd – Network status monitor				
SYNOPSIS	<code>/usr/lib/nfs/statd</code>				
DESCRIPTION	<p>statd is an intermediate version of the status monitor. It interacts with lockd(1M) to provide the crash and recovery functions for the locking services on NFS. statd keeps track of the clients with processes which hold locks on a server. When the server reboots after a crash, statd sends a message to the statd on each client indicating that the server has rebooted. The client statd processes then inform the lockd on the client that the server has rebooted. The client lockd then attempts to reclaim the lock(s) from the server.</p> <p>statd on the client host also informs the statd on the server(s) holding locks for the client when the client has rebooted. In this case, the statd on the server informs its lockd that all locks held by the rebooting client should be released, allowing other processes to lock those files.</p>				
FILES	<p><code>/var/statmon/sm</code> Lists hosts and network addresses to be contacted after a reboot.</p> <p><code>/var/statmon/sm.bak</code> Lists hosts and network addresses that could not be contacted after last reboot.</p> <p><code>/var/statmon/state</code> Includes a number which changes during a reboot.</p> <p><code>/usr/include/rpcsvc/sm_inter.x</code> Contains the rpcgen source code for the interface services provided by the statd daemon.</p>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>statd must be started with a UID of 0, a sensitivity label of ADMIN_LOW, and a clearance of ADMIN_HIGH. It must be started from the Trusted Path and must have these privileges: net_mac_read, net_privaddr, and net_upgrade_sl.</p> <p>statd creates the directory <code>/var/statmon/sm</code> and all its files, the directory <code>/var/statmon/sm.bak</code> and all its files, and the file <code>/var/statmon/state</code> at the sensitivity label ADMIN_LOW. The directories <code>/var/statmon/sm</code> and <code>/var/statmon/sm.bak</code> are created with the mode 700.</p>				
Trusted Solaris 8 4/01 Reference Manual	lockd(1M)				
See Also	attributes(5)				
Reference Manual NOTES	The crash of a server is only detected upon its recovery.				

su(1M)

NAME	su – become another user
SYNOPSIS	su [-] [<i>username</i> [<i>arg...</i>]]
DESCRIPTION	<p>The su command enables one to become another user without logging off. The default <i>username</i> is <i>root</i>. If the <i>username</i> is a role, the current user must have been assigned the role, but role assumption is generally not useful with this command. Roles usually require trusted path and the <i>admin_low</i> label.</p> <p>To use su, the appropriate password must be supplied unless su inherits the <i>proc_setid</i> privilege. By default, authentication must be done through the trusted path, as is the case when running in a role workspace. If the su command is executed without the trusted path, an authorization check is made if the option <i>su_auth_check_on</i> is specified in <i>pam.conf</i>(4) for the su service module <i>pam_tp_auth</i>(5).</p> <p>The authorization checked by this module, <i>solaris.login.su</i>, should be assigned to the target <i>username</i> rather than the current user.</p> <p>If authentication is successful, su creates a new shell process that has the real and effective user ID, group IDs, and supplementary group list set to those of the specified <i>username</i>.</p> <p>Any additional arguments given on the command line are passed to the new shell. When using programs such as <i>sh</i>, an <i>arg</i> of the form <i>-c string</i> executes <i>string</i> using the shell and an <i>arg</i> of <i>-r</i> gives the user a restricted shell.</p> <p>The following statements are true if the login shell is <i>/usr/bin/sh</i> or an empty string (which defaults to <i>/usr/bin/sh</i>) in the specific user's password file entry. If the first argument to su is a dash (-), the environment will be changed to what would be expected if the user actually logged in as the specified user. Otherwise, the environment is passed along, with the exception of <i>\$PATH</i>, which is controlled by <i>PATH</i> and <i>SUPATH</i> in <i>/etc/default/su</i>. Additionally, the user's project ID is set if the dash argument is present. See <i>settaskid</i>(2).</p> <p>All attempts to become another user using su are logged in the log file <i>/var/adm/sulog</i> (see <i>sulog</i>(4)).</p>
SECURITY	<p>su uses <i>pam</i>(3PAM) for authentication and account management. The PAM configuration policy, listed through <i>/etc/pam.conf</i>, specifies the modules to be used for su. Here is a partial <i>pam.conf</i> file with entries for the su command using the UNIX authentication and account management.</p> <pre>su auth requisite /usr/lib/security/pam_unix.so.1 su_auth_check_on su auth sufficient /usr/lib/security/pam_tp_auth.so.1 su account requisite /usr/lib/security/pam_roles.so.1 su account required /usr/lib/security/pam_unix.so.1 su account required /usr/lib/security/pam_tsol.so.1</pre>

	<p>If there are no entries for the <code>su</code> service, then the entries for the "other" service will be used. If multiple authentication modules are listed, then the user may be prompted for multiple passwords.</p>	
EXAMPLES	EXAMPLE 1 Becoming User <code>bin</code> While Retaining Your Previously Exported Environment	
	<p>To become user <code>bin</code> while retaining your previously exported environment, execute:</p> <pre>example% su bin</pre>	
	EXAMPLE 2 Becoming User <code>bin</code> and Changing to <code>bin</code> 's Login Environment	
	<p>To become user <code>bin</code> but change the environment to what would be expected if <code>bin</code> had originally logged in, execute:</p> <pre>example% su - bin</pre>	
	EXAMPLE 3 Executing command with user <code>bin</code> 's Environment and Permissions	
	<p>To execute command with the temporary environment and permissions of user <code>bin</code>, type:</p> <pre>example% su - bin -c "command args"</pre>	
ENVIRONMENT VARIABLES	<p>Variables with <code>LD_</code> prefix are removed for security reasons. Thus, <code>su bin</code> will not retain previously exported variables with <code>LD_</code> prefix while becoming user <code>bin</code>.</p>	
	<p>If any of the <code>LC_*</code> variables (<code>LC_CTYPE</code>, <code>LC_MESSAGES</code>, <code>LC_TIME</code>, <code>LC_COLLATE</code>, <code>LC_NUMERIC</code>, and <code>LC_MONETARY</code>) (see <code>environ(5)</code>) are not set in the environment, the operational behavior of <code>su</code> for each corresponding locale category is determined by the value of the <code>LANG</code> environment variable. If <code>LC_ALL</code> is set, its contents are used to override both the <code>LANG</code> and the other <code>LC_*</code> variables. If none of the above variables are set in the environment, the "C" (U.S. style) locale determines how <code>su</code> behaves.</p>	
	<code>LC_CTYPE</code>	Determines how <code>su</code> handles characters. When <code>LC_CTYPE</code> is set to a valid value, <code>su</code> can display and handle text and filenames containing valid characters for that locale. <code>su</code> can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. <code>su</code> can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.
	<code>LC_MESSAGES</code>	Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).
FILES	<code>\$HOME/.profile</code>	user's login commands for <code>sh</code> and <code>ksh</code>
	<code>/etc/passwd</code>	system's password file

su(1M)

/etc/profile	system-wide sh and ksh login commands
/var/adm/sulog	log file
/etc/default/su	the default parameters in this file are:
SULOG	If defined, all attempts to su to another user are logged in the indicated file.
CONSOLE	If defined, all attempts to su to root are logged on the console.
PATH	Default path. (/usr/bin:)
SUPATH	Default path for a user invoking su to root. (/usr/sbin:/usr/bin)
SYSLOG	Determines whether the syslog(3C) LOG_AUTH facility should be used to log all su attempts. LOG_NOTICE messages are generated for su's to root, LOG_INFO messages are generated for su's to other users, and LOG_CRIT messages are generated for failed su attempts.
SLEEPTIME	If present, sets the number of seconds to wait before login failure is printed to the screen and another login attempt is allowed. Default is 4 seconds. Minimum is 0 seconds. Maximum is 5 seconds.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

Trusted Solaris 8
4/01 Reference
Manual

The specified *username* may need the the authorization `solaris.login.su` if the trusted path policy is enabled. Users should use the Trusted Path menu to assume a role.

login(1), roles(1), pam_tp_auth(5)
csh(1), env(1), ksh(1), sh(1), syslogd(1M), settaskid(2), pam(3PAM),
syslog(3C), pam.conf(4), passwd(4), sulog(4), attributes(5), environ(5),
pam_unix(5)

NAME	swap – Swap administrative interface
SYNOPSIS	<pre> /usr/sbin/swap -a swapname [swaplow] [swaplen] /usr/sbin/swap -d swapname [swaplow] /usr/sbin/swap -l /usr/sbin/swap -s </pre>
DESCRIPTION	The swap utility provides a method of adding, deleting, and monitoring the system swap areas used by the memory manager.
OPTIONS	<p>The following options are supported:</p> <p>-a swapname Add the specified swap area. This option requires appropriate privilege. <i>swapname</i> is the name of the swap file: for example, <code>/dev/dsk/c0t0d0s1</code> or a regular file. <i>swaplow</i> is the offset in 512-byte blocks into the file where the swap area should begin. <i>swaplen</i> is the desired length of the swap area in 512-byte blocks. The value of <i>swaplen</i> can not be less than 16. For example, if <i>n</i> blocks are specified, then (<i>n</i>–1) blocks would be the actual swap length. <i>swaplen</i> must be at least one page in length. One page of memory is equivalent to eight 512-byte blocks. The size of a page of memory can be determined by using the <code>pagesize</code> command. See <code>pagesize(1)</code>. Since the first page of a swap file is automatically skipped, and a swap file needs to be at least one page in length, the minimum size should be a factor of 2 <code>pagesize</code> bytes. The size of a page of memory is machine dependent.</p> <p><i>swaplow</i> + <i>swaplen</i> must be less than or equal to the size of the swap file. If <i>swaplen</i> is not specified, an area will be added starting at <i>swaplow</i> and extending to the end of the designated file. If neither <i>swaplow</i> nor <i>swaplen</i> are specified, the whole file will be used except for the first page. Swap areas are normally added automatically during system startup by the <code>/sbin/swapadd</code> script. This script adds all swap areas which have been specified in the <code>/etc/vfstab</code> file; for the syntax of these specifications, see <code>vfstab(4)</code>.</p> <p>To use an NFS or local file-system <i>swapname</i>, you should first create a file using <code>mkfile(1M)</code>. A local file-system swap file can now be added to the running system by just running the <code>swap -a</code> command. For NFS mounted swap files, the server needs to export the file. Do this by performing the following steps:</p> <ol style="list-style-type: none"> 1. Add the following line to <code>/etc/dfs/dfstab</code>: <pre>share -F nfs -o rw=clientname,root=clientname path-to-swap-file</pre> 2. Run <code>shareall(1M)</code>. 3. Have the client add the following lines to <code>/etc/vfstab</code>: <pre>server:path-to-swap-file - local-path-to-swap-file nfs - - - local-path-to-swap-file - - swap - - -</pre>

swap(1M)

4. Have the client run mount:

```
# mount local-path-to-swap-file
```

5. The client can then run swap -a to add the swap space:

```
# swap -a local-path-to-swap-file
```

-d *swapname*

Delete the specified swap area. The -d *swapname* option requires appropriate privilege. *swapname* is the name of the swap file: for example, /dev/dsk/c0t0d0s1 or a regular file. *swaplow* is the offset in 512-byte blocks into the swap area to be deleted. If *swaplow* is not specified, the area will be deleted starting at the second page. When the command completes, swap blocks can no longer be allocated from this area and all swap blocks previously in use in this swap area have been moved to other swap areas.

- l List the status of all the swap areas. The output has five columns:

path	The path name for the swap area.
dev	The major/minor device number in decimal if it is a block special device; zeroes otherwise.
swaplo	The <i>swaplow</i> value for the area in 512-byte blocks.
blocks	The <i>swapplen</i> value for the area in 512-byte blocks.
free	The number of 512-byte blocks in this area that are not currently allocated.

The list does not include swap space in the form of physical memory because this space is not associated with a particular swap area.

If swap -l is run while *swapname* is in the process of being deleted (by swap -d), the string INDEL will appear in a sixth column of the swap stats.

- s Print summary information about total swap space usage and availability:

allocated	The total amount of swap space in bytes currently allocated for use as backing store.
reserved	The total amount of swap space in bytes not currently allocated, but claimed by memory mappings for possible future use.
used	The total amount of swap space in bytes that is either allocated or reserved.
available	The total swap space in bytes that is currently available for future reservation and allocation.

swap(1M)

These numbers include swap space from all configured swap areas as listed by the -l option, as well swap space in the form of physical memory.

USAGE Only the first 2 Gbyte of a block device larger than 2 Gbyte in size can be used for swap in swapfs on a 32-bit operating system. With a 64-bit operating system, a block device larger than 2 Gbyte can be fully utilized for swap up to $2^{63}-1$ bytes.

ENVIRONMENT VARIABLES See environ(5) for descriptions of the following environment variables that affect the execution of swap: LC_CTYPE and LC_MESSAGE.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES When used with the -a or -d option, this command needs the sys_mount privilege to succeed.

Trusted Solaris 8 4/01 Reference Manual shareall(1M), vfstab(4)

Solaris 9.0 Reference Manual pagesize(1), mkfile(1M), getpagesize(3C), attributes(5), environ(5), largefile(5)

WARNINGS No check is done to see if a swap area being added overlaps with an existing file system.

sysdef(1M)

NAME	sysdef – Output system definition
SYNOPSIS	<pre>/usr/sbin/sysdef [-n <i>namelist</i>]</pre> <pre>/usr/sbin/sysdef [-h] [-d] [-D]</pre>
DESCRIPTION	<p>outputs the current system definition in tabular form. It lists all hardware devices, as well as pseudo devices, system devices, loadable modules, and the values of selected kernel tunable parameters.</p> <p>It generates the output by analyzing the named bootable operating system file (<i>namelist</i>) and extracting the configuration information from it.</p> <p>The default system <i>namelist</i> is /dev/kmem.</p> <p>The sysdef utility needs the file_mac_read privilege to succeed.</p>
OPTIONS	<p>-n <i>namelist</i> Specifies a <i>namelist</i> other than the default (/dev/kmem). The <i>namelist</i> specified must be a valid bootable operating system.</p> <p>-h Prints the identifier of the current host in hexadecimal. This numeric value is unique across all Sun hosts.</p> <p>-d The output includes the configuration of system peripherals formatted as a device tree.</p> <p>-D For each system peripheral in the device tree, display the name of the device driver used to manage the peripheral.</p>
EXAMPLES	<p>EXAMPLE 1 Sample sysdef output format</p> <p>The following example displays the format of the sysdef -d output:</p> <pre>example% sysdef -d Node 'Sun 4/60', unit #0 (no driver) Node 'options', unit #0 (no driver) Node 'zs', unit #0 Node 'zs', unit #1 Node 'fd', unit #0 Node 'audio', unit #0 Node 'sbus', unit #0 Node 'dma', unit #0 Node 'esp', unit #0 Node 'st', unit #1 (no driver) Node 'st', unit #0 Node 'sd', unit #2 Node 'sd', unit #1 Node 'sd', unit #0 Node 'le', unit #0 Node 'bwtwo', unit #0 Node 'auxiliary-io', unit #0 Node 'interrupt-enable', unit #0 Node 'memory-error', unit #0 Node 'counter-timer', unit #0 Node 'eeprom', unit #0</pre>

EXAMPLE 1 Sample sysdef output format *(Continued)*

FILES /dev/kmem Default operating system image.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu (32-bit)
	SUNWcsxu (64-bit)

SUMMARY OF TRUSTED SQUARES The sysdef utility needs the file_mac_read privilege to succeed.

Trusted Squares 4/01 Reference Manual prtconf(1M)

Changes to Solaris 4/01 Reference Manual hostid(1), nlist(3ELF), attributes(5)

sysh(1M)

NAME	sysh – system shell						
SYNOPSIS	sysh [-acefhiknpPrstuvx] [<i>argument...</i>]						
DESCRIPTION	sysh, the system shell, is a modified version of the Bourne shell, sh(1). sysh is used to control the use of privileges in commands run from the rc scripts. sysh allows any command to be executed but consults profiles for the privileges, user ID (UID), group ID (GID), and sensitivity label (SL) with which the command is to be run.						
USAGE	<p>Refer to the sh(1) man page for a complete usage description. The sysh command adds the setprof command.</p> <p>To list profiles and privileges that are being used by any command in a profile shell, use the smprofile(1) command. See EXAMPLES on the smprofile page for examples of using smprofile list.</p>						
Commands	setprof [<i>profilename</i>]	<p>The setprof command is used to specify a profile other than the boot profile. A profile with a space in its name must be specified within single quotes. For example:</p> <pre>setprof 'custom boot'</pre> <p>Commands in sysh scripts must be specified in a profile only if they need to run with non-default attributes. The default attributes are:</p> <pre>label ADMIN_LOW clearance ADMIN_LOW uid 0 gid 0</pre>					
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td rowspan="2">Availability</td><td>SUNWtsr</td></tr><tr><td>SUNWtsu</td></tr></tbody></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsr	SUNWtsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWtsr						
	SUNWtsu						
Trusted Solaris 8 4/01 Reference Manual	smprofile(1M), exec_attr(4), prof_attr(4)						
See Also	sh(1), attributes(5)						
WARNINGS	If sysh finds that a command needs privileges that sysh does not inherit, a warning message is printed and the command is run with no privileges.						
NOTES	These interfaces are uncommitted. Although they are not expected to change between minor releases of the Trusted Solaris environment, they may.						

tbootparam(1M)

NAME tbootparam – Send a request to rpc.tbootparamd to inform it that a host is in normal (labeled) state now

SYNOPSIS **/usr/sbin/tbootparam** *server_host client_host*
/usr/sbin/tbootparam *client_host*

DESCRIPTION The first form informs the server *server_host* that the host *client_host* is now in the normal state.

The second form broadcasts a message to all `rpc.tbootparamd` processes listening that the host *client_host* is now in the normal state.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWtsu

Trusted Solaris 8
4/01 Reference
Manual
SUNWtsu
Reference Manual

`rpc.tbootparamd(1M)`, `chstate(2)`.
`attributes(5)`

telinit(1M)

NAME	init, telinit – Process control initialization								
SYNOPSIS	<pre>/sbin/init [0123456abcQqSs]</pre> <pre>/etc/telinit [0123456abcQqSs]</pre>								
DESCRIPTION	init is a general process spawner. Its primary role is to create processes from information stored in the file <code>/etc/inittab</code> .								
Run Level Defined	At any given time, the system is in one of eight possible run levels. A run level is a software configuration under which only a selected group of processes exists. Processes spawned by <code>init</code> for each of these run levels are defined in <code>/etc/inittab</code> . <code>init</code> can be in one of eight run levels, 0–6 and S or s (S and s are identical). The run level changes when a privileged user runs <code>/sbin/init</code> . This sends appropriate signals to the original <code>init</code> spawned by the operating system at boot time, saying which run level to invoke.								
init and System Booting	<p>When the system is booted, <code>init</code> is invoked and the following occurs. First, it reads <code>/etc/default/init</code> to set environment variables. This is typically where TZ (time zone) and locale-related environments such as LANG or LC_CTYPE get set.</p> <p><code>init</code> then looks in <code>/etc/inittab</code> for the <code>initdefault</code> entry [see <code>inittab(4)</code>]. If the <code>initdefault</code> entry:</p> <table> <tr> <td>exists</td><td><code>init</code> usually uses the run level specified in that entry as the initial run level to enter.</td></tr> <tr> <td>does not exist</td><td><code>/etc/inittab</code>, <code>init</code> asks the user to enter a run level from the system console.</td></tr> <tr> <td>S or s</td><td><code>init</code> goes to the single-user state. In this state, the system console device (<code>/dev/console</code>) is opened for reading and writing and the command <code>/sbin/su</code>, (see <code>su(1M)</code>), is invoked. Use either <code>init</code> or <code>telinit</code> to change the run level of the system. Note that if the shell is terminated (using an end-of-file), <code>init</code> only re-initializes to the single-user state if <code>/etc/inittab</code> does not exist.</td></tr> <tr> <td>0–6</td><td><code>init</code> enters the corresponding run level. Run levels 0, 5, and 6 are reserved states for shutting the system down. Run levels 2, 3, and 4 are available as multi-user operating states.</td></tr> </table> <p>If this is the first time since power up that <code>init</code> has entered a run level other than single-user state, <code>init</code> first scans <code>/etc/inittab</code> for <code>boot</code> and <code>bootwait</code> entries (see <code>inittab(4)</code>). These entries are performed before any other processing of <code>/etc/inittab</code> takes place, providing that the run level entered matches that of the entry. In this way any special initialization of the operating system, such as mounting</p>	exists	<code>init</code> usually uses the run level specified in that entry as the initial run level to enter.	does not exist	<code>/etc/inittab</code> , <code>init</code> asks the user to enter a run level from the system console.	S or s	<code>init</code> goes to the single-user state. In this state, the system console device (<code>/dev/console</code>) is opened for reading and writing and the command <code>/sbin/su</code> , (see <code>su(1M)</code>), is invoked. Use either <code>init</code> or <code>telinit</code> to change the run level of the system. Note that if the shell is terminated (using an end-of-file), <code>init</code> only re-initializes to the single-user state if <code>/etc/inittab</code> does not exist.	0–6	<code>init</code> enters the corresponding run level. Run levels 0, 5, and 6 are reserved states for shutting the system down. Run levels 2, 3, and 4 are available as multi-user operating states.
exists	<code>init</code> usually uses the run level specified in that entry as the initial run level to enter.								
does not exist	<code>/etc/inittab</code> , <code>init</code> asks the user to enter a run level from the system console.								
S or s	<code>init</code> goes to the single-user state. In this state, the system console device (<code>/dev/console</code>) is opened for reading and writing and the command <code>/sbin/su</code> , (see <code>su(1M)</code>), is invoked. Use either <code>init</code> or <code>telinit</code> to change the run level of the system. Note that if the shell is terminated (using an end-of-file), <code>init</code> only re-initializes to the single-user state if <code>/etc/inittab</code> does not exist.								
0–6	<code>init</code> enters the corresponding run level. Run levels 0, 5, and 6 are reserved states for shutting the system down. Run levels 2, 3, and 4 are available as multi-user operating states.								

	file systems, can take place before users are allowed onto the system. <code>init</code> then scans <code>/etc/inittab</code> and executes all other entries that are to be processed for that run level.												
	To spawn each process in <code>/etc/inittab</code> , <code>init</code> reads each entry and for each entry that should be respawned, it forks a child process. After it has spawned all of the processes specified by <code>/etc/inittab</code> , <code>init</code> waits for one of its descendant processes to die, a powerfail signal, or a signal from another <code>init</code> or <code>telinit</code> process to change the system's run level. When one of these conditions occurs, <code>init</code> re-examines <code>/etc/inittab</code> .												
inittab Additions	New entries can be added to <code>/etc/inittab</code> at any time; however, <code>init</code> still waits for one of the above three conditions to occur before re-examining <code>/etc/inittab</code> . To get around this, <code>init Q</code> or <code>init q</code> command wakes <code>init</code> to re-examine <code>/etc/inittab</code> immediately.												
	When <code>init</code> comes up at boot time and whenever the system changes from the single-user state to another run state, <code>init</code> sets the <code>ioctl(2)</code> states of the console to those modes saved in the file <code>/etc/iotl.syscon</code> . <code>init</code> writes this file whenever the single-user state is entered.												
Run Level Changes	When a run level change request is made, <code>init</code> sends the warning signal (<code>SIGTERM</code>) to all processes that are undefined in the target run level. <code>init</code> waits five seconds before forcibly terminating these processes by sending a kill signal (<code>SIGKILL</code>).												
	When <code>init</code> receives a signal telling it that a process it spawned has died, it records the fact and the reason it died in <code>/var/adm/utmpx</code> and <code>/var/adm/wtmpx</code> if it exists (see <code>who(1)</code>). A history of the processes spawned is kept in <code>/var/adm/wtmpx</code> .												
	If <code>init</code> receives a powerfail signal (<code>SIGPWR</code>) it scans <code>/etc/inittab</code> for special entries of the type <code>powerfail</code> and <code>powerwait</code> . These entries are invoked (if the run levels permit) before any further processing takes place. In this way <code>init</code> can perform various cleanup and recording functions during the powerdown of the operating system.												
/etc/defaults/init File	Default values can be set for the following flags in <code>/etc/default/init</code> . For example: <code>TZ=US/Pacific</code>												
	<table> <tr> <td><code>TZ</code></td><td>Either specifies the timezone information (see <code>ctime(3C)</code>) or the name of a timezone information file <code>/usr/share/lib/zoneinfo</code>.</td></tr> <tr> <td><code>LC_CTYPE</code></td><td>Character characterization information.</td></tr> <tr> <td><code>LC_MESSAGES</code></td><td>Message translation.</td></tr> <tr> <td><code>LC_MONETARY</code></td><td>Monetary formatting information.</td></tr> <tr> <td><code>LC_NUMERIC</code></td><td>Numeric formatting information.</td></tr> <tr> <td><code>LC_TIME</code></td><td>Time formatting information.</td></tr> </table>	<code>TZ</code>	Either specifies the timezone information (see <code>ctime(3C)</code>) or the name of a timezone information file <code>/usr/share/lib/zoneinfo</code> .	<code>LC_CTYPE</code>	Character characterization information.	<code>LC_MESSAGES</code>	Message translation.	<code>LC_MONETARY</code>	Monetary formatting information.	<code>LC_NUMERIC</code>	Numeric formatting information.	<code>LC_TIME</code>	Time formatting information.
<code>TZ</code>	Either specifies the timezone information (see <code>ctime(3C)</code>) or the name of a timezone information file <code>/usr/share/lib/zoneinfo</code> .												
<code>LC_CTYPE</code>	Character characterization information.												
<code>LC_MESSAGES</code>	Message translation.												
<code>LC_MONETARY</code>	Monetary formatting information.												
<code>LC_NUMERIC</code>	Numeric formatting information.												
<code>LC_TIME</code>	Time formatting information.												

telinit(1M)

	LC_ALL	If set, all other LC_* environmental variables take on this value.
	LANG	If LC_ALL is not set, and any particular LC_* is also not set, the value of LANG is used for that particular environmental variable.
telinit	telinit, which is linked to /sbin/init, is used to direct the actions of init. It takes a one-character argument and signals init to take the appropriate action.	
SECURITY	init uses pam(3PAM) for session management. The PAM configuration policy, listed through /etc/pam.conf, specifies the session management module to be used for init. Here is a partial pam.conf file with entries for init using the UNIX session management module.	
	<pre>init session required /usr/lib/security/pam_unix.so.1</pre>	
	If there are no entries for the init service, then the entries for the "other" service will be used.	
OPTIONS	0	Go into firmware.
	1	Put the system in system administrator mode. All local file systems are mounted. Only a small set of essential kernel processes are left running. This mode is for administrative tasks such as installing optional utility packages. All files are accessible and no users are logged in on the system.
	2	Put the system in multi-user mode. All multi-user environment terminal processes and daemons are spawned. This state is commonly referred to as the multi-user state.
	3	Extend multi-user mode by making local resources available over the network.
	4	Is available to be defined as an alternative multi-user environment configuration. It is not necessary for system operation and is usually not used.
	5	Shut the machine down so that it is safe to remove the power. Have the machine remove power, if possible.
	6	Stop the operating system and reboot to the state defined by the initdefault entry in /etc/inittab.
	a, b, c	Process only those /etc/inittab entries having the a, b, or c run level set. These are pseudo-states, which may be defined to run certain commands, but which do not cause the current run level to change.
	Q, q	Re-examine /etc/inittab.
	S, s	Enter single-user mode. This is the only run level that doesn't require the existence of a properly formatted /etc/inittab file. If this file does not exist, then by default, the only legal run level

that `init` can enter is the single-user mode. When in single-user mode, the filesystems required for basic system operation will be mounted. When the system comes down to single-user mode, these file systems will remain mounted (even if provided by a remote file server), and any other local filesystems will also be left mounted. During the transition down to single-user mode, all processes started by `init` or `init.d` scripts that should only be running in multi-user mode are killed. In addition, any process that has a `utmpx` entry will be killed. This last condition insures that all port monitors started by the SAC are killed and all services started by these port monitors, including `ttymon` login services, are killed.

FILES	<code>/etc/inittab</code>	Controls process dispatching by <code>init</code> .
	<code>/var/adm/utmpx</code>	User access and administration information
	<code>/var/adm/wtmpx</code>	History of user access and administration information
	<code>/etc/ioctl.syscon</code>	System console states.
	<code>/dev/console</code>	System console device.
	<code>/etc/default/init</code>	Environment variables.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

**SUMMARY OF
TRUSTED
SOLARIS
CHANGES**
Trusted Solaris 3
4/01 Reference
Manual
SunOS 5.8

`init` requires privilege to run in the Trusted Solaris environment.

`login(1)`, `kill(2)`, `inittab(4)`

`sh(1)`, `stty(1)`, `who(1)`, `shutdown(1M)`, `su(1M)`, `ttymon(1M)`, `ioctl(2)`, `ctime(3C)`, `pam(3PAM)`, `pam.conf(4)`, `utmpx(4)`, `attributes(5)`, `pam_unix(5)`, `termio(7I)`

DIAGNOSTICS

If `init` finds that it is respawning an entry from `/etc/inittab` more than ten times in two minutes, assumes that there is an error in the command string in the entry, and generates an error message on the system console. It will then refuse to respawn this entry until either five minutes has elapsed or it receives a signal from a user-spawned `init` or `telinit`. This prevents `init` from eating up system resources when someone makes a typographical error in the `inittab` file, or a program is removed that is referenced in `/etc/inittab`.

NOTES

`init` and `telinit` can be run only by a privileged user.

telinit(1M)

The `S` or `s` state must not be used indiscriminately in `/etc/inittab`. When modifying this file, it is best to avoid adding this state to any line other than `initdefault`.

If a default state is not specified in the `initdefault` entry in `/etc/inittab`, state `6` is entered. Consequently, the system will loop by going to firmware and rebooting continuously.

If the `utmpx` file cannot be created when booting the system, the system will boot to state `"s"` regardless of the state specified in the `initdefault` entry in `/etc/inittab`. This can occur if the `/var` file system is not accessible.

NAME	in.tftpd, tftpd – Internet Trivial File Transfer Protocol server				
SYNOPSIS	in.tftpd [-s] [<i>homedir</i>]				
DESCRIPTION	<p>tftpd is a server that supports the Internet Trivial File Transfer Protocol (TFTP). This server is normally started by inetd(1M) and operates at the port indicated in the tftp Internet service description in the /etc/inetd.conf file. By default, the entry for in.tftpd in etc/inetd.conf is commented out. To make in.tftpd operational, the comment character(s) must be deleted from the file. See inetd.conf(4).</p> <p>Before responding to a request, the server attempts to change its current directory to <i>homedir</i>; the default directory is /tftpboot.</p> <p>The use of tftp does not require an account or password on the remote system. Due to the lack of authentication information, in.tftpd will allow only publicly readable files to be accessed. Files may be written only if they already exist and are publicly writable. Note that this extends the concept of “public” to include all users on all hosts that can be reached through the network; this may not be appropriate on all systems, and its implications should be considered before enabling this service.</p> <p>in.tftpd runs with the user ID and group ID set to [GU] ID_NOBODY under the assumption that no files exist with that owner or group. However, nothing checks this assumption or enforces this restriction.</p>				
OPTIONS	<p>-s Secure. When specified, the directory change to <i>homedir</i> must succeed. The daemon also changes its root directory to <i>homedir</i>.</p>				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>in.tftpd should be started from the trusted path with a UID of 0; it must inherit the proc_chroot, proc_owner, and proc_setid privileges.</p> <p>/etc/inetd.conf Configuration file for inetd.</p>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
Trusted Solaris 8 4/01 Reference Manual	<p>inetd(1M)</p> <p>tftp(1), inetd.conf(4), netconfig(4), attributes(5), ip6(7P)</p> <p>Sollins, K.R., <i>The TFTP Protocol (Revision 2)</i>, RFC 783, Network Information Center, SRI International, Menlo Park, California, June 1981.</p>				

tnchkdb(1M)

NAME	tnchkdb – check file syntax of trusted network databases
SYNOPSIS	<pre> /usr/sbin/tnchkdb /usr/sbin/tnchkdb -t [pathname] /usr/sbin/tnchkdb -h [pathname] /usr/sbin/tnchkdb -t [t_pathname] -h [h_pathname] /usr/sbin/tnchkdb -i [pathname] </pre>
DESCRIPTION	<p>tnchkdb checks the syntax of the tnrhttp(4), tnrhdb(4), or tnidb(4) databases at <i>pathname</i>. (<i>pathname</i> is the full pathname and filename of the file.) If no database is specified, all three databases in /etc/security/tsol are checked. tnchkdb returns an exit status of 0 (true) and no output if the file is syntactically and semantically correct. Otherwise, tnchkdb returns a nonzero (false) exit status and writes an error diagnostic to the standard output file. tnchkdb also examines the label and DAC information on the specified database files and reports mismatches as WARNINGS rather than ERRORS.</p> <p>tnchkdb can be run at any sensitivity label that dominates the sensitivity label of the database file. This restriction can be overridden by the file_mac_read privilege.</p>
OPTIONS	<pre> -t [pathname] Check <i>pathname</i> for proper tnrhttp syntax. If the <i>pathname</i> is not specified, then check /etc/security/tsol/tnrhttp. -h [pathname] Check <i>pathname</i> for proper tnrhdb syntax. If the <i>pathname</i> is not specified, then check /etc/security/tsol/tnrhdb. -t [t_pathname] -h [h_pathname] Check <i>t_pathname</i> for proper tnrhttp syntax and check <i>h_pathname</i> for proper tnrhdb syntax. This option complains about template names assigned in tnrhdb but not defined in tnrhttp. If the <i>pathname</i> is not specified, then check /etc/security/tsol/tnrhttp for the -t option and /etc/security/tsol/tnrhdb for the -h option. -i [pathname] Check <i>pathname</i> for proper tnidb syntax. If the <i>pathname</i> is not specified, then check /etc/security/tsol/tnidb. </pre>
EXAMPLES	<p>EXAMPLE 1 Example Error Message When Checking tnrhttp and tnrhdb</p> <p>The tnchkdb command prints an error message if the tnrhdb entry does not exactly match its tnrhttp template entry. In the following example, a space after 192.168.113.170:tsol in the tnrhdb file causes an error. Note that the \$ are included to indicate the end of the lines, but do not exist in the file.</p> <pre> % grep tsol /etc/security/tsol/tnrhdb # Assume that template tsol is defined in the tnrhttp database.\$ 192.168.113.170:tsol \$ </pre>

EXAMPLE 1 Example Error Message When Checking tnrhtp and tnrhdb*(Continued)*

```
% tnchkdb -t -h
checking /etc/security/tsol/tnrhtp ...
checking /etc/security/tsol/tnrhdb ...
Error: Unknown template name: tsol
done.
```

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWtsu

FILES

```
/etc/security/tsol/tnidb
    Trusted network interface-control database

/etc/security/tsol/tnrhdb
    Trusted network remote-host database

/etc/security/tsol/tnrhtp
    Trusted network remote-host templates
```

Trusted Solaris 8
4/01 Reference
Manual
SUNWtsu
Reference Manual
NOTES

tnd(1M), tnctl(1M), tnidb(4), tnrhdb(4), tnrhtp(4)

attributes(5)

It is possible to have inconsistent but valid configurations of tnrhtp and tnrhdb, since NIS+ may be used to supply missing templates.

tnctl(1M)

NAME	tnctl – Configure Trusted Solaris network-daemon control parameters										
SYNOPSIS	<pre> /usr/sbin/tnctl [-v] [-d <i>debug_level</i>] [-p <i>poll-interval</i>] [-i <i>interface_name</i>] [-h <i>host_name</i>] [-t <i>template_name</i>] [-b <i>ip_address</i>] [-B <i>ip_address</i>] /usr/sbin/tnctl -I <i>tnidb_path</i> /usr/sbin/tnctl -T <i>tnrhtp_path</i> /usr/sbin/tnctl -H <i>tnrhdb_path</i> </pre>										
DESCRIPTION	<p>tnctl provides an interface to send control and configuration messages either to the kernel directly or to tnd(1M).</p> <p>If a local trusted-networking database file is modified, the administrator should issue tnchkdb(1M) to check the syntax, and must also issue tnctl to reload the kernel caches.</p> <p>tnctl must be started from the trusted path; and for the -i, -t, -h, -b, -B, -I, -T, and -H options, it must have the sys_net_config privilege. tnctl can be run at any sensitivity label, except the -h and -H options, which need to run at ADMIN_LOW. The file_mac_read privilege can be used to override this policy.</p>										
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWtsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu						
ATTRIBUTE TYPE	ATTRIBUTE VALUE										
Availability	SUNWtsu										
OPTIONS	<table> <tr> <td>-v</td><td>Turn on verbose mode.</td></tr> <tr> <td>-d <i>debug_level</i></td><td>Turn on debugging for tnd to the level specified by <i>debug_level</i>. <i>debug_level</i> may be 1 or 2; however, currently no distinction is made between the two values. The debug output goes to the log file specified on the tnd command line, or by default to /var/tsxol/tndlog.</td></tr> <tr> <td>-p <i>poll-interval</i></td><td>Set poll interval to <i>poll-interval</i> seconds. The valid range is 0 to 2147483647; a zero value causes tnd to poll the name service databases immediately and then revert to the original poll-interval. This may be useful when changes to tnrhdb(4) or tnrhtp(4) databases are to be made effective immediately.</td></tr> <tr> <td>-i <i>interface_name</i></td><td>Update the kernel-interface cache on the specified <i>interface_name</i>. If the entry does not exist in the database, return an error message.</td></tr> <tr> <td>-h <i>hostname</i></td><td>Update the kernel remote-host cache on the specified <i>hostname</i>. If the entry does not exist in the database, delete the entry from the cache.</td></tr> </table>	-v	Turn on verbose mode.	-d <i>debug_level</i>	Turn on debugging for tnd to the level specified by <i>debug_level</i> . <i>debug_level</i> may be 1 or 2; however, currently no distinction is made between the two values. The debug output goes to the log file specified on the tnd command line, or by default to /var/tsxol/tndlog.	-p <i>poll-interval</i>	Set poll interval to <i>poll-interval</i> seconds. The valid range is 0 to 2147483647; a zero value causes tnd to poll the name service databases immediately and then revert to the original poll-interval. This may be useful when changes to tnrhdb(4) or tnrhtp(4) databases are to be made effective immediately.	-i <i>interface_name</i>	Update the kernel-interface cache on the specified <i>interface_name</i> . If the entry does not exist in the database, return an error message.	-h <i>hostname</i>	Update the kernel remote-host cache on the specified <i>hostname</i> . If the entry does not exist in the database, delete the entry from the cache.
-v	Turn on verbose mode.										
-d <i>debug_level</i>	Turn on debugging for tnd to the level specified by <i>debug_level</i> . <i>debug_level</i> may be 1 or 2; however, currently no distinction is made between the two values. The debug output goes to the log file specified on the tnd command line, or by default to /var/tsxol/tndlog.										
-p <i>poll-interval</i>	Set poll interval to <i>poll-interval</i> seconds. The valid range is 0 to 2147483647; a zero value causes tnd to poll the name service databases immediately and then revert to the original poll-interval. This may be useful when changes to tnrhdb(4) or tnrhtp(4) databases are to be made effective immediately.										
-i <i>interface_name</i>	Update the kernel-interface cache on the specified <i>interface_name</i> . If the entry does not exist in the database, return an error message.										
-h <i>hostname</i>	Update the kernel remote-host cache on the specified <i>hostname</i> . If the entry does not exist in the database, delete the entry from the cache.										

- t *template_name* Update the kernel template cache on the specified *template_name*. If the entry does not exist in the database, return an error message. See WARNINGS about the risks of changing a template when the network is up.
- I *tnidb_path* Load all entries in the *tnidb_path* file into the kernel cache. *tnidb_path* is the full pathname plus filename of the file.
- T *tnrhtp_path* Load all entries in the file *tnrhtp_path* into the kernel cache. *tnrhtp_path* is the full pathname plus filename of the file.
- H *tnrhdb_path* Load all entries in the *tnrhdb_path* file into the kernel cache. *tnrhdb_path* is the full pathname plus filename of the file.
- b *ip_address* Add a remote broadcast address.
- B *ip_address* Delete a remote broadcast address.

FILES /etc/security/tsol/tnidb
Trusted network interface-control database

/etc/security/tsol/tnrhdb
Trusted network remote-host database

/etc/security/tsol/tnrhtp
Trusted network remote-host templates

/etc/nsswitch.conf
Configuration file for the name service switch

Trusted Solaris 8 tninfo(1M), tnd(1M), tnchfdb(1M), nsswitch.conf(4), tnidb(4), tnrhdb(4),
4/01 Reference tnrhtp(4)
Manual

SunOS 5.8 attributes(5)
Reference Manual

NOTES Currently, only level-1 debugging is supported.

WARNINGS Changing a template while the network is up can change the security view of an undetermined number of hosts.

tnd(1M)

NAME	tnd – trusted network daemon								
SYNOPSIS	/usr/sbin/tnd [-d <i>debug_level</i>] [-f <i>logfile</i>] [-p <i>poll-interval</i>] [-n]								
DESCRIPTION	<p>The tnd (trusted network daemon) initializes the kernel with trusted network databases and also reloads the databases on demand. tnd also services requests for tnrhdb(4) templates from tninfo(1M) and the kernel. tnd is started at the beginning of the boot process.</p> <p>tnd loads these databases into the kernel: the remote host database, tnrhdb(4); the remote-host template database, tnrhtp(4); and the interface database, tnidb(4). These databases and their effect on the trusted network are described in their respective man pages. When tnrhdb(4) and tnrhtp(4) and the associated NIS+ tables are changed, tnd also updates the local kernel cache at the predetermined interval.</p> <p>tnd logs its debugging information in a log file (by default, /var/tsol/tndlog) which is set by using the -f option.</p> <p>If a local trusted networking database file is modified, the administrator should issue a tnchkdb(1M) to check the syntax, and must issue a tnctl to reload the kernel caches.</p> <p>tnd must be started from the Trusted Path and inherit these privileges to run: net_privaddr, net_mac_read, net_upgrade_sl, sys_net_config, proc_setclr, proc_setsl. tnd is intended to be started from an rc script and to run at the ADMIN_LOW sensitivity label.</p>								
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWtsu</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu				
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Availability	SUNWtsu								
OPTIONS	<table><tr><td>-d <i>debug_level</i></td><td>Turn on debugging to the level specified by <i>debug_level</i>. <i>debug_level</i> may be 1 or 2; however, currently no distinction is made between the two values. If log file is not specified with the -f option, use /var/tsol/tndlog.</td></tr><tr><td>-f <i>logfile</i></td><td>Set logfile path to <i>logfile</i> for writing debugging information. If <i>logfile</i> already exists, append debugging information to it.</td></tr><tr><td>-p <i>poll-interval</i></td><td>Set poll interval to <i>poll-interval</i> seconds. The default <i>poll-interval</i> is 1800 seconds (30 minutes). The inet svc script uses this flag to set the poll-interval to 120 seconds.</td></tr><tr><td>-n</td><td>Disable polling of name services when starting up. This option is used when tnd is started in an rc script and it is not intended for use when tnd is started otherwise.</td></tr></table>	-d <i>debug_level</i>	Turn on debugging to the level specified by <i>debug_level</i> . <i>debug_level</i> may be 1 or 2; however, currently no distinction is made between the two values. If log file is not specified with the -f option, use /var/tsol/tndlog.	-f <i>logfile</i>	Set logfile path to <i>logfile</i> for writing debugging information. If <i>logfile</i> already exists, append debugging information to it.	-p <i>poll-interval</i>	Set poll interval to <i>poll-interval</i> seconds. The default <i>poll-interval</i> is 1800 seconds (30 minutes). The inet svc script uses this flag to set the poll-interval to 120 seconds.	-n	Disable polling of name services when starting up. This option is used when tnd is started in an rc script and it is not intended for use when tnd is started otherwise.
-d <i>debug_level</i>	Turn on debugging to the level specified by <i>debug_level</i> . <i>debug_level</i> may be 1 or 2; however, currently no distinction is made between the two values. If log file is not specified with the -f option, use /var/tsol/tndlog.								
-f <i>logfile</i>	Set logfile path to <i>logfile</i> for writing debugging information. If <i>logfile</i> already exists, append debugging information to it.								
-p <i>poll-interval</i>	Set poll interval to <i>poll-interval</i> seconds. The default <i>poll-interval</i> is 1800 seconds (30 minutes). The inet svc script uses this flag to set the poll-interval to 120 seconds.								
-n	Disable polling of name services when starting up. This option is used when tnd is started in an rc script and it is not intended for use when tnd is started otherwise.								

tnd(1M)

FILES	/etc/security/tsol/tnidb
	Trusted network interface-control database
	/etc/security/tsol/tnrhdb
	Trusted network remote-host database
	/etc/security/tsol/tnrhtp
	Trusted network remote-host templates
Trusted Solaris 8 4/01 Reference Manual	/var/tsol/tndlog
	Log of tnd debugging information
	/etc/nsswitch.conf
	Configuration file for the name service switch
	SunOS 5.8 Reference Manual
	tnchkdb(1M), tninfo(1M), tnctl(1M), tnidb(4), tnrhdb(4), tnrhtp(4), tndlog(4), nsswitch.conf(4)
	attributes(5)

tninfo(1M)

NAME	tninfo – Print information and statistics about kernel-level network															
SYNOPSIS	/usr/sbin/tninfo [-skcv] [-i [if_name]] [-h [hostname]] [-t [template_name]]															
DESCRIPTION	<p>tninfo provides an interface to retrieve and display kernel-level network information and statistics.</p> <p>tninfo is intended to be run at ADMIN_HIGH and effective user ID 0. The exceptions are the -h, -i , and -t options with an argument, which can be run at any label as any user.</p> <p>For those options that have restrictions, the restrictions can be overridden by these privileges: file_mac_read, sys_trans_label, and file_dac_read. The tninfo executable should be maintained with a sensitivity label of ADMIN_LOW.</p>															
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:															
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWtsu</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu										
ATTRIBUTE TYPE	ATTRIBUTE VALUE															
Availability	SUNWtsu															
OPTIONS	<table><tr><td>-s</td><td>Print the default security structures associated with each socket or stream.</td></tr><tr><td>-k</td><td>Print the network statistics. This is the default option.</td></tr><tr><td>-c</td><td>Print the cache statistics.</td></tr><tr><td>-v</td><td>Turn on verbose mode.</td></tr><tr><td>-i [if_name]</td><td>Display the security structure for the specified interface in the kernel cache. The output should reflect what is specified in the tnidb database. If if_name is not specified, display the entire interface cache.</td></tr><tr><td>-h [hostname]</td><td>Display the security structure for the specified host in the kernel remote-host cache. The output should reflect what is specified in the tnrhdb and tnrhtp databases.</td></tr><tr><td>-t [template_name]</td><td>Display the structure associated with the specified template_name. The output should reflect what is specified in the tnrhtp database. If template_name is not specified, display the entire remote-host template cache. If a field within an entry is not specified (for example, def_uid=empty;), then that field will not be displayed.</td></tr></table>		-s	Print the default security structures associated with each socket or stream.	-k	Print the network statistics. This is the default option.	-c	Print the cache statistics.	-v	Turn on verbose mode.	-i [if_name]	Display the security structure for the specified interface in the kernel cache. The output should reflect what is specified in the tnidb database. If if_name is not specified, display the entire interface cache.	-h [hostname]	Display the security structure for the specified host in the kernel remote-host cache. The output should reflect what is specified in the tnrhdb and tnrhtp databases.	-t [template_name]	Display the structure associated with the specified template_name. The output should reflect what is specified in the tnrhtp database. If template_name is not specified, display the entire remote-host template cache. If a field within an entry is not specified (for example, def_uid=empty;), then that field will not be displayed.
-s	Print the default security structures associated with each socket or stream.															
-k	Print the network statistics. This is the default option.															
-c	Print the cache statistics.															
-v	Turn on verbose mode.															
-i [if_name]	Display the security structure for the specified interface in the kernel cache. The output should reflect what is specified in the tnidb database. If if_name is not specified, display the entire interface cache.															
-h [hostname]	Display the security structure for the specified host in the kernel remote-host cache. The output should reflect what is specified in the tnrhdb and tnrhtp databases.															
-t [template_name]	Display the structure associated with the specified template_name. The output should reflect what is specified in the tnrhtp database. If template_name is not specified, display the entire remote-host template cache. If a field within an entry is not specified (for example, def_uid=empty;), then that field will not be displayed.															
FILES	/etc/security/tsol/tnidb	Trusted network interface-control database														
	/etc/security/tsol/tnrhdb	Trusted network remote-host database														

tninfo(1M)

/etc/security/tsol/tnrhtp Trusted network remote-host templates

tnd(1M), tnctl(1M), tnidb(4), tnrhdb(4), tnrltp(4)

attributes(5)

The kernel's tables can change while tninfo is examining them; the result is incorrect or partial displays.

tokmapctl(1M)

NAME	tokmapctl – Configure token-mapping daemon					
SYNOPSIS	tokmapctl [-H <i>hostname</i>] [-P <i>satmp_port</i>] [-s <i>timeout</i>] [-r <i>retries</i>] [-R <i>retry_interval</i>] [-I [<i>cachesize</i>]] [-F [<i>hostname</i>]] [-m <i>meter_type</i>] [-d <i>level</i>] [-D <i>level</i>] [-l <i>logfile</i>] [-M <i>hostname</i>] [-x]					
DESCRIPTION	<p>tokmapctl provides an interface to send control and configuration requests to a tokmapd process.</p> <p>tokmapctl must be started from the trusted path and must inherit the net_privaddr and net_mac_read privileges. tokmapctl should be run at sensitivity label ADMIN_HIGH.</p>					
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:					
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWtsu</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE					
Availability	SUNWtsu					
OPTIONS	<p>-H <i>hostname</i> Send the control and configuration requests to the tokmapd process on host <i>hostname</i>. If this option is not specified, the request is sent to the tokmapd process on the local host.</p> <p>-P <i>port</i> Send the requests to tokmapd on port number <i>port</i>. This option is intended for debugging only. If this option is not specified, requests are sent to port 90.</p> <p>-s <i>timeout</i> Tell tokmapd to use <i>timeout</i> seconds as its timeout period before retrying a request that has been sent to another token-mapping server but has received no reply. The default is 5 seconds.</p> <p>-r <i>retries</i> Tell tokmapd to use <i>retries</i> as the maximum number of times to retry requests to other token-mapping servers. The default is 5 retries.</p> <p>-R <i>retry_interval</i> Tell tokmapd to use <i>retry_interval</i> milliseconds as its interval between checks for the need to retry requests to other token-mapping servers. The default interval is 100 milliseconds.</p> <p>-I [<i>cachesize</i>] Tell tokmapd to reinitialize its token store. If it is specified, <i>cachesize</i> is used to set the size of the token store in-memory cache. <i>cachesize</i> specifies how many entries of each attribute type to keep in the cache. The default is 10.</p> <p>-F [<i>hostname</i>] Tell tokmapd to flush all tokens for <i>hostname</i> from its token store. If <i>hostname</i> is omitted, tokmapd flushes all tokens for remote hosts.</p> <p>-m <i>meter_type</i> Fetch and display metering data from tokmapd. The allowable values for <i>meter_type</i> are <i>hostlist</i>, <i>general</i>, <i>store</i>, and <i>all</i>.</p>					

	tokmapctl(1M)
	Multiple <code>-m</code> options may be specified to request multiple types of metering data; specify type <code>all</code> to fetch and display all the meter types.
<code>-d level</code>	Set <code>tokmapd</code> debugging level to <i>level</i> . Debugging level 1 produces minimal output showing when messages are sent and received. Level 3 shows the contents of the headers of messages. Level 5 shows detailed information including buffer addresses and contents. Levels above 5 show additional internal information.
<code>-D level</code>	Set <code>tokmapctl</code> debugging level to <i>level</i> . Debugging level 1 produces minimal output showing when messages are sent and received. Level 3 shows the contents of the headers of messages. Level 5 shows detailed information including buffer addresses and contents. Levels above 5 show additional internal information.
<code>-l logfile</code>	Tell <code>tokmapd</code> to write its debugging output to <i>logfile</i> .
<code>-M hostname</code>	Fetch and display metering data from <code>tokmapd</code> for its token-mapping exchanges with host <i>hostname</i> .
<code>-x</code>	Send a request for an orderly shutdown and exit to <code>tokmapd</code> .
	<code>tokmapd(1M)</code>
	<code>attributes(5)</code>
	If the token store becomes too large, use the <code>-I</code> option of <code>tokmapctl</code> to make <code>tokmapd</code> delete the current token store and reinitialize.
	These interfaces are uncommitted. Although they are not expected to change between minor releases of the Trusted Solaris environment, they may.

tokmapd(1M)

NAME	tokmapd – Token-mapping daemon													
SYNOPSIS	/usr/sbin/tokmapd [-d level] [-l logfile] [-c cachesize] [-P satmp_port] [-p kernel_port] [-s timeout] [-r retries] [-R retry_interval] [-f path]													
DESCRIPTION	<p>tokmapd implements the SATMP token-mapping protocol to support the labeling of information transferred over the trusted network. The information is labeled using tokens that represent attribute values. tokmapd is responsible for mapping tokens to attribute values and vice versa. tokmapd accepts token-mapping requests from the kernel and from token-mapping servers on other hosts.</p> <p>tokmapd must be started from the trusted path and must inherit the net_privaddr, proc_setclr, and proc_setsl privileges. tokmapd should be run at sensitivity label ADMIN_HIGH.</p> <p>If tokmapd is stopped and its on-disk cache reinitialized or removed, the machine should be rebooted.</p>													
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:													
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWtsu</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu								
ATTRIBUTE TYPE	ATTRIBUTE VALUE													
Availability	SUNWtsu													
OPTIONS	<table><tr><td>-d level</td><td>Set tokmapd debugging level to level. Debugging level 1 produces minimal output showing when messages are sent and received. Level 3 shows the contents of the headers of messages. Level 5 shows detailed information including buffer addresses and contents. Levels above 5 show additional internal information.</td></tr><tr><td>-l logfile</td><td>Write any debugging output to logfile. If logfile already exists, the debugging output is appended to it. If this option is not specified, the default logfile /var/tsxol/tokmapdlog is used.</td></tr><tr><td>-c cachesize</td><td>Set the size of the token store in-memory cache to cachesize. cachesize specifies how many entries of each attribute type to keep in the cache. The default is 10.</td></tr><tr><td>-P satmp_port</td><td>Listen on satmp_port for SATMP and tokmapctl requests. This option is intended for debugging only. If this option is not specified, port 90 is used.</td></tr><tr><td>-p kernel_port</td><td>Listen on kernel_port for token-mapping requests from the kernel. This option is intended for debugging only. If this option is not specified, port 10800 is used.</td></tr><tr><td>-s timeout</td><td>Use timeout seconds as the timeout period before retrying a request that has been sent to another token-mapping server but has received no reply. If this option is not specified, a timeout interval of 5 seconds is used.</td></tr></table>		-d level	Set tokmapd debugging level to level. Debugging level 1 produces minimal output showing when messages are sent and received. Level 3 shows the contents of the headers of messages. Level 5 shows detailed information including buffer addresses and contents. Levels above 5 show additional internal information.	-l logfile	Write any debugging output to logfile. If logfile already exists, the debugging output is appended to it. If this option is not specified, the default logfile /var/tsxol/tokmapdlog is used.	-c cachesize	Set the size of the token store in-memory cache to cachesize. cachesize specifies how many entries of each attribute type to keep in the cache. The default is 10.	-P satmp_port	Listen on satmp_port for SATMP and tokmapctl requests. This option is intended for debugging only. If this option is not specified, port 90 is used.	-p kernel_port	Listen on kernel_port for token-mapping requests from the kernel. This option is intended for debugging only. If this option is not specified, port 10800 is used.	-s timeout	Use timeout seconds as the timeout period before retrying a request that has been sent to another token-mapping server but has received no reply. If this option is not specified, a timeout interval of 5 seconds is used.
-d level	Set tokmapd debugging level to level. Debugging level 1 produces minimal output showing when messages are sent and received. Level 3 shows the contents of the headers of messages. Level 5 shows detailed information including buffer addresses and contents. Levels above 5 show additional internal information.													
-l logfile	Write any debugging output to logfile. If logfile already exists, the debugging output is appended to it. If this option is not specified, the default logfile /var/tsxol/tokmapdlog is used.													
-c cachesize	Set the size of the token store in-memory cache to cachesize. cachesize specifies how many entries of each attribute type to keep in the cache. The default is 10.													
-P satmp_port	Listen on satmp_port for SATMP and tokmapctl requests. This option is intended for debugging only. If this option is not specified, port 90 is used.													
-p kernel_port	Listen on kernel_port for token-mapping requests from the kernel. This option is intended for debugging only. If this option is not specified, port 10800 is used.													
-s timeout	Use timeout seconds as the timeout period before retrying a request that has been sent to another token-mapping server but has received no reply. If this option is not specified, a timeout interval of 5 seconds is used.													

		tokmapd(1M)
	-r <i>retries</i>	Resend requests to other token-mapping servers a maximum of <i>retries</i> times. If this option is not specified, a retry limit of 5 is used.
	-R <i>retry_interval</i>	Use <i>retry_interval</i> milliseconds as the interval between checks for the need to do retries. The default interval is 100 milliseconds.
	-f <i>path</i>	Place the token store and host-list files in the <i>path</i> directory. If this option is not specified, the files are stored in /etc/security/tsol.
FILES	/etc/security/tsol/tokenadb.pag	Token store file
	/etc/security/tsol/tokenadb.dir	Token store file
	/etc/security/tsol/tokenadb.ir	Token store file
	/etc/security/tsol/tokenadb.hosts	Token store file
	/var/tsol/tokmapdlog	Logfile of debugging output
	tokmapctl(1M)	
	attributes(5)	
	The token store is checked for consistency each time tokmapd is started. If the token store was not properly flushed to disk at the last shutdown, or if other inconsistencies are found, the token-store contents are deleted and the token store is reinitialized.	
	These interfaces are uncommitted. Although they are not expected to change between minor releases of the Trusted Solaris environment, they may.	

traceroute(1M)

NAME	traceroute – Print the route packets take to network host
SYNOPSIS	traceroute [-adFIlnSvx] [-A <i>addr_family</i>] [-c <i>traffic_class</i>] [-f <i>first_hop</i>] [-g <i>gateway</i> [-g <i>gateway...</i>] -r] [-i <i>iface</i>] [-L <i>flow_label</i>] [-m <i>max_hop</i>] [-P <i>pause_sec</i>] [-p <i>port</i>] [-Q <i>max_timeout</i>] [-q <i>nqueries</i>] [-s <i>src_addr</i>] [-t <i>tos</i>] [-w <i>wait_time</i>] host [<i>packetlen</i>]
DESCRIPTION	<p>The Internet is a large and complex aggregation of network hardware, connected together by gateways. Tracking the route a packet follows can be difficult. The utility <code>traceroute</code> traces the route that an IP packet follows to another internet host. In the Trusted Solaris environment, <code>traceroute</code> must be run with the <code>net_rawaccess</code> privilege.</p> <p>The <code>traceroute</code> utility utilizes the both the IPv4 and IPv6 protocols. Use the <code>-A</code> option to override the default behavior. <code>traceroute</code> uses the IPv4 protocol <i>tll</i> (time to live) field or the IPv6 field <i>hop limit</i>. It attempts to elicit an ICMP or ICMP6 <code>TIME_EXCEEDED</code> response from each <i>gateway</i> along the path, and a <code>PORT_UNREACHABLE</code>(or <code>ECHO_REPLY</code> if <code>-I</code> is used) response from the destination host. It starts by sending probes with a <i>tll</i> or <i>hop limit</i> of 1 and increases by 1 until it either gets to the host, or it hits the maximum <i>max_hop</i>. The default maximum <i>max_hop</i> is 30 hops, but this can be set by the <code>-m</code> option.</p> <p>Three probes are sent at each <i>tll</i> (<i>hop limit</i>) setting, and a line is printed showing the <i>tll</i> (<i>hop limit</i>), the hostname and the address of the gateway, and the <i>rtt</i> (round trip time) of each probe. The number of probes may be specifically set using the <code>-q</code> option. If the probe answers come from different gateways, the hostname and the address of each responding system will be printed. If there is no response within a 5 second timeout interval, a "*" is printed for that probe. The <code>-w</code> option may be used to set the timeout interval. Other possible annotations that may appear after the time are:</p> <ul style="list-style-type: none"> ! the <i>tll</i> (<i>hop limit</i>) value in the received packet is <= 1. !H host unreachable. !X communication administratively prohibited. < !N > ICMP (ICMP6) unreachable code N. <p>The following annotations appear only for IPv4:</p> <ul style="list-style-type: none"> !F fragmentation needed. This should never occur. If this is seen, the associated gateway is broken. !N network unreachable. !P protocol unreachable. !S source route failed. This should never occur. If this is seen, the associated gateway is broken. !T unreachable for the specified <i>tos</i> (type-of-service). !U source host isolated or precedence problem.

The following annotations appear only for IPv6:

- !A host unreachable for a reason other than lack of an entry in the routing table.
- !B packet too big.
- !E destination is not a neighbor.
- !R unrecognized next header.

If almost all the probes result in some kind of unreachable code, then `tracert` gives up and exits.

The destination *host* is not supposed to process the UDP probe packets, so the destination *port* default is set to an unlikely value. However, if some application on the destination is using that value, the value of *port* can be changed with the `-p` option.

The only mandatory parameter is the destination *host* name or IP number. The default probe datagram length is 40 bytes (60 bytes for IPv6), but this may be increased by specifying a packet length (in bytes) after the destination *host* name.

All integer arguments to `tracert` can be specified in either decimal or hexadecimal notation. For example, *packetlen* can be specified either as 256 or 0x100.

OPTIONS

`-A addr_family`

Specify the address family of the target host. *addr_family* can be either `inet` or `inet6`. Address family determines which protocol to use. For an argument of `inet`, IPv4 is used. For `inet6`, IPv6 is used.

By default, if the name of a host is provided, not the literal IP address, and a valid IPv6 address exists in the name service database, `tracert` will use this address. Otherwise, if the name service database contains an IPv4 address, it will try the IPv4 address.

Specify the address family `inet` or `inet6` to override the default behavior. If the argument specified is `inet`, `tracert` will use the IPv4 address associated with the hostname. If none exists, `tracert` will state that the host is unknown and exit. It will not try to determine if an IPv6 address exists in the name service database.

If the specified argument is `inet6`, `tracert` will use the IPv6 address that is associated with the hostname. If none exists, `tracert` will state that the host is unknown and exit.

traceroute(1M)

-a	Probe all of the addresses of a multi-homed destination. The output looks like <code>traceroute</code> has been run once for each IP address of the destination. If this option is used together with <code>-A</code> , <code>traceroute</code> probes only the addresses that are of the specified address family. While probing one of the addresses of the destination, user can skip to the next address by sending a <code>SIGINT</code> , or exit <code>traceroute</code> by sending a <code>SIGQUIT</code> signal. See <code>signal(5)</code>
-c <i>traffic_class</i>	Specify the traffic class of probe packets. The value must be an integer in the range from 0 to 255. Gateways along the path may route the probe packet differently depending upon the value of <i>traffic_class</i> set in the probe packet. This option is valid only on IPv6.
-d	Set the <code>SO_DEBUG</code> socket option.
-F	Set the "don't fragment" bit. This option is valid only on IPv4.
-f <i>first_hop</i>	Set the starting <i>ttl</i> (<i>hop limit</i>) value to <i>first_hop</i> , to override the default value 1. <code>traceroute</code> skips processing for those intermediate gateways which are less than <i>first_hop</i> hops away.
-g <i>gateway</i>	Specify a loose source route <i>gateway</i> . The user can specify more than one <i>gateway</i> by using <code>-g</code> for each gateway. The maximum number of gateways is 8 for IPv4 and 127 for IPv6. Note that some factors such as the link MTU can further limit the number of gateways for IPv6. This option cannot be used with the <code>-r</code> option.
-I	Use ICMP (ICMP6) ECHO instead of UDP datagrams.
-i <i>iface</i>	For IPv4, this option specifies a network interface to obtain the source IP address. This is normally only useful on a multi-homed host. The <code>-s</code> option is also another way to do this. For IPv6, it specifies the network interface on which probe packets are transmitted. The argument can be either an interface index, for example, 1, 2, or an interface name, for example, <code>le0</code> , <code>hme0</code> .
-L <i>flow_label</i>	Specify the flow label of probe packets. The value must be an integer in the range from 0 to 1048575. This option is valid only on IPv6.
-l	Print the value of the <i>ttl</i> (<i>hop limit</i>) field in each packet received.

traceroute(1M)

<code>-m max_hop</code>	Set the maximum <i>ttl</i> (<i>hop limit</i>) used in outgoing probe packets. The default is 30 hops, which is the same default used for TCP connections.
<code>-n</code>	Print hop addresses numerically rather than symbolically and numerically. This saves a nameserver address-to-name lookup for each gateway found on the path.
<code>-P pause_sec</code>	Specify a delay, in seconds, to pause between probe packets. This may be necessary if the final destination does not accept undeliverable packets in bursts. By default, <code>traceroute</code> sends the next probe as soon as it has received a reply. Note that <i>pause_sec</i> is a real number.
<code>-p port</code>	Set the base UDP <i>port</i> number used in probes. The default is 33434. <code>traceroute</code> hopes that nothing is listening on UDP <i>ports</i> $(base + (nhops - 1) * nqueries)$ to $(base + (nhops * nqueries) - 1)$ at the destination host, so that an ICMP (ICMP6) <code>PORT_UNREACHABLE</code> message will be returned to terminate the route tracing. If something is listening on a <i>port</i> in the default range, this option can be used to select an unused <i>port</i> range. <i>nhops</i> is defined as the number of hops between the source and the destination.
<code>-Q max_timeout</code>	Stop probing this hop after <i>max_timeout</i> consecutive timeouts are detected. The default value is 5. Useful in combination with the <code>-q</code> option if you have specified a large <i>nqueries</i> probe count.
<code>-q nqueries</code>	Set the desired number of probe queries. The default is 3.
<code>-r</code>	Bypass the normal routing tables and send directly to a host on an attached network. If the host is not on a directly-attached network, an error is returned. This option can be used to send probes to a local host through an interface that has been dropped by the router daemon. See <code>in.routed(1M)</code> . You cannot use this option if the <code>-g</code> option is used.
<code>-s src_addr</code>	Use the following address, which usually is given as a literal IP address, not a hostname, as the source address in outgoing probe packets. On multi-homed hosts, those with more than one IP address, this option can be used to force the source address to be something other than the IP address <code>traceroute</code> picks by default. If

traceroute(1M)

	<p>the IP address is not one of this machine's interface addresses, an error is returned and nothing is sent. For IPv4, when used together with the <code>-i</code> option, the given IP address should be configured on the specified interface. Otherwise, an error will be returned. In the case of IPv6, the interface name and the source address do not have to match.</p>
	<p><code>-t tos</code> Set the <i>tos</i>(type-of-service) in probe packets to the specified value. The default is zero. The value must be an integer in the range from 0 to 255. Gateways along the path may route the probe packet differently depending upon the <i>tos</i> value set in the probe packet. This option is valid only on IPv4.</p>
	<p><code>-v</code> Verbose output. For each hop, the size and the destination of the response packets is displayed. Also ICMP (ICMP6) packets received other than <code>TIME_EXCEEDED</code> and <code>UNREACHABLE</code> are listed as well.</p>
	<p><code>-w waittime</code> Set the time, in seconds, to wait for a response to a probe. The default is 5 seconds.</p>
	<p><code>-x</code> Prevent <code>traceroute</code> from calculating checksums. Note that checksums are usually required for the last hop when using ICMP ECHO probes. This option is valid only on IPv4. See the <code>-I</code> option.</p>
OPERANDS	<p>The following operands are supported:</p> <p><i>host</i> The network host.</p>
EXAMPLES	<p>EXAMPLE 1 Using the <code>traceroute</code> Utility For a Host Which has Only IPv4 Addresses</p> <p>In the following examples, <code>traceroute</code> is tracking the route to host <code>sanfrancisco</code>, which has only IPv4 addresses in the name service database. Therefore <code>traceroute</code> uses only IPv4 addresses. The following shows the 7-hop path that a packet would follow from the host <code>istanbul</code> to the host <code>sanfrancisco</code>.</p> <pre>istanbul% traceroute sanfrancisco traceroute: Warning: Multiple interfaces found; using 172.31.86.247 @ 1e0 traceroute to sanfrancisco (172.29.64.39), 30 hops max, 40 byte packets 1 frbldg7c-86 (172.31.86.1) 1.516 ms 1.283 ms 1.362 ms 2 bldg1a-001 (172.31.1.211) 2.277 ms 1.773 ms 2.186 ms 3 bldg4-bldg1 (172.30.4.42) 1.978 ms 1.986 ms 13.996 ms 4 bldg6-bldg4 (172.30.4.49) 2.655 ms 3.042 ms 2.344 ms 5 ferbldg11a-001 (172.29.1.236) 2.636 ms 3.432 ms 3.830 ms 6 frbldg12b-153 (172.29.153.72) 3.452 ms 3.146 ms 2.962 ms 7 sanfrancisco (172.29.64.39) 3.430 ms 3.312 ms 3.451 ms</pre>

EXAMPLE 1 Using the traceroute Utility For a Host Which has Only IPv4 Addresses
(Continued)

EXAMPLE 2 Using the traceroute Utility With Source Routing

The following example shows the path of a packet that goes from istanbul to sanfrancisco through the hosts cairo and paris, as specified by the -g option. The -I option makes traceroute send ICMP ECHO probes to the host sanfrancisco. The -i option sets the source address to the IP address configured on the interface qe0.

```
istanbul% traceroute -g cairo -g paris -i qe0 -q 1 -I sanfrancisco
traceroute to sanfrancisco (172.29.64.39), 30 hops max, 56 byte packets
 1  frbldg7c-86 (172.31.86.1)  2.012 ms
 2  flrbldg7u (172.31.17.131)  4.960 ms
 3  cairo (192.168.163.175)  4.894 ms
 4  flrbldg7u (172.31.17.131)  3.475 ms
 5  frbldg7c-017 (172.31.17.83)  4.126 ms
 6  paris (172.31.86.31)  4.086 ms
 7  frbldg7b-82 (172.31.82.1)  6.454 ms
 8  bldg1a-001 (172.31.1.211)  6.541 ms
 9  bldg6-bldg4 (172.30.4.49)  6.518 ms
10  ferbldg11a-001 (172.29.1.236)  9.108 ms
11  frbldg12b-153 (172.29.153.72)  9.634 ms
12  sanfrancisco (172.29.64.39)  14.631 ms
```

EXIT STATUS The following exit values are returned:

0 Successful operation.

>0 An error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

The traceroute utility must be run with the net_rawaccess privilege.

An administrative role must run this utility. By default, traceroute is in the Network Management profile.

Trusted Solaris 8
4/01 Reference Manual

netstat(1M)
ping(1M),attributes(5)

traceroute(1M)

WARNINGS	This utility is intended for use in network testing, measurement and management. It should be used primarily for manual fault isolation. Because of the load it could impose on the network, it is unwise to use <code>traceroute</code> during normal operations or from automated scripts.
-----------------	--

NAME	uadmin – administrative control				
SYNOPSIS	<pre>/usr/sbin/uadmin cmd fcn [mdep] /sbin/uadmin cmd fcn [mdep]</pre>				
DESCRIPTION	<p>The uadmin command provides control for basic administrative functions. This command is tightly coupled to the system administration procedures and is not intended for general use.</p> <p>Both the <i>cmd</i> (command) and <i>fcn</i> (function) arguments are converted to integers and passed to the uadmin system call. The optional <i>mdep</i> (machine dependent) argument is only available for the <i>cmd</i> values of 1 (A_REBOOT) or 2 (A_SHUTDOWN), to pass a single string of boot arguments to the uadmin system call. For any other <i>cmd</i> value, no <i>mdep</i> command-line argument is allowed.</p> <p>When passing an <i>mdep</i> value that contains whitespaces, the string must be grouped together as a single argument enclosed within quotes (for example, uadmin 1 1 "-s kernel/unix").</p>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>The privileges needed for this command to succeed depend on <i>cmd</i> and <i>fcn</i>. For A_REMOUNT, A_SHUTDOWN, A_REBOOT, and A_FREEZE, the necessary privilege is sys_boot. For A_SWAPCTL ADD, and A_SWAPCTL REMOVE, the necessary privilege is sys_mount.</p>				
Trusted Solaris 8 4/01 Reference Manual	<p>uadmin(2)</p> <p>attributes(5)</p>				

h2>umount(1M)

NAME	mount, umount – mount or unmount file systems and remote resources
SYNOPSIS	<pre> mount [-p -v] mount [-F <i>FSType</i>] [<i>generic_options</i>] [-o <i>specific_options</i>] [-O] <i>special</i> <i>mount_point</i> mount [-F <i>FSType</i>] [<i>generic_options</i>] [-o <i>specific_options</i>] [-O] [-S <i>attribute_list</i>] <i>special mount_point</i> mount -a [-F <i>FSType</i>] [-V] [<i>current_options</i>] [-o <i>specific_options</i>] [-S <i>attribute_list</i>] [<i>mount_point...</i>] umount [-f] [-V] [-o <i>specific_options</i>] <i>special</i> <i>mount_point</i> umount -a [-f] [-V] [-o <i>specific_options</i>] [<i>mount_point...</i>] umount -a [-V] [-o <i>specific_options</i>] [<i>mount_point...</i>] </pre>
DESCRIPTION	<p>mount attaches a file system to the file system hierarchy at the <i>mount_point</i>, which is the pathname of a directory. If <i>mount_point</i> has any contents prior to the mount operation, these are hidden until the file system is unmounted.</p> <p>umount unmounts a currently mounted file system, which may be specified either as a <i>mount_point</i> or as <i>special</i>, the device on which the file system resides.</p> <p>The table of currently mounted file systems can be found by examining the mounted file system information file. This is provided by a file system that is usually mounted on <i>/etc/mnttab</i>. The mounted file system information is described in <i>mnttab</i>(4). Mounting a file system adds an entry to the mount table; a <i>umount</i> command removes an entry from the table.</p> <p>When invoked with both the <i>special</i> and <i>mount_point</i> arguments and the <i>-F</i> option, mount validates all arguments except for <i>special</i> and invokes the appropriate <i>FSType</i>-specific mount module. If invoked with no arguments, mount lists all the mounted file systems recorded in the mount table, <i>/etc/mnttab</i>. If invoked with a partial argument list (with only one of <i>special</i> or <i>mount_point</i>, or with both <i>special</i> or <i>mount_point</i> specified but not <i>FSType</i>), mount will search <i>/etc/vfstab</i> for an entry that will supply the missing arguments. If no entry is found, and the <i>special</i> argument starts with <i>"/"</i>, the default local file system type specified in <i>/etc/default/fs</i> will be used. Otherwise the default remote file system type will be used. The default remote file system type is determined by the first entry in the <i>/etc/dfs/fstypes</i> file. After filling in missing arguments, mount will invoke the <i>FSType</i>-specific mount module.</p> <p>The <i>-o</i> and <i>-S</i> options can be used to assign any or all of the following mount-time security attributes to the named file system when appropriate: a sensitivity label, forced privilege(s), allowed privilege(s), a filesystem label range, or an MLD prefix. If <i>-o</i> or <i>-S</i> options are not used, mount also searches <i>/etc/security/tsol/vfstab_adjunct</i> for any security attributes that may be specified there for the file system being mounted.</p>

Mount-time security attributes should be specified for file systems whose objects do not support the Trusted Solaris extended security attributes, such as sensitivity labels. When a required attribute is not specified at mount-time, a default value is applied. The defaults are described in the `OPTIONS` section, where the keywords are defined for the `-S` option.

File system types `UFS`, `TMPFS`, and `NFS` (from a Trusted Solaris server) have a full set of Trusted Solaris extended security attributes already defined. (See the `getfsattr(1M)` man page for how to get attributes on mounted file systems). Because the attributes can be changed on these file systems *after* they are mounted, they are called *variable* file systems. For example, the sensitivity label on a file in a variable file system can be changed by an authorized user. The security attributes on a variable file system can be overridden at mount time, but individual objects in the file system retain any attributes that were originally set on the objects.

File systems that do not support the Trusted Solaris extended security attributes are called *fixed* because any attributes assigned to them (either at mount time or by default) cannot be changed. For example, the sensitivity label specified at mount time for a fixed-attribute file system cannot be changed on any of the objects in that file system. An object that is moved or copied from the fixed file system to a variable file system can be changed after the move.

Mount-time security attributes override existing security attributes on a file system. However, mount-time attributes never override security attributes on the files and directories within the file system.

Without privilege, `mount` can be used to list mounted file systems and resources. To be able to mount and unmount, the `mount` command must have the `sys_mount` privilege. The `umount` command must have the `sys_mount` privilege. Because mounting a `UFS` file system enables/disables logging, it requires the `sys_fs_config` privilege. Mandatory and discretionary read access is required both to the mount point and to the device being mounted; otherwise, MAC or DAC override privileges are required as described in `Intro(2)`. To succeed in all cases with no error side effects, the `mount` command needs: `file_mac_read`, `file_dac_read`, `file_mac_write`, `file_dac_write`, `file_mac_search`, `file_dac_search`, `net_privaddr`, `proc_setsl`, `sys_fs_config`, `sys_mount`, and `sys_trans_label`. To succeed in all cases, `umount` needs: `file_mac_read`, `file_dac_read`, `file_mac_write`, `file_dac_write`, `file_mac_search`, and `file_dac_search`.

When mounting a `UFS` file system, `mount` should assert the `sys_fs_config` privilege. Otherwise, the mount succeeds, but logging is not enabled/disabled, `errno` is set to `EPERM`, and the user sees an error message.

OPTIONS

`-F FSType`

Used to specify the `FSType` on which to operate. The `FSType` must be specified or must be determinable from `/etc/vfstab`, or by consulting `/etc/default/fs` or `/etc/dfs/fstypes`.

umount(1M)

-a [*mount_points*. . .]

Perform mount or umount operations in parallel, when possible.

If mount points are not specified, mount will mount all file systems whose `/etc/vfstab` "mount at boot" field is "yes". If mount points are specified, then `/etc/vfstab` "mount at boot" field will be ignored.

If mount points are specified, umount will only unmount those mount points. If none is specified, then umount will attempt to unmount all filesystems in `/etc/mnttab`, with the exception of certain system required file systems: `/`, `/usr`, `/var`, `/proc`, `/dev/fd`, and `/tmp`.

-f

Forcibly unmount a file system.

Without this option, umount does not allow a file system to be unmounted if a file on the file system is busy. Using this option can cause data loss for open files; programs which access files after the file system has been unmounted will get an error (EIO).

-p

Print the list of mounted file systems in the `/etc/vfstab` format. Must be the only option specified.

-v

Print the list of mounted file systems in verbose format. Must be the only option specified.

-V

Echo the complete command line, but do not execute the command. umount generates a command line by using the options and arguments provided by the user and adding to them information derived from `/etc/mnttab`. This option should be used to verify and validate the command line.

generic_options

Options that are commonly supported by most *FSType*-specific command modules. The following options are available:

-m

Mount the file system without making an entry in `/etc/mnttab`.

-g

Globally mount the file system. On a clustered system, this globally mounts the file system on all nodes of the cluster. On a non-clustered system this has no effect.

-o

Specify *FSType*-specific options in a comma separated (without spaces) list of suboptions and keyword-attribute pairs for interpretation by the *FSType*-specific module of the command. (See `mount_ufs(1M)`)

- O
Overlay mount. Allow the file system to be mounted over an existing mount point, making the underlying file system inaccessible. If a mount is attempted on a pre-existing mount point without setting this flag, the mount will fail, producing the error “device busy”.
- r
Mount the file system read-only.
- S *attribute_list*
Specify in *attribute_list* a quoted semicolon-separated list of security attributes to associate with the filesystem mount. Each attribute is specified with a value assigned to a keyword in semicolon-separated fields. All keywords are optional and follow the format:
- keyword=value where *keyword* is one of the following:
- | | |
|------------|---|
| slabel | Sets the sensitivity label for all objects in the file system. Specify the sensitivity label in hexadecimal or text format. |
| forced | Specify one or more forced privileges for all executable files in the file system. Specify symbolic privilege name(s) in a comma-separated list (such as: forced=file_audit, file_chown;) or use all to indicate all privileges. Using none or omitting the keyword results in no forced privileges being applied. See <code>priv_desc(4)</code> . Any forced privileges must be a subset of the allowed privileges. |
| allowed | Specify one or more allowed privilege(s) for all executable files in the file system. Specify symbolic privilege names in a comma-separated list (such as: allowed=file_audit, file_chown;) or use all to indicate all privileges. Using none or omitting the keyword results in no allowed privileges being applied. See <code>priv_desc(4)</code> for names of privileges. Any allowed privilege(s) must be a superset of the forced privileges. |
| low_range | Specify the lower bound of the file system label range as a sensitivity label in text format. |
| hi_range | Specify the upper bound of the file system label range as a sensitivity label in text format. |
| mld_prefix | Set a prefix to be used in the adorned names of multilevel directories. (See multilevel directories in the DEFINITIONS in <code>Intro(2)</code> for more about the MLD prefix.) Specify the value in text format (such as: .MLD. or .hidden.). On unlabeled (fixed attribute) file systems, the prefix generally has no useful effect—with the exception that an <code>mld_prefix</code> should be supplied if a variable filesystem is being mounted on the unlabeled filesystem and the root of the variable filesystem |

umount(1M)

is an MLD.

Any of the above keywords may be omitted.

Note – The semicolon separators between keyword/value pairs and any brackets used to specify sensitivity labels must be commented out so that the separators and brackets can be interpreted properly by the shell.

When a keyword appears without an attribute value or when a keyword is missing, a default value is assigned to that attribute. The default values for fixed attribute file systems are:

slabel	The default sensitivity label of a fixed file system being mounted from a local device (such as a hard disk, floppy, or CD-ROM) is the sensitivity label of the device. For an allocated device, the file system is assigned the sensitivity label at which the device was allocated.
forced	None
allowed	None
low_range	ADMIN_LOW
hi_range	ADMIN_HIGH
mld_prefix	None

For example, the assignment of `forced=`; results in the default of "none" being applied.

Note – Most of the keyword=value pairs used to specify security attributes with the `-S` option can be entered directly under the `-o` option—with one caveat. Since mount options are comma-separated, any security attribute specified with a keyword followed by multiple values separated by commas is not allowed after `-o`. See Example 2.

USAGE See `largefile(5)` for the description of the behavior of `mount` and `umount` when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).

EXAMPLES **EXAMPLE 1** Assigning Security Attributes with `-o`

In this example, the `-o` is used to assign security attributes.

```
% mount -F tmpfs -o allowed=all,slabel=c swap /mnt
```

EXAMPLE 2 Assigning Security Attributes with `-S`

Trusted Solaris security attributes that are separated with commas cannot be passed to the `-o` option. Therefore, use the `-S` option.

```
% mount -F tmpfs -S "allowed=all;forced=proc_tranquil,proc_dumpcore" \\  
swap /mnt
```

EXAMPLE 2 Assigning Security Attributes with -S (Continued)

These security attributes cannot be entered with the -o option since the comma separator in the privileges list would be interpreted as the start of a new option.

FILES

/etc/mnttab	Mount table
/etc/default/fs	Default local file system type. Default values can be set for the following flags in /etc/default/fs. For example: LOCAL=ufs
	LOCAL: The default partition for a command if no <i>FSType</i> is specified.
/etc/vfstab	List of default parameters for each file system.
/etc/security/tsol/vfstab_adjunct	Mount-time attributes for file systems.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

Trusted Solaris security policy applies when mounting and unmounting file systems.

Mount-time security attributes may be specified by using mount with the -o or -S option on the command line or by specifying the attributes in the vfstab_adjunct file. Mount-time security attributes override existing security attributes on a file system. However, they never override security attributes on the files and directories within the file system. When access-control decisions are made, security attributes on a file or directory take precedence over security attributes specified either at the filesystem level or at mount time.

Except when merely listing mounted file systems and resources, mount must run with the sys_mount privilege. umount also must run with the sys_mount privilege. To succeed in all cases, mount needs: file_mac_read, file_dac_read, file_mac_write, file_dac_write, file_mac_search, file_dac_search, net_privaddr, proc_setsl, sys_mount, and sys_trans_label.

When mounting a UFS file system, mount should assert the sys_fs_config privilege. Otherwise, the mount succeeds, but logging is not enabled/disabled, errno is set to EPERM, and the user sees an error message.

Trusted Solaris 8 4/01 Reference Manual

getfsattr(1M), getmldadorn(1), mount_hsf(1M), mount_nfs(1M), mount_pcfs(1M), mount_tmpfs(1M), mount_ufs(1M), mountall(1M), setfsattr(1M), mnttab(4), priv_desc(4), vfstab(4), vfstab_adjunct(4)

umount(1M)

Trusted Solaris Administrator's Procedures

**SunOS 5.8
Reference Manual**

mount_cachefs(1M), default_fs(4), attributes(5), largefile(5), lofs(7FS),
pcfs(7FS)

NOTES

If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on the directory to which the symbolic link refers, rather than on top of the symbolic link itself.

NAME	mountall, umountall – Mount, unmount multiple file systems
SYNOPSIS	mountall [-F <i>FSType</i>] [-l -r] [<i>file_system_table</i>] umountall [-k] [-s] [-F <i>FSType</i>] [-l -r] umountall [-k] [-s] [-h <i>host</i>]
DESCRIPTION	<p>mountall is used to mount file systems specified in a file system table. The file system table must be in <code>vfstab(4)</code> format. If no <i>file_system_table</i> is specified, <code>/etc/vfstab</code> will be used. If <code>'-'</code> is specified as <i>file_system_table</i>, mountall will read the file system table from the standard input. mountall only mounts those file systems with the mount at boot field set to <code>yes</code> in the <i>file_system_table</i>.</p> <p>Each file system which has an <code>fsckdev</code> entry specified in the file system table will be checked using <code>fsck(1M)</code> in order to determine if it may be safely mounted. If the file system does not appear mountable, it is fixed using <code>fsck</code> before the mount is attempted. File systems with a <code>'-'</code> entry in the <code>fsckdev</code> field will be mounted without first being checked.</p> <p>umountall causes all mounted file systems except <code>root</code>, <code>/usr</code>, <code>/var</code>, <code>/var/adm</code>, <code>/var/run</code>, <code>/proc</code>, and <code>/dev/fd</code> to be unmounted. If the <i>FSType</i> is specified, mountall and umountall limit their actions to the <i>FSType</i> specified. There is no guarantee that umountall will unmount <i>busy</i> filesystems, even if the <code>-k</code> option is specified.</p> <p>mountall and umountall must run with the <code>sys_mount</code> privilege.</p> <p>Mandatory and discretionary read access are required to both the mount point and the device being mounted; to override MAC and DAC restrictions requires privilege as described in <code>Intro(2)</code>. To succeed in all cases, the mountall and umountall commands need the privileges: <code>file_mac_read</code>, <code>file_dac_read</code>, <code>file_mac_write</code>, <code>file_dac_write</code>, <code>file_mac_search</code>, <code>file_dac_search</code>, <code>net_privaddr</code>, <code>proc_setsl</code>, <code>sys_mount</code>, and <code>sys_trans_label</code>.</p> <p>When mounting a UFS file system, mount should assert the <code>sys_fs_config</code> privilege. Otherwise, the mount succeeds, but logging is not enabled/disabled, <code>errno</code> is set to <code>EPERM</code>, and the user sees an error message.</p>
OPTIONS	<p><code>-F</code> Specify the <i>FSType</i> of the file system to be mounted or unmounted.</p> <p><code>-h host</code> Unmount all file systems listed in <code>/etc/mnttab</code> that are remote-mounted from host.</p> <p><code>-k</code> Use the <code>fuser -k mount-point</code> command. See the <code>fuser(1M)</code> for details. The <code>-k</code> option sends the <code>SIGKILL</code> signal to each process using the file. As this option spawns kills for each process, the kill messages may not show up immediately. There is no guarantee that umountall will unmount <i>busy</i> filesystems, even if the <code>-k</code> option is specified.</p> <p><code>-l</code> Limit the action to local file systems.</p>

umountall(1M)

- r Limit the action to remote file system types.
- s Do not perform the umount operation in parallel.

FILES /etc/mnttab mounted file system table
/etc/vfstab table of file system defaults

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

SUMMARY OF TRUSTED SOLARIS CHANGES

Trusted Solaris security policy applies when mounting and unmounting file systems. mountall and umountall must run with the sys_mount privilege.

Mandatory and discretionary read access are required to both the mount point and the device being mounted; to override MAC and DAC restrictions requires privilege as described in Intro(2). To succeed in all cases, the mountall and umountall commands need the privileges: file_mac_read, file_dac_read, file_mac_write, file_dac_write, file_mac_search, file_dac_search, net_privaddr, proc_setsl, sys_mount, and sys_trans_label.

When mounting a UFS file system, mount should assert the sys_fs_config privilege. Otherwise, the mount succeeds, but logging is not enabled/disabled, errno is set to EPERM, and the user sees an error message.

Mount-time security attributes may be specified in the vfstab_adjunct file.

Trusted Solaris 8 4/01 Reference Manual
Mount
Reference Manual
DIAGNOSTICS

mount(1M), mnttab(4), vfstab(4), vfstab_adjunct(4)
fsck(1M), fuser(1M), attributes(5)

No messages are printed if the file systems are mountable and clean.

Error and warning messages come from fsck(1M) and mount(1M).

NAME	unshare – Make local resource unavailable for mounting by remote systems				
SYNOPSIS	unshare [-F <i>FSType</i>] [-o <i>specific_options</i>] [<i>pathname</i> <i>resourcename</i>]				
DESCRIPTION	The unshare command makes a shared local resource unavailable as file system type <i>FSType</i> . If the option -F <i>FSType</i> is omitted, then the first file system type listed in file <i>/etc/dfs/fstypes</i> will be used as the default. <i>Specific_options</i> , as well as the semantics of <i>resourcename</i> , are specific to particular distributed file systems.				
OPTIONS	<div> <div>-F <i>FSType</i></div> <div>Specify the file system type.</div> </div> <div> <div>-o <i>specific_options</i></div> <div>Specify options specific to the file system provided by the -F option.</div> </div>				
SUMMARY OF TRUSTED SOLARIS CHANGES	The unshare command must be run with an effective UID of 0. If the file being unshared is of the type NFS, this command must have the <code>sys_nfs</code> privilege to succeed. If this command has the <code>file_mac_write</code> privilege, it can be run any sensitivity label other than <code>ADMIN_LOW</code> . To succeed in all cases, this command needs the <code>file_mac_read</code> and <code>file_mac_search</code> privileges.				
FILES	<div> <div><i>/etc/dfs/fstypes</i></div> <div>List of system types, NFS by default.</div> </div> <div> <div><i>/etc/dfs/sharetab</i></div> <div>System record of shared file systems.</div> </div>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
Trusted Solaris 8 4/01 Reference Manual	<code>share(1M)</code> , <code>shareall(1M)</code>				
UNIX 95 Reference Manual	<code>attributes(5)</code>				
NOTES	If <i>pathname</i> or <i>resourcename</i> is not found in the shared information, an error message will be sent to standard error.				

unshareall(1M)

NAME	shareall, unshareall – Share, unshare multiple resources				
SYNOPSIS	shareall [-F <i>FSType</i> [, <i>FSType</i> ...]] [- <i>file</i>] unshareall [-F <i>FSType</i> [, <i>FSType</i> ...]]				
DESCRIPTION	<p>When used with no arguments, shareall shares all resources from <i>file</i>, which contains a list of share command lines. If the operand is a hyphen (-), then the share command lines are obtained from the standard input. Otherwise, if neither a <i>file</i> nor a hyphen is specified, then the file <i>/etc/dfs/dfstab</i> is used as the default.</p> <p>Resources may be shared by specific file system types by specifying the file systems in a comma-separated list as an argument to -F.</p> <p>unshareall unshares all currently shared resources. Without a -F flag, it unshares resources for all distributed file system types.</p>				
OPTIONS	-F <i>FSType</i> Specify file system type. Defaults to the first entry in <i>/etc/dfs/fstypes</i> .				
FILES	<i>/etc/dfs/dfstab</i> List of share commands to be executed at boot time.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes: <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> </thead> <tbody> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>The shareall and unshareall commands must be run with an effective UID of 0. If any file being shared or unshared is of the type NFS, then the command requires the sys_nfs privilege [see share_nfs(1M)]. If the file <i>/etc/dfs/sharetab</i> does not exist, the shareall command will create the file; thus, the shareall command must be run at the sensitivity level of ADMIN_LOW. If the file <i>/etc/dfs/sharetab</i> exists, then the shareall and unshareall commands can be run at any other sensitivity level if they have the file_mac_write privilege. To succeed in all cases, the commands need the file_mac_read and file_mac_search privileges.</p>				
Trusted Solaris 8 4/01 Reference Manual	share(1M), unshare(1M) attributes(5)				

NAME	unshare_nfs – Make local NFS file systems unavailable for mounting by remote systems				
SYNOPSIS	unshare [-F nfs] <i>pathname</i>				
DESCRIPTION	The unshare command makes local file systems unavailable for mounting by remote systems. The shared file system must correspond to a line with NFS as the <i>FSType</i> in the file <i>/etc/dfs/sharetab</i> .				
OPTIONS	<p>The following options are supported:</p> <p>-F This option may be omitted if NFS is the first file system type listed in the file <i>/etc/dfs/fstypes</i>.</p>				
FILES	<p><i>/etc/dfs/fstypes</i> List of system types, NFS by default.</p> <p><i>/etc/dfs/sharetab</i> System record of shared file systems.</p>				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWcsu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsu				
SUMMARY OF TRUSTED SOLARIS CHANGES Trusted Solaris 8 4/01 Reference Manual SunOS 5.6 Reference Manual	<p>The <i>sys_nfs</i> privilege is required to run this command, which must be run as UID 0 at label [ADMIN_LOW].</p> <p>share(1M)</p> <p>attributes(5)</p> <p>If the file system being unshared is a symbolic link to a valid pathname, the canonical path (the path which the symbolic link follows) will be unshared.</p> <p>For example, if <i>/export/foo</i> is a symbolic link to <i>/export/bar</i> (<i>/export/foo</i> -> <i>/export/bar</i>), the following unshare command will result in <i>/export/bar</i> as the unshared pathname (and not <i>/export/foo</i>).</p> <p>example# unshare -F nfs /export/foo</p>				
NOTES					

updatehome(1M)

NAME	updatehome – Update the home directory copy and link files for the current label				
SYNOPSIS	updatehome [-cirs]				
DESCRIPTION	<p>updatehome reads the user's minimum-label copy and link-control files (<code>.copy_files</code> and <code>.link_files</code>), which contain a list of files to be copied and symbolically linked from the user's minimum-label home directory to the user's home directory at the current label.</p> <p>The Trusted Solaris <code>dtsession</code> performs an <code>updatehome</code> whenever a newly labeled workspace is created so that the user's favorite files are available for use. For example, the user probably wants a symlink to such files as <code>.profile</code>, <code>.login</code>, <code>.cshrc</code>, <code>.exrc</code>, <code>.mailrc</code>, and <code>~/bin</code>. <code>updatehome</code> provides a convenient mechanism for accomplishing this symlink. The user may add files to those to be copied (<code>.copy_files</code>) and to those to be symbolically linked (<code>.link_files</code>).</p>				
ATTRIBUTES	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p> <table><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWtsu</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWtsu				
OPTIONS	<p>-c Replace existing home-directory copies at the current label. (The default is to skip over existing copies.)</p> <p>-i Ignore errors encountered. (The default aborts on error.)</p> <p>-r Replace existing home-directory copies or symbolic links at the current label. This option implies options -c and -s. (The default is to skip over existing copies or symbolic links.)</p> <p>-s Replace existing home-directory symbolic links at the current label. (The default is to skip over existing symbolic links.)</p>				
RETURN VALUES	Upon success, <code>updatehome</code> returns 0. Upon failure, <code>updatehome</code> returns 1 and writes diagnostic messages to standard error.				
EXAMPLES	<p>EXAMPLE 1 A Sample <code>copy_files</code></p> <p>The files listed in <code>.copy_files</code> can be modified at every user's label.</p> <pre>.cshrc .mailrc .netscape/bookmarks.html</pre> <p>EXAMPLE 2 A Sample <code>link_files</code></p> <p>The files listed in <code>.link_files</code> can be modified at the lowest label, and the changes will propagate to the other labels available to the user.</p>				

EXAMPLE 2 A Sample link_files *(Continued)*

```
~/bin
.netscape/preferences
.xrc
.rhosts
```

FILES	\$HOME/.copy_files	List of files to be copied
	\$HOME/.link_files	List of files to be symbolically linked
	attributes(5)	

writeaudit(1M)

NAME	writeaudit – write an audit record																				
SYNOPSIS	writeaudit <i>event</i> [-a <i>type:value</i> ...] [-f <i>type:filename</i> ...]																				
DESCRIPTION	For a specified event, this command writes an audit record containing zero or more attributes. If no AW_RETURN attribute is specified, a successful return attribute (0,0) will be included in the audit record. Multiple -a or -f options can be specified on a single writeaudit call.																				
FIELDS	<i>event</i> The name of the event to record in the audit record. This option must always be present. The name must be defined in audit_event file. See audit_event(4).																				
OPTIONS	<p>-a <i>type:value</i> Add an attribute to the audit record. The <i>type</i> must be AW_DATA, AW_INADDR, AW_OPAQUE, AW_PATH, AW_RETURN, AW_SLABEL, or AW_TEXT. Valid formats for <i>value</i> are described below.</p> <p>-f <i>type:filename</i> Add an attribute to the audit record. The <i>type</i> must be AW_DATA, AW_INADDR, AW_OPAQUE, AW_PATH, AW_RETURN, AW_SLABEL, or AW_TEXT. The <i>value</i> is read from the file <i>filename</i>. Valid formats for <i>value</i> are described below.</p>																				
AW_DATA Format	<p>AW_DATA: <i>printformat</i> :<i>items</i><i>size</i> :<i>number</i><i>items</i> :<i>item</i>1: . . . <i>item</i>N</p> <p>The <i>printformat</i> field must be one of these:</p> <table> <tr><td>AWD_BINARY</td><td>Print data in binary</td></tr> <tr><td>AWD_OCTAL</td><td>Print data in octal</td></tr> <tr><td>AWD_DECIMAL</td><td>Print data in decimal</td></tr> <tr><td>AWD_HEX</td><td>Print data in hex</td></tr> <tr><td>AWD_STRING</td><td>Print data as a string</td></tr> </table> <p>The <i>items</i><i>size</i> field must be one of these:</p> <table> <tr><td>AWD_BYTE</td><td>Data is in units of bytes</td></tr> <tr><td>AWD_CHAR</td><td>Data is in units of chars (1 byte)</td></tr> <tr><td>AWD_SHORT</td><td>Data is in units of shorts (2 bytes)</td></tr> <tr><td>AWD_INT</td><td>Data is in units of ints (4 bytes)</td></tr> <tr><td>AWD_LONG</td><td>Data is in units of longs (4 bytes)</td></tr> </table> <p><i>number</i><i>items</i> specifies the number of items to be printed and must be an integer in the range 1-255.</p> <p><i>item</i>1 through <i>item</i>N specify the data fields to be printed and must be entered in hex (for example, 0xffff), octal (for example, 0777), or decimal.</p>	AWD_BINARY	Print data in binary	AWD_OCTAL	Print data in octal	AWD_DECIMAL	Print data in decimal	AWD_HEX	Print data in hex	AWD_STRING	Print data as a string	AWD_BYTE	Data is in units of bytes	AWD_CHAR	Data is in units of chars (1 byte)	AWD_SHORT	Data is in units of shorts (2 bytes)	AWD_INT	Data is in units of ints (4 bytes)	AWD_LONG	Data is in units of longs (4 bytes)
AWD_BINARY	Print data in binary																				
AWD_OCTAL	Print data in octal																				
AWD_DECIMAL	Print data in decimal																				
AWD_HEX	Print data in hex																				
AWD_STRING	Print data as a string																				
AWD_BYTE	Data is in units of bytes																				
AWD_CHAR	Data is in units of chars (1 byte)																				
AWD_SHORT	Data is in units of shorts (2 bytes)																				
AWD_INT	Data is in units of ints (4 bytes)																				
AWD_LONG	Data is in units of longs (4 bytes)																				
AW_INADDR Format	AW_INADDR: <i>hostname</i>																				

writeaudit(1M)

	<p><i>hostname</i> must be a valid hostname (for example, hamlet), or a standard IP address (for example, 129.150.117.44).</p>				
AW_OPAQUE Format	<p><i>AW_OPAQUE: numberitems : item1: . . . itemN</i></p> <p><i>numberitems</i> specifies the number of items to be printed and must be an integer in the range 1-255.</p> <p><i>item1</i> through <i>itemN</i> specify the fields to be printed and must be input in hex (for example, 0xfff), octal (for example, 0777), or decimal. Each field must not exceed 1 byte in length.</p>				
AW_PATH Format	<p><i>AW_PATH: path</i></p> <p><i>path</i> is a text string (for example, /usr/bin/).</p>				
AW_RETURN Format	<p><i>AW_RETURN: status_value : return_value</i></p> <p><i>status_value</i> identifies the error status of the call and must be an integer in the range 0-255.</p> <p><i>return_value</i> identifies the call return value and must be an integer in the range 0-255.</p>				
AW_SLABEL Format	<p><i>AW_SLABEL: sensitivity_label</i></p> <p><i>sensitivity_label</i> must be a valid character-coded sensitivity label; for example, S AB or 0x7fffffffffffffffffffffffffffffffff\ ffffffffffffffffffffffffffffffffff</p>				
AW_TEXT Format	<p><i>AW_TEXT: string</i></p> <p><i>string</i> must be a text string; for example, successful change.</p>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>SUNWtsu</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWtsu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWtsu				
EXAMPLES	<p>EXAMPLE 1 Write Event Records</p> <p>For the event, write an AUE_event record containing the string successful change:</p> <pre>writeaudit AUE_event -a AW_TEXT: "successful change"</pre> <p>For the event, read the text string from the file eventfile and write an AUE_event record (the file eventfile might, for example, contain the string successful change):</p>				

writeaudit(1M)

Trusted Solaris 8
4/01 Reference
Manual
NOTES

EXAMPLE 1 Write Event Records (Continued)

```
writeaudit AUE_event -f \ AW_TEXT:eventfile -a AW_RETURN:-1:4
```

For the event, write an AUE_event record containing the specified arbitrary data:

```
writeaudit AUE_event -a\  
AW_DATA:AWD_DECIMAL:AWD_BYTE:5:1:2:3:4:5
```

```
audit(2), auditwrite(3TSOL), audit_event(4)
```

```
attributes(5)
```

This command must have the `proc_audit_appl` privilege in its set of effective privileges. To translate labels (for example, *type* `AW_SLABEL`) that dominate the process's sensitivity label, this command must have the `priv_sys_trans_label` privilege in its set of effective privileges.

These interfaces are uncommitted. Although they are not expected to change between minor releases of the Trusted Solaris environment, they may.

NAME	ypbind – NIS binder process
SYNOPSIS	<code>/usr/lib/netsvc/yp/ypbind</code> [-broadcast -u -ypset -ypsetme]
DESCRIPTION	<p>NIS provides a simple network lookup service consisting of databases and processes. The databases are stored at the machine that runs an NIS server process. The programmatic interface to NIS is described in <code>ypclnt(3NSL)</code>. Administrative tools are described in <code>ypinit(1)</code>, <code>ypwhich(1)</code>, and <code>ypset(1)</code>. Tools to see the contents of NIS maps are described in <code>ypcat(1)</code>, and <code>ypmatch(1)</code>.</p> <p><code>ypbind</code> is a daemon process that is activated at system startup time from the startup script <code>/etc/init.d/rpc</code>. By default, it is invoked as <code>ypbind -broadcast</code>. <code>ypbind</code> runs on all client machines that are set up to use NIS. See <code>sysidtool(1)</code>. The function of <code>ypbind</code> is to remember information that lets all NIS client processes on a node communicate with some NIS server process. <code>ypbind</code> must run on every machine which has NIS client processes. The NIS server may or may not be running on the same node, but must be running somewhere on the network. If the NIS server is a NIS+ server in NIS (YP) compatibility mode, see the NOTES section of the <code>ypfiles(4)</code> man page for more information. <code>ypbind</code> must be run from the Trusted Path and inherit the <code>file_dac_write</code>, <code>net_broadcast</code>, <code>net_downgrade_sl</code>, <code>net_upgrade_sl</code>, <code>net_mac_read</code>, <code>net_privaddr</code>, <code>proc_setclr</code>, and <code>proc_setsl</code> privileges.</p> <p>The information <code>ypbind</code> remembers is called a <i>binding</i> — the association of a domain name with a NIS server. The process of binding is driven by client requests. As a request for an unbound domain comes in, if started with the <code>-broadcast</code> option, the <code>ypbind</code> process broadcasts on the net trying to find an NIS server, either a <code>ypserv</code> process serving the domain or an <code>rpc.nisd</code> process in "YP-compatibility mode" serving NIS+ directory with name the same as (case sensitive) the domain in the client request. Since the binding is established by broadcasting, there must be at least one NIS server on the net. If started without the <code>-broadcast</code> option, <code>ypbind</code> process steps through the list of NIS servers that was created by <code>ypinit -c</code> for the requested domain. There must be an NIS server process on at least one of the hosts in the NIS servers file. All the hosts in the NIS servers file must be listed in either the <code>/etc/hosts</code> or <code>/etc/inet/ipnodes</code> files along with their IP addresses. Once a domain is bound by <code>ypbind</code>, that same binding is given to every client process on the node. The <code>ypbind</code> process on the local node or a remote node may be queried for the binding of a particular domain by using the <code>ypwhich(1)</code> command.</p> <p>If <code>ypbind</code> is unable to speak to the NIS server process it is bound to, it marks the domain as unbound, tells the client process that the domain is unbound, and tries to bind the domain once again. Requests received for an unbound domain will wait until the requested domain is bound. In general, a bound domain is marked as unbound when the node running the NIS server crashes or gets overloaded. In such a case, <code>ypbind</code> will try to bind to another NIS server using the process described above. <code>ypbind</code> also accepts requests to set its binding for a particular domain. The request is usually generated by the <code>ypset(1)</code> command. In order for <code>ypset</code> to work, <code>ypbind</code> must have been invoked with flags <code>-ypset</code> or <code>-ypsetme</code>.</p>

ypbind(1M)

OPTIONS**-broadcast**

Send a broadcast datagram using UDP/IP that requests the information needed to bind to a specific NIS server. This option is analogous to ypbind with no options in earlier Sun releases and is recommended for ease of use. Use of this option is strongly discouraged for security reasons.

-u

Allow non-Trusted Solaris TCB machines to bind to ypserv(1M).

-ypset

Allow users from any remote machine to change the binding by means of the ypset command. By default, no one can change the binding. Use of this option is strongly discouraged for security reasons.

-ypsetme

Only allow a user with the NAMESERVICE_CONFIG_AUTH authorization on the local machine to change the binding to a desired server by means of the ypset command. ypbind can verify the caller indeed has the NAMESERVICE_CONFIG_AUTH authorization by accepting such requests only on the loopback transport. By default, no external process can change the binding.

SUMMARY OF TRUSTED SOLARIS CHANGES

ypbind must be run from the Trusted Path and inherit the file_dac_write, net_broadcast, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl privileges.

The -ypsetme option only allows a user with the NAMESERVICE_CONFIG_AUTH authorization on the local machine to change the binding to a desired server by means of the ypset command.

FILES

/var/yp/binding/ypdomain/ypservers

/etc/inet/hosts

/etc/inet/ipnodes

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWnisu

**Trusted Solaris 8
4/01 Reference
Manual**

ifconfig(1M), rpc.nisd(1M), ypserv(1M)

ypcat(1), ypmatch(1), ypwhich(1), ypinit(1M), ypset(1M), ypclnt(3NSL), hosts(4), ipnodes(4), ypfiles(4), attributes(5)

NOTES ypbind supports multiple domains. The ypbind process can maintain bindings to several domains and their servers, the default domain is the one specified by the domainname(1M) command at startup time.

The -broadcast option works only on the UDP transport. It is insecure since it trusts "any" machine on the net that responds to the broadcast request and poses itself as an NIS server.

yppasswdd(1M)

NAME	rpc.yppasswdd, yppasswdd – Server for modifying NIS password file
SYNOPSIS	<pre>/usr/lib/netsvc/yp/rpc.yppasswdd [-D <i>directory</i>] [-nogecos] [-noshell] [-nopw] [-m <i>argument1 argument2...</i>] [-u] /usr/lib/netsvc/yp/rpc.yppasswdd [<i>passwordfile</i> [<i>adjunctfile</i>]] [-nogecos] [-noshell] [-nopw] [-m <i>argument1 argument2...</i>] [-u]</pre>
DESCRIPTION	<p>rpc.yppasswdd is a server that handles password change requests from yppasswd(1). It changes a password entry in the passwd, shadow, and security/passwd.adjunct files. The passwd and shadow files provide the basis for the passwd.byname and passwd.byuid maps. The passwd.adjunct file provides the basis for the passwd.adjunct.byname and passwd.adjunct.byuid maps. Entries in the passwd, shadow or passwd.adjunct files are changed only if the password presented by yppasswd(1) matches the encrypted password of the entry. All password files are located in the PWDIR directory. rpc.yppasswdd must be run from the Trusted Path and inherit the net_broadcast, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl privileges.</p> <p>If the -D option is given, the passwd, shadow, or passwd.adjunct files are placed under the directory path that is the argument to -D.</p> <p>If the -noshell, -nogecos or -nopw options are given, these fields cannot be changed remotely using chfn, chsh, or passwd(1).</p> <p>If the -m option is given, a make(1) is performed in /var/yp after any of the passwd, shadow, or passwd.adjunct files are modified. All arguments following the flag are passed to make.</p> <p>If the -u option is given, updates from non-Trusted Solaris TCB clients are allowed.</p> <p>The second of the listed syntaxes is provided only for backward compatibility. If the second syntax is used, the <i>passwordfile</i> is the full pathname of the password file and <i>adjunctfile</i> is the full pathname of the optional passwd.adjunct file. If a shadow file is found in the same directory as <i>passwordfile</i>, the <i>shadowfile</i> is used as described above. Use of this syntax and the discovery of a <i>shadowfile</i> file generates diagnostic output. The daemon, however, starts normally.</p> <p>The first and second syntaxes are mutually exclusive. You cannot specify the full pathname of the passwd, passwd.adjunct files and use the -D option at the same time.</p> <p>The daemon is started automatically on the master server of the passwd map by ypstart(1), which is invoked at boot time by the /etc/init.d/rpc script.</p> <p>The server does not insist on the presence of a shadow file unless there is no -D option present or the directory named with the -D option is /etc. In addition, a passwd.adjunct file is not necessary. If the -D option is given, the server attempts to find a passwd.adjunct file in the security subdirectory of the named directory.</p>

For example, in the presence of “-D /var/yp” the server checks for a “/var/yp/security/passwd.adjunct” file.

If only a passwd file exists, then the encrypted password is expected in the second field. If both a passwd and a passwd.adjunct file exist, the encrypted password is expected in the second field of the adjunct file with ##username in the second field of the passwd file. If all three files are in use, the encrypted password is expected in the shadow file. Any deviation causes a password update to fail.

If you remove or add a shadow or passwd.adjunct file after rpc.yppasswdd has started, you must stop and restart the daemon to enable it to recognize the change. See ypstart(1) for information on restarting the daemon.

SUMMARY OF TRUSTED SOLARIS CHANGES

rpc.yppasswdd must be run from the Trusted Path and inherit the net_broadcast, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl privileges.

If the -u option is given, updates from non-Trusted Solaris TCB clients are allowed.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWypu

Trusted Solaris 8 4/01 Reference Manual NOTES

passwd(1), inetd(1M), shadow(4)

make(1), yppasswd(1), ypmake(1), passwd(4), ypfiles(4), attributes(5)

If make has not been installed and the -m option is given, the daemon outputs a warning and proceeds, effectively ignoring the -m flag.

When using the -D option, you should make sure that the PWDIR of the /var/yp/Makefile is set accordingly.

The second listed syntax is supplied only for backward compatibility and might be removed in a future release of this daemon.

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications PLC, and cannot be used without permission.

ypserv(1M)

NAME	ypserv, ypxfrd – NIS server and binder processes
SYNOPSIS	/usr/lib/netsvc/yp/ypserv [-duv] /usr/lib/netsvc/yp/ypxfrd
DESCRIPTION	<p>The Network Information Service (NIS) provides a simple network lookup service consisting of databases and processes. The databases are ndbm files in a directory tree rooted at /var/yp. See dbm_clearerr(3C). These files are described in ypfiles(4). The processes are /usr/lib/netsvc/yp/ypserv, the NIS database lookup server, and /usr/lib/netsvc/yp/ypbind, the NIS binder. The programmatic interface to the NIS service is described in ypclnt(3NSL). Administrative tools are described in yppoll(1M), yppush(1M), ypset(1M), ypxfr(1M), and ypwhich(1). Tools to see the contents of NIS maps are described in ypcat(1), and ypmatch(1). Database generation and maintenance tools are described in ypinit(1M), ypmake(1M), and makedbm(1M).</p> <p>The ypserv utility is a daemon process typically activated at system startup from /etc/init.d/rpc. Alternatively, it can be started on the command line using ypstart(1M) from the Trusted Path with these privileges: file_dac_read, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl. ypserv runs only on NIS server machines with a complete NIS database. You can halt all NIS services using the ypstop(1M) command run from the Trusted Path with the proc_owner privilege.</p> <p>The ypxfrd utility transfers entire NIS maps in an efficient manner. For systems that use this daemon, map transfers are 10 to 100 times faster, depending on the map. To use this daemon, be sure ypxfrd is running on the master server. See /usr/lib/netsvc/yp/ypstart. ypxfr attempts to use ypxfrd first. If that fails, it prints a warning, then uses the older transfer method. ypxfrd must be run from the Trusted Path and inherit the net_broadcast, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl privileges.</p> <p>The ypserv daemon's primary function is to look up information in its local database of NIS maps.</p> <p>The operations performed by ypserv are defined for the implementor by the <i>YP Protocol Specification</i>, and for the programmer by the header file <rpcsvc/yp_prot.h>.</p> <p>Communication to and from ypserv is by means of RPC calls. Lookup functions are described in ypclnt(3NSL), and are supplied as C-callable functions in the libns1(3LIB) library. There are four lookup functions, all of which are performed on a specified map within some NIS domain: yp_match(3NSL), yp_first(3NSL), yp_next(3NSL), and yp_all(3NSL). The yp_match operation takes a key, and returns the associated value. The yp_first operation returns the first key-value pair from the map, and yp_next can be used to enumerate the remainder. yp_all ships the entire map to the requester as the response to a single RPC request.</p>

A number of special keys in the DBM files can alter the way in which ypserv operates. The keys of interest are:

YP_INTERDOMAIN	The presence of this key causes ypserv to forward to a DNS server host lookups that cannot be satisfied by the DBM files.
YP_SECURE	This key causes ypserv to answer only questions coming from clients on reserved ports.
YP_MULTI_hostname	This is a special key in the form, YP_MULTI_hostname addr1,...,addrN. A client looking for hostname receives the “closest” address.

Two other functions supply information about the map, rather than map entries: yp_order(3NSL), and yp_master(3NSL). In fact, both order number and master name exist in the map as key-value pairs, but the server will not return either through the normal lookup functions. If you examine the map with makedbm(1M), however, they are visible. Other functions are used within the NIS service subsystem itself, and are not of general interest to NIS clients. They include do_you_serve_this_domain?, transfer_map, and reinitialize_internal_state.

ypserv	-d	The NIS service should go to the DNS for more host information. This requires the existence of a correct /etc/resolv.conf file pointing at a machine running in.named(1M). This option turns on DNS forwarding regardless of whether or not the YP_INTERDOMAIN flag is set in the hosts maps. See makedbm(1M). In the absence of an /etc/resolv.conf file, ypserv complains, but ignores the -d option.
	-u	Allow non-Trusted Solaris TCB machines to be clients.
	-v	Operate in the verbose mode, printing diagnostic messages to stderr.

SUMMARY OF TRUSTED SOLARIS CHANGES

ypserv can be started on the command line using ypstart(1M) from the Trusted Path with these privileges: file_dac_read, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl. The ypstop(1M) command must run from the Trusted Path with the proc_owner privilege

Upon startup, ypxfrd must inherit the net_broadcast, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl privileges.

-u allows non-Trusted Solaris TCB machines to be clients.

FILES

/var/yp/securenets	Defines the hosts and networks that are granted access to information in the served domain; it is read at startup time by both ypserv and ypxfrd.
/etc/init.d/rpc	Startup file that starts up basic RPC services and NIS by calling ypstart(1M). If the

ypserv(1M)

`/var/yp/ypserv.log` file exists when `ypserv` starts up, log information is written to it when error conditions arise. The file `/var/yp/binding/domainname/ypservers` is used to list the NIS server hosts that `ypbind` can bind to.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWypu

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

`in.named(1M)`, `ypbind(1M)`, `ypxfr(1M)`

`ypcat(1)`, `ypmatch(1)`, `ypwhich(1)`, `domainname(1M)`, `makedbm(1M)`, `ypinit(1M)`, `ypmake(1M)`, `ypoll(1M)`, `yppush(1M)`, `ypset(1M)`, `ypstart(1M)`, `ypstop(1M)`, `dbm_clearerr(3C)`, `ypclnt(3NSL)`, `libnsl(3LIB)`, `securenets(4)`, `ypfiles(4)`, `attributes(5)`

Network Interfaces Programmer's Guide

System Administration Guide, Volume 1

NOTES `ypserv` supports multiple domains. The `ypserv` process determines the domains it serves by looking for directories of the same name in the directory `/var/yp`. It replies to all broadcasts requesting yp service for that domain.

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications PLC, and must not be used without permission.

NAME	rpc.ypupdated, ypupdated – server for changing NIS information				
SYNOPSIS	/usr/lib/netsvc/yp/rpc.ypupdated [-isu]				
DESCRIPTION	<p>ypupdated is a daemon that updates information in the Network Information Service (NIS). ypupdated consults the updaters(4) file in the /var/yp directory to determine which NIS maps should be updated and how to change them. ypupdated must be run from the Trusted Path and inherit the net_broadcast, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl privileges.</p> <p>By default, the daemon requires the most secure method of authentication available to it, either DES (secure) or UNIX (insecure).</p>				
OPTIONS	<p>-i Accept RPC calls with the insecure AUTH_UNIX credentials. This allows programmatic updating of the NIS maps in all networks.</p> <p>-s Accept only calls authenticated using the secure RPC mechanism (AUTH_DES authentication). This disables programmatic updating of the NIS maps unless the network supports these calls.</p> <p>-u Allow updates from non-Trusted Solaris TCB clients.</p>				
SUMMARY OF TRUSTED SOLARIS CHANGES	<p>ypupdated must be run from the Trusted Path and inherit the net_broadcast, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl privileges.</p> <p>-u allows updates from non-Trusted Solaris TCB clients.</p>				
FILES	/var/yp/updaters Configuration file for rpc.updated command.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWypu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWypu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWypu				
SunOS 5.8 Reference Manual	<p>keyserv(1M), updaters(4), attributes(5)</p> <p><i>System Administration Guide, Volume I</i></p>				
NOTES	The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two services remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications PLC, and must not be used without permission.				

ypxfr(1M)

NAME	ypxfr, ypxfr_1perday, ypxfr_1perhour, ypxfr_2perday – transfer NIS map from a NIS server to host
SYNOPSIS	<code>/usr/lib/netsvc/yp/ypxfr</code> [-c] [-f] [-C <i>tid prog server</i>] [-d <i>ypdomain</i>] [-h <i>host</i>] [-s <i>ypdomain</i>] <i>mapname</i>
DESCRIPTION	<p>The <code>ypxfr</code> command moves an NIS map in the default domain for the local host to the local host by making use of normal NIS services. It creates a temporary map in the directory <code>/var/yp/ypdomain</code> (this directory must already exist; <i>ypdomain</i> is the default domain for the local host), fills it by enumerating the map's entries, fetches the map parameters (master and order number), and loads them. It then deletes any old versions of the map and moves the temporary map to the real <i>name</i>. <code>ypxfr</code> must be run from the Trusted Path and inherit the <code>net_privaddr</code> privilege.</p> <p>If run interactively, <code>ypxfr</code> writes its output to the terminal. However, if it is started without a controlling terminal, and if the log file <code>/var/yp/ypxfr.log</code> exists, it appends all its output to that file. Since <code>ypxfr</code> is most often run from the privileged user's <code>crontab</code> file, or by <code>ypserv</code>, the log file can retain a record of what was attempted, and what the results were.</p> <p>For consistency between servers, <code>ypxfr</code> should be run periodically for every map in the NIS data base. Different maps change at different rates: a map might not change for months at a time, for instance, and can therefore be checked only once a day. Some maps might change several times per day. In such a case, you might want to check hourly for updates. A <code>crontab(1)</code> entry can be used to automatically perform periodic updates. Rather than having a separate <code>crontab</code> entry for each map, you can group commands to update several maps in a shell script. Examples (mnemonically named) are in <code>/usr/sbin/yp: ypxfr_1perday, ypxfr_2perday, and ypxfr_1perhour</code>.</p> <p>Refer to <code>ypfiles(4)</code> for an overview of the NIS name service.</p>
OPTIONS	<p>-c Do not send a "Clear current map" request to the local <code>ypserv</code> process. Use this flag if <code>ypserv</code> is not running locally at the time you are running <code>ypxfr</code>. Otherwise, <code>ypxfr</code> complains that it cannot communicate with the local <code>ypserv</code>, and the transfer fails.</p> <p>-f Force the transfer to occur even if the version at the master is not more recent than the local version.</p> <p>-C <i>tid prog server</i> This option is for use <i>only</i> by <code>ypserv</code>. When <code>ypserv</code> starts <code>ypxfr</code>, it specifies that <code>ypxfr</code> should call back a <code>yppush</code> process at the host <i>server</i>, registered as program number <i>prog</i>, and waiting for a response to transaction <i>tid</i>.</p> <p>-d <i>ypdomain</i> Specify a domain other than the default domain.</p> <p>-h <i>host</i> Get the map from <i>host</i>, regardless of the master. If <i>host</i> is not specified, <code>ypxfr</code> asks the NIS service for the</p>

	ypxfr(1M)				
	name of the master, and tries to get the map from there. <i>host</i> must be a valid host name.				
	-s <i>ypdomain</i> Specify a source domain from which to transfer a map that should be the same across domains.				
SUMMARY OF TRUSTED SOLARIS FILES CHANGES	ypxfr must be run from the Trusted Path and inherit the net_privaddr privilege. /var/yp/ypxfr.log Log file /usr/lib/netsvc/yp/ypxfr_1perday Script to run one transfer per day, for use with cron(1M) /usr/lib/netsvc/yp/ypxfr_2perday Script to run two transfer per day, for use with cron(1M) /usr/lib/netsvc/yp/ypxfr_1perhour Script for hourly transfers of volatile maps /var/yp/ <i>ypdomain</i> NIS domain /usr/spool/cron/crontabs/root Privileged user's crontab file				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
ypxfr Only	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWnisu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWnisu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWnisu				
ypxfr_1perday, ypxfr_1perhour, and ypxfr_2perday	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWypu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWypu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWypu				
Trusted Solaris 8 4/01 Reference Manual	crontab(1), cron(1M), ypserv(1M) ypinit(1M), yppush(1M), ypfiles(4), attributes(5)				

ypxfr_1perday(1M)

NAME	ypxfr, ypxfr_1perday, ypxfr_1perhour, ypxfr_2perday – transfer NIS map from a NIS server to host
SYNOPSIS	<code>/usr/lib/netsvc/yp/ypxfr [-c] [-f] [-C <i>tid prog server</i>] [-d <i>ypdomain</i>] [-h <i>host</i>] [-s <i>ypdomain</i>] <i>mapname</i></code>
DESCRIPTION	<p>The <code>ypxfr</code> command moves an NIS map in the default domain for the local host to the local host by making use of normal NIS services. It creates a temporary map in the directory <code>/var/yp/ypdomain</code> (this directory must already exist; <i>ypdomain</i> is the default domain for the local host), fills it by enumerating the map's entries, fetches the map parameters (master and order number), and loads them. It then deletes any old versions of the map and moves the temporary map to the real <i>name</i>. <code>ypxfr</code> must be run from the Trusted Path and inherit the <code>net_privaddr</code> privilege.</p> <p>If run interactively, <code>ypxfr</code> writes its output to the terminal. However, if it is started without a controlling terminal, and if the log file <code>/var/yp/ypxfr.log</code> exists, it appends all its output to that file. Since <code>ypxfr</code> is most often run from the privileged user's <code>crontab</code> file, or by <code>ypserv</code>, the log file can retain a record of what was attempted, and what the results were.</p> <p>For consistency between servers, <code>ypxfr</code> should be run periodically for every map in the NIS data base. Different maps change at different rates: a map might not change for months at a time, for instance, and can therefore be checked only once a day. Some maps might change several times per day. In such a case, you might want to check hourly for updates. A <code>crontab(1)</code> entry can be used to automatically perform periodic updates. Rather than having a separate <code>crontab</code> entry for each map, you can group commands to update several maps in a shell script. Examples (mnemonically named) are in <code>/usr/sbin/yp: ypxfr_1perday, ypxfr_2perday, and ypxfr_1perhour</code>.</p> <p>Refer to <code>ypfiles(4)</code> for an overview of the NIS name service.</p>
OPTIONS	<p><code>-c</code> Do not send a "Clear current map" request to the local <code>ypserv</code> process. Use this flag if <code>ypserv</code> is not running locally at the time you are running <code>ypxfr</code>. Otherwise, <code>ypxfr</code> complains that it cannot communicate with the local <code>ypserv</code>, and the transfer fails.</p> <p><code>-f</code> Force the transfer to occur even if the version at the master is not more recent than the local version.</p> <p><code>-C <i>tid prog server</i></code> This option is for use <i>only</i> by <code>ypserv</code>. When <code>ypserv</code> starts <code>ypxfr</code>, it specifies that <code>ypxfr</code> should call back a <code>yppush</code> process at the host <i>server</i>, registered as program number <i>prog</i>, and waiting for a response to transaction <i>tid</i>.</p> <p><code>-d <i>ypdomain</i></code> Specify a domain other than the default domain.</p> <p><code>-h <i>host</i></code> Get the map from <i>host</i>, regardless of the master. If <i>host</i> is not specified, <code>ypxfr</code> asks the NIS service for the</p>

	ypxfr_1perday(1M)				
	name of the master, and tries to get the map from there. <i>host</i> must be a valid host name.				
	-s <i>ypdomain</i> Specify a source domain from which to transfer a map that should be the same across domains.				
SUMMARY OF TRUSTED SOLARIS FILES CHANGES	ypxfr must be run from the Trusted Path and inherit the net_privaddr privilege. /var/yp/ypxfr.log Log file /usr/lib/netsvc/yp/ypxfr_1perday Script to run one transfer per day, for use with cron(1M) /usr/lib/netsvc/yp/ypxfr_2perday Script to run two transfer per day, for use with cron(1M) /usr/lib/netsvc/yp/ypxfr_1perhour Script for hourly transfers of volatile maps /var/yp/ <i>ypdomain</i> NIS domain /usr/spool/cron/crontabs/root Privileged user's crontab file				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
ypxfr Only	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWnisu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWnisu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWnisu				
ypxfr_1perday, ypxfr_1perhour, and ypxfr_2perday	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWypu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWypu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWypu				
Trusted Solaris 8 4/01 Reference Manual SunOS 5.8 Reference Manual	crontab(1), cron(1M), ypserv(1M) ypinit(1M), yppush(1M), ypfiles(4), attributes(5)				

ypxfr_1perhour(1M)

NAME	ypxfr, ypxfr_1perday, ypxfr_1perhour, ypxfr_2perday – transfer NIS map from a NIS server to host
SYNOPSIS	<code>/usr/lib/netsvc/yp/ypxfr</code> [-c] [-f] [-C <i>tid prog server</i>] [-d <i>ypdomain</i>] [-h <i>host</i>] [-s <i>ypdomain</i>] <i>mapname</i>
DESCRIPTION	<p>The <code>ypxfr</code> command moves an NIS map in the default domain for the local host to the local host by making use of normal NIS services. It creates a temporary map in the directory <code>/var/yp/ypdomain</code> (this directory must already exist; <i>ypdomain</i> is the default domain for the local host), fills it by enumerating the map's entries, fetches the map parameters (master and order number), and loads them. It then deletes any old versions of the map and moves the temporary map to the real <i>name</i>. <code>ypxfr</code> must be run from the Trusted Path and inherit the <code>net_privaddr</code> privilege.</p> <p>If run interactively, <code>ypxfr</code> writes its output to the terminal. However, if it is started without a controlling terminal, and if the log file <code>/var/yp/ypxfr.log</code> exists, it appends all its output to that file. Since <code>ypxfr</code> is most often run from the privileged user's <code>crontab</code> file, or by <code>ypserv</code>, the log file can retain a record of what was attempted, and what the results were.</p> <p>For consistency between servers, <code>ypxfr</code> should be run periodically for every map in the NIS data base. Different maps change at different rates: a map might not change for months at a time, for instance, and can therefore be checked only once a day. Some maps might change several times per day. In such a case, you might want to check hourly for updates. A <code>crontab(1)</code> entry can be used to automatically perform periodic updates. Rather than having a separate <code>crontab</code> entry for each map, you can group commands to update several maps in a shell script. Examples (mnemonically named) are in <code>/usr/sbin/yp: ypxfr_1perday, ypxfr_2perday, and ypxfr_1perhour</code>.</p> <p>Refer to <code>ypfiles(4)</code> for an overview of the NIS name service.</p>
OPTIONS	<p>-c Do not send a "Clear current map" request to the local <code>ypserv</code> process. Use this flag if <code>ypserv</code> is not running locally at the time you are running <code>ypxfr</code>. Otherwise, <code>ypxfr</code> complains that it cannot communicate with the local <code>ypserv</code>, and the transfer fails.</p> <p>-f Force the transfer to occur even if the version at the master is not more recent than the local version.</p> <p>-C <i>tid prog server</i> This option is for use <i>only</i> by <code>ypserv</code>. When <code>ypserv</code> starts <code>ypxfr</code>, it specifies that <code>ypxfr</code> should call back a <code>yppush</code> process at the host <i>server</i>, registered as program number <i>prog</i>, and waiting for a response to transaction <i>tid</i>.</p> <p>-d <i>ypdomain</i> Specify a domain other than the default domain.</p> <p>-h <i>host</i> Get the map from <i>host</i>, regardless of the master. If <i>host</i> is not specified, <code>ypxfr</code> asks the NIS service for the</p>

	ypxfr_1perhour(1M)				
	name of the master, and tries to get the map from there. <i>host</i> must be a valid host name.				
	-s <i>ypdomain</i> Specify a source domain from which to transfer a map that should be the same across domains.				
SUMMARY OF TRUSTED SOLARIS FILES CHANGES	ypxfr must be run from the Trusted Path and inherit the <code>net_privaddr</code> privilege. /var/yp/ypxfr.log Log file /usr/lib/netsvc/yp/ypxfr_1perday Script to run one transfer per day, for use with <code>cron(1M)</code> /usr/lib/netsvc/yp/ypxfr_2perday Script to run two transfer per day, for use with <code>cron(1M)</code> /usr/lib/netsvc/yp/ypxfr_1perhour Script for hourly transfers of volatile maps /var/yp/ <i>ypdomain</i> NIS domain /usr/spool/cron/crontabs/root Privileged user's crontab file				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
ypxfr Only	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWnisu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWnisu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWnisu				
ypxfr_1perday, ypxfr_1perhour, and ypxfr_2perday	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWypu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWypu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWypu				
Trusted Solaris 8 4/01 Reference Manual SunOS 5.8 Reference Manual	<code>crontab(1)</code> , <code>cron(1M)</code> , <code>ypserv(1M)</code> <code>ypinit(1M)</code> , <code>yppush(1M)</code> , <code>ypfiles(4)</code> , <code>attributes(5)</code>				

ypxfr_2perday(1M)

NAME	ypxfr, ypxfr_1perday, ypxfr_1perhour, ypxfr_2perday – transfer NIS map from a NIS server to host
SYNOPSIS	<code>/usr/lib/netsvc/yp/ypxfr</code> [-c] [-f] [-C <i>tid prog server</i>] [-d <i>ypdomain</i>] [-h <i>host</i>] [-s <i>ypdomain</i>] <i>mapname</i>
DESCRIPTION	<p>The <code>ypxfr</code> command moves an NIS map in the default domain for the local host to the local host by making use of normal NIS services. It creates a temporary map in the directory <code>/var/yp/ypdomain</code> (this directory must already exist; <i>ypdomain</i> is the default domain for the local host), fills it by enumerating the map's entries, fetches the map parameters (master and order number), and loads them. It then deletes any old versions of the map and moves the temporary map to the real <i>name</i>. <code>ypxfr</code> must be run from the Trusted Path and inherit the <code>net_privaddr</code> privilege.</p> <p>If run interactively, <code>ypxfr</code> writes its output to the terminal. However, if it is started without a controlling terminal, and if the log file <code>/var/yp/ypxfr.log</code> exists, it appends all its output to that file. Since <code>ypxfr</code> is most often run from the privileged user's <code>crontab</code> file, or by <code>ypserv</code>, the log file can retain a record of what was attempted, and what the results were.</p> <p>For consistency between servers, <code>ypxfr</code> should be run periodically for every map in the NIS data base. Different maps change at different rates: a map might not change for months at a time, for instance, and can therefore be checked only once a day. Some maps might change several times per day. In such a case, you might want to check hourly for updates. A <code>crontab(1)</code> entry can be used to automatically perform periodic updates. Rather than having a separate <code>crontab</code> entry for each map, you can group commands to update several maps in a shell script. Examples (mnemonically named) are in <code>/usr/sbin/yp: ypxfr_1perday</code>, <code>ypxfr_2perday</code>, and <code>ypxfr_1perhour</code>.</p> <p>Refer to <code>ypfiles(4)</code> for an overview of the NIS name service.</p>
OPTIONS	<p>-c Do not send a "Clear current map" request to the local <code>ypserv</code> process. Use this flag if <code>ypserv</code> is not running locally at the time you are running <code>ypxfr</code>. Otherwise, <code>ypxfr</code> complains that it cannot communicate with the local <code>ypserv</code>, and the transfer fails.</p> <p>-f Force the transfer to occur even if the version at the master is not more recent than the local version.</p> <p>-C <i>tid prog server</i> This option is for use <i>only</i> by <code>ypserv</code>. When <code>ypserv</code> starts <code>ypxfr</code>, it specifies that <code>ypxfr</code> should call back a <code>yppush</code> process at the host <i>server</i>, registered as program number <i>prog</i>, and waiting for a response to transaction <i>tid</i>.</p> <p>-d <i>ypdomain</i> Specify a domain other than the default domain.</p> <p>-h <i>host</i> Get the map from <i>host</i>, regardless of the master. If <i>host</i> is not specified, <code>ypxfr</code> asks the NIS service for the</p>

	ypxfr_2perday(1M)				
	name of the master, and tries to get the map from there. <i>host</i> must be a valid host name.				
	-s <i>ypdomain</i> Specify a source domain from which to transfer a map that should be the same across domains.				
SUMMARY OF TRUSTED SOLARIS FILES CHANGES	ypxfr must be run from the Trusted Path and inherit the net_privaddr privilege. /var/yp/ypxfr.log Log file /usr/lib/netsvc/yp/ypxfr_1perday Script to run one transfer per day, for use with cron(1M) /usr/lib/netsvc/yp/ypxfr_2perday Script to run two transfer per day, for use with cron(1M) /usr/lib/netsvc/yp/ypxfr_1perhour Script for hourly transfers of volatile maps /var/yp/ypdomain NIS domain /usr/spool/cron/crontabs/root Privileged user's crontab file				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
ypxfr Only	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWnisu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWnisu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWnisu				
ypxfr_1perday, ypxfr_1perhour, and ypxfr_2perday	<table> <tr> <th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr> <tr> <td>Availability</td><td>SUNWypu</td></tr> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWypu
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWypu				
Trusted Solaris 8 4/01 Reference Manual	crontab(1), cron(1M), ypserv(1M) ypinit(1M), yppush(1M), ypfiles(4), attributes(5)				

ypxfrd(1M)

NAME	ypserv, ypxfrd – NIS server and binder processes
SYNOPSIS	<pre>/usr/lib/netsvc/yp/ypserv [-duv] /usr/lib/netsvc/yp/ypxfrd</pre>
DESCRIPTION	<p>The Network Information Service (NIS) provides a simple network lookup service consisting of databases and processes. The databases are ndbm files in a directory tree rooted at /var/yp. See dbm_clearerr(3C). These files are described in ypfiles(4). The processes are /usr/lib/netsvc/yp/ypserv, the NIS database lookup server, and /usr/lib/netsvc/yp/ypbind, the NIS binder. The programmatic interface to the NIS service is described in ypclnt(3NSL). Administrative tools are described in yppoll(1M), yppush(1M), ypset(1M), ypxfr(1M), and ypwhich(1). Tools to see the contents of NIS maps are described in ypcat(1), and ypmatch(1). Database generation and maintenance tools are described in ypinit(1M), ypmake(1M), and makedbm(1M).</p> <p>The ypserv utility is a daemon process typically activated at system startup from /etc/init.d/rpc. Alternatively, it can be started on the command line using ypstart(1M) from the Trusted Path with these privileges: file_dac_read, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl. ypserv runs only on NIS server machines with a complete NIS database. You can halt all NIS services using the ypstop(1M) command run from the Trusted Path with the proc_owner privilege.</p> <p>The ypxfrd utility transfers entire NIS maps in an efficient manner. For systems that use this daemon, map transfers are 10 to 100 times faster, depending on the map. To use this daemon, be sure ypxfrd is running on the master server. See /usr/lib/netsvc/yp/ypstart. ypxfr attempts to use ypxfrd first. If that fails, it prints a warning, then uses the older transfer method. ypxfrd must be run from the Trusted Path and inherit the net_broadcast, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl privileges.</p> <p>The ypserv daemon's primary function is to look up information in its local database of NIS maps.</p> <p>The operations performed by ypserv are defined for the implementor by the <i>YP Protocol Specification</i>, and for the programmer by the header file <rpcsvc/yp_prot.h>.</p> <p>Communication to and from ypserv is by means of RPC calls. Lookup functions are described in ypclnt(3NSL), and are supplied as C-callable functions in the libnsl(3LIB) library. There are four lookup functions, all of which are performed on a specified map within some NIS domain: yp_match(3NSL), yp_first(3NSL), yp_next(3NSL), and yp_all(3NSL). The yp_match operation takes a key, and returns the associated value. The yp_first operation returns the first key-value pair from the map, and yp_next can be used to enumerate the remainder. yp_all ships the entire map to the requester as the response to a single RPC request.</p>

A number of special keys in the DBM files can alter the way in which ypserv operates. The keys of interest are:

YP_INTERDOMAIN	The presence of this key causes ypserv to forward to a DNS server host lookups that cannot be satisfied by the DBM files.
YP_SECURE	This key causes ypserv to answer only questions coming from clients on reserved ports.
YP_MULTI_hostname	This is a special key in the form, YP_MULTI_hostname addr1,...,addrN. A client looking for hostname receives the “closest” address.

Two other functions supply information about the map, rather than map entries: yp_order(3NSL), and yp_master(3NSL). In fact, both order number and master name exist in the map as key-value pairs, but the server will not return either through the normal lookup functions. If you examine the map with makedbm(1M), however, they are visible. Other functions are used within the NIS service subsystem itself, and are not of general interest to NIS clients. They include do_you_serve_this_domain?, transfer_map, and reinitialize_internal_state.

ypserv	-d	The NIS service should go to the DNS for more host information. This requires the existence of a correct /etc/resolv.conf file pointing at a machine running in.named(1M). This option turns on DNS forwarding regardless of whether or not the YP_INTERDOMAIN flag is set in the hosts maps. See makedbm(1M). In the absence of an /etc/resolv.conf file, ypserv complains, but ignores the -d option.
	-u	Allow non-Trusted Solaris TCB machines to be clients.
	-v	Operate in the verbose mode, printing diagnostic messages to stderr.

SUMMARY OF TRUSTED SOLARIS CHANGES

ypserv can be started on the command line using ypstart(1M) from the Trusted Path with these privileges: file_dac_read, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl. The ypstop(1M) command must run from the Trusted Path with the proc_owner privilege

Upon startup, ypxfrd must inherit the net_broadcast, net_downgrade_sl, net_upgrade_sl, net_mac_read, net_privaddr, proc_setclr, and proc_setsl privileges.

-u allows non-Trusted Solaris TCB machines to be clients.

FILES

/var/yp/securenets	Defines the hosts and networks that are granted access to information in the served domain; it is read at startup time by both ypserv and ypxfrd.
/etc/init.d/rpc	Startup file that starts up basic RPC services and NIS by calling ypstart(1M). If the

ypxfrd(1M)

`/var/yp/ypserv.log` file exists when `ypserv` starts up, log information is written to it when error conditions arise. The file `/var/yp/binding/domainname/ypservers` is used to list the NIS server hosts that `ypbind` can bind to.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWypu

Trusted Solaris 8
4/01 Reference
Manual
SunOS 5.8
Reference Manual

`in.named(1M)`, `ypbind(1M)`, `ypxfr(1M)`

`ypcat(1)`, `ypmatch(1)`, `ypwhich(1)`, `domainname(1M)`, `makedbm(1M)`, `ypinit(1M)`, `ypmake(1M)`, `yppoll(1M)`, `yppush(1M)`, `ypset(1M)`, `ypstart(1M)`, `ypstop(1M)`, `dbm_clearerr(3C)`, `ypclnt(3NSL)`, `libnsl(3LIB)`, `securenets(4)`, `ypfiles(4)`, `attributes(5)`

Network Interfaces Programmer's Guide

System Administration Guide, Volume 1

NOTES

`ypserv` supports multiple domains. The `ypserv` process determines the domains it serves by looking for directories of the same name in the directory `/var/yp`. It replies to all broadcasts requesting yp service for that domain.

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications PLC, and must not be used without permission.

Index

A

accept — accept print requests, 22, 30, 402
add a new device driver to the system —
 add_drv, 26
add_allocatable
 add entries to allocation databases, 24
add_drv — add a new device driver to the
 system, 26
Address resolution display and control —
 arp, 32
arp — address resolution display and
 control, 32
atohexlabel — convert a character-coded label
 to its hexadecimal equivalent, 34
audit — maintain audit trail, 35
audit record, create — writeaudit, 620
audit records
 select or merge from audit trail files —
 auditreduce, 43
audit statistics report — auditstat, 52
audit trail file
 select records from — auditreduce, 43
auditconfig — get and set kernel audit
 parameters, 37
auditd — audit daemon, 41
auditing commands
 — audit, 35
 — auditconfig, 37
 — auditd, 41
 — auditreduce, 43
 — audit_startup, 51
 — auditstat, 52

auditing commands (*continued*)

 — audit_warn, 54
 — praudit, 386
auditreduce — select or merge audit records
 from audit trail files, 43
audit_startup shell script, 51
auditstat — display kernel audit statistics, 52
audit_warn — audit daemon warning
 script, 54
autofs
 automatically mount file systems —
 automount, 56
 mount/unmount request server —
 automountd, 64
automount — automatically mount file
 systems, 56
automountd — autofs mount/unmount request
 server, 64
autopush — configures lists of automatically
 pushed STREAMS modules, 65

B

boot parameter server — rpc.bootparamd, 67,
 427
broadcast message
 write to all users over a network —
 rwall, 448

C

CD-ROM

- managing — rmmount, 410

- mounting — rmmount, 410

change processor operational status —

- psradm, 392

chk_encodings — check label encodings file

- syntax, 70

chroot — change root directory for a

- command, 71

CIM Boot Manager

- starting, 168

CIM Object Manager

- stopping, 168

configure device policy — devpolicy, 84

configure system wide IPsec policy —

- ipseconf, 211

control and query bindings of processes to

- processors — pbind, 378

convert a character-coded label to its

- hexadecimal equivalent — atohexlabel, 34

convert a hexadecimal label to its

- character-coded equivalent —

- hextoalabel, 135

core file administration — coreadm, 73

coreadm — core file administration, 73

create audit record — writeaudit, 620

cron — clock daemon, 77

D

daemons

- clock daemon — cron, 77

- Internet Trivial File Transfer Protocol —

- in.tftpd, 210, 585

- network router discovery daemon —

- in.rdisc, 194, 398

- network status monitor — statd, 569

- NFS — nfsd, 341

- NIS+ service — rpc.nisd, 355, 434

- remote shell server — in.rshd, 205, 444

- server which returns peer process

- information — rpc.sprayd, 428

date

- set system date from a remote host —

- rdate, 397

devfsadm — administer /dev and

- /devices, 79, 81

devfsadmd — devfsadm daemon, 79, 81

device_clean — device clean programs, 83

device_maps

- display entries — dminfo, 93

devices

- add to allocation databases —

- add_allocatable, 24

/devices directory

- administer — devfsadm, 79, 81

devices

- display access control entries from

- device_maps, 93

- remove a device driver from the system —

- rem_drv, 404

- remove from allocation databases—

- remove_allocatable, 405

/devices directory

- configure — drvconfig, 95

devpolicy — configure device policy, 84

dfmounts — displays information on resources

- shared through DFS, 85

DFS

- display information on resources shared —

- dfmounts, 85

- list available resources from remote or local

- systems — dfshares, 87

dfshares — list available resources from remote

- or local systems, 87

disk usage

- summary — du, 98

disks

- partitioning and maintenance utility —

- format, 107

dispadmin — process scheduler

- administration, 89

display

- file system security attributes —

- getfsattr, 133

- system configuration information —

- prtconf, 388

display file system security attributes —

- getfsattr, 132

Distributed File System — see DFS, 85

dl_booting — inform the kernel that a machine

- is in the state of disklessly booting, 91, 92

dl_restore — inform the kernel that a machine is
in the normal state, 91, 92
dminfo — display device_maps entries, 93
drvconfig — configure /devices, 95
du — summarize disk usage, 98

E

EEPROM display and load program —
eeprom, 101

F

file system

- change the dynamic parameters —
setfsattr, 339, 466
- loopback — mount, 277, 606
- mount file systems and remote resources —
mount, 277, 284, 606, 613
- mount ufs — mount_ufs, 305
- report processes using file or file structure —
fuser, 130
- share multiple resources — shareall, 471,
616
- unmount — umount, 277, 284, 606, 613
- unshare multiple resources —
unshareall, 471, 616

File Transfer Protocol

- server — in.ftpd, 120, 153

floppy

- managing — rmmount, 410
- mounting — rmmount, 410

format — disk partitioning and maintenance
utility, 107

fsdb_ufs — ufs file system debugger, 111

- Commands, 114
- Expressions, 112
- Formatted Output, 117
- Inode Commands, 115

FTP

- daemon on remote host — in.ftpd, 120, 153

fuser — identify processes using file or file
structure, 130

G

getfsattr — display file system security
attributes, 132, 133

H

halt — stop the processor, 134, 385
hextoalabel — convert a hexadecimal label to its
character-coded equivalent, 135
hsfs
mount — mount_hsfs, 288

I

ICMP

- router discovery daemon — in.rdisc, 194,
398

ifconfig — configure network interface
parameters, 136

inetd — Internet services daemon, 150

in.ftpd — File Transfer Protocol daemon on
remote host, 120, 153

init — process control initialization, 163, 580
/etc/defaults/init file, 164, 581
init and System Booting, 163, 580
inittab Additions, 164, 581
Run Level Changes, 164, 581
Run Level Defined, 163, 580
telinit, 165, 582

init.wbem — start/stop CIM Boot
Manager, 168

in.named — internet domain name server, 171,
310

in.rarpd — Reverse Address Resolution
Protocol Server, 192, 395

in.rdisc — ICMP router discovery
daemon, 194, 398

in.rexecd — remote execution server, 196, 406

in.rlogind — remote login server, 198, 408

in.routed — network routing daemon, 200, 420

install — install commands, 208

Internet

- File Transfer Protocol daemon on remote
host — in.ftpd, 120, 153

Internet (*continued*)

- ICMP router discovery daemon —
 - in.rdisc, 194, 398
- network routing daemon — in.routed, 200, 420
- query domain name servers —
 - nslookup, 369
- RARP server — in.rarpd, 192, 395
- services daemon — inetd, 150
- Trivial File Transfer Protocol server —
 - in.tftpd, 210, 585
- Internet Control Message Protocol
 - in.rdisc, 194, 398
- Internet Protocol
 - to Ethernet addresses — arp, 32
- in.tftpd — Internet Trivial File Transfer Protocol server, 210, 585
- ipsecconf — configure system wide IPsec policy, 211
- ipseckey — manually manipulate an IPsec Security Association Database (SABD), 226

K

- kernel
 - load a module — modload, 275
 - unload a module — modunload, 276

L

- lockd — network lock daemon, 235
- loopback file system
 - mount file systems and remote resources —
 - mount, 277, 606
- LP print services
 - administer filters — lpfilter, 248
 - administer forms — lpforms, 254
 - configure — lpadmin, 236
 - register remote systems — lpssystem, 266
 - set printing queue priorities — lpusers, 267
- lpadmin — configure LP print service, 236
- lpfilter — administer filters used with LP print service, 248
- lpforms — administer forms used with LP print service, 254

- lpforms — administer forms used with the LP print service
 - Adding or Changing a Form, 254
 - Allowing and Denying Access to a Form, 257
- lpforms — administer forms used with the Deleting a Form, 256
- lpforms — administer forms used with the LP print service
 - Listing Form Attributes, 257
 - Listing the Current Alert, 259
 - Removing an Alert Definition, 259
 - Setting an Alert to Mount a Form, 257
 - Terminating an Active Alert, 259
- lpmove — moves print requests that are queued, 261
- lpsched — start the LP print service, 263
- lpshut — stop the LP print service, 265
- lpssystem — register remote systems with LP print service, 266
- lpusers — set printing queue priorities, 267

M

- mail delivery server — sendmail, 449
- make device_allocate and device_maps entries —
 - mkdevdb, 269, 271, 273
- make local NFS file systems available for mounting by remote systems —
 - share_nfs, 472
- manage bulk operations on user accounts —
 - smmultiuser, 512
- manage email alias entries — smmaillist, 508
- manage entries in the exec_attr database —
 - smexec, 493
- manage entries in the hosts database —
 - smhost, 503
- manage entries in the interface database —
 - smnetidb, 517
- manage entries in the network template database — smnettmpl, 521
- manage entries in the networks database —
 - smnetwork, 526
- manage group entries — smgroup, 499
- manage jobs in the crontab database —
 - smcron, 487

- manage profiles in the `prof_attr` and `exec_attr` databases — `smprofile`, 530
- manage roles and users in role accounts — `smrole`, 536
- manage user entries — `smuser`, 545
- manually manipulate an IPsec Security Association Database (SABD) — `ipseckey`, 226
- `mkdevalloc` — make device_allocate entries, 269, 271, 273
- `mkdevdb` — make device_allocate and device_maps entries, 269, 271, 273
- `mkdevmaps` — make device_maps entries, 269, 271, 273
- `modload` — load a kernel module, 275
- `modunload` — unload a kernel module, 276
- `mount` — mount file systems and remote resources, 277, 606
- `mount`
 - show all remote mounts — `showmount`, 481
- `mount_hfs` file systems — `mount_hfs`, 288
- `mount_pcfs` file systems — `mount_pcfs`, 300
- `mountall` — mount multiple filesystems, 284, 613
- `mountd` — NFS mount request server, 286
- `mount_hfs` — mount hfs file systems, 288
- `mount_nfs` — mount remote NFS resources, 291
- `mount_pcfs` — mount pcfs file systems, 300
- `mount_tmpfs` — mount tmpfs, 302
- `mount_ufs` — mount ufs, 305

N

- name service cache daemon — `nsd`, 367
- `named` — internet domain name server, 171, 310
- `ndd` — get and set driver configuration parameters, 331
- `netstat` — display network status, 333
 - Active Sockets (First Form), 334
 - DHCP Interface Information (Seventh Form), 337
 - Extended Metric Routing Table (Fifth Form), 337

- `netstat` — display network status (*continued*)
 - Interface Status (Third Form), 336
 - Multicast Routing Tables (Sixth Form), 337
 - Network Data Structures (Second Form), 335
 - Routing Table (Fourth Form), 336
 - TCP Sockets, 335
- network routing daemon — `in.routed`, 200, 420
- network
 - lock daemon — `lockd`, 235
 - network interface parameters
 - configure — `ifconfig`, 136
 - network packets capture and inspection — `snoop`, 556
 - network status, display — `netstat`, 333
- `newsecfs` — set security attributes on a newly created file system, 339, 466
- NFS
 - crash and recovery functions for locking services — `statd`, 569
 - daemon — `nfsd`, 341
 - display statistics — `nfsstat`, 343
 - make local NFS filesystem unavailable for mounting by remote systems — `unshare_nfs`, 617
 - `mount` — `mount_nfs`, 291
 - mount request server — `mountd`, 286
- `nfsd` — NFS daemon, 341
- `nfsstat` — display NFS statistics, 343
- NIS
 - binder process — `ybind`, 623
- NIS+
 - initialize a domain to store system administration information—`nissetup`, 366
 - `nissetup` — initialize a NIS+ domain to serve clients, 366
 - service daemon — `rpc.nisd`, 355, 434
- NIS
 - transfer NIS map from a NIS server to host — `ypxfr`, 632, 634, 636, 638
- NIS+
 - utility to cache location information about NIS+ servers — `nis_cachemgr`, 348
- NIS+ credentials for NIS+ principals
 - initialize — `nisclient`, 350

- NIS+ password update daemon
 - nispasswd, 359, 438
 - rpc.nispasswd, 359, 438
- NIS server and binder processes
 - ypserv, 628, 640
 - ypxfrd, 628, 640
- nis_cachemgr — NIS+ utility to cache location information about NIS+ servers, 348
- nisclient — initialize NIS+ credentials for NIS+ principals, 350
- nispawdd — NIS+ password update daemon, 359, 438
- nispopulate — populate the NIS+ tables in a NIS+ domain, 361
- nissetup — initialize a domain to serve clients, 366
- nscd — name service cache daemon, 367
- nslookup — query Internet domain name servers, 369

O

- output system definition
 - display current — sysdef, 576
- override privilege, 19

P

- pbind — control and query bindings of processes to processors, 378
 - Binding processes, 378
 - Querying Bindings, 379
 - Unbinding a process, 379
- pcfs
 - mount — mount_pcfs, 300
- pkgchk — check package installation accuracy, 382
- populate the NIS+ tables in a NIS+ domain — nispopulate, 361
- poweroff — stop the processor, 134, 385
- praudit — display audit trail, 386
- print queue
 - accept or reject requests — accept, reject, 22, 402

- print requests
 - accept or reject — accept, reject, 22, 402
- print service, LP — lpmove, 261
- printer filters
 - add and change — lpfilter, 248
 - list attributes — lpfilter, 248
 - remove — lpfilter, 248
- printer forms
 - add or change — lpforms, 254
 - delete — lpforms, 256
 - list attributes — lpforms, 257
 - listing the current alert — lpforms, 259
 - provide access — lpforms, 257
 - removing an alert definition — lpforms, 259
 - setting an alert to mount a form — lpforms, 257
 - terminating an active alert — lpforms, 259
- printers
 - add and change printers — lpadmin, 236
 - define alerts for printer faults — lpadmin, 236
 - mount printer wheels — lpadmin, 236
 - remove printers — lpadmin, 236
 - set or change system default destination — lpadmin, 236
 - setting priorities — lpusers, 267
- privilege
 - override, 19
 - required, 18
- process scheduler
 - administration — dispadmin, 89
- processes
 - initialization — init, 163, 580
 - using file or file structure — fuser, 130
- programming tools
 - install — install commands, 208
- PROM monitor program
 - display and load program — eeprom, 101
- prtconf — print system configuration information, 388
- psradm — change processor operational status, 392

Q

- quick halt — halt, 134, 385

R

RARP

server — in.rarpd, 192, 395

rarpd — DARPA Reverse Address Resolution Protocol server, 192, 395

rdate — set system date from a remote host, 397

rdisc — network router discovery daemon, 194, 398

reboot — restart the operating system, 400

reject — reject print requests, 22, 402

rem_drv — remove a device driver from the system, 404

remote execution server — in.rexecd, 196, 198, 406, 408

remote login server — in.rlogind
rlogind, 198, 408

remote resources

mount or unmount — mount, 277, 606

mount NFS — mount_nfs, 291

remote system

make local resource unavailable for mounting — unshare, 615

register with LP print service —
lpssystem, 266

set system date — rdate, 397

shell server — in.rshd, 205, 444

removable media mounter for CD-ROM and floppy — rmmount, 410

remove_allocatable

remove entries to allocation databases, 405

required privilege, 18

Reverse Address Resolution Protocol
— in.rarpd, 192, 395

rlogind — remote login server, 198, 408

rmmount — removable media mounter for CD-ROM and floppy, 410

root directory

change for a command — chroot, 71

route — manually manipulate routing tables, 413

routed — network routing daemon, 200, 420

RPC

NIS+ service daemon — rpc.nisd, 355, 434

program number to universal addresses
mapping — rpcbind, 425

report information — rpcinfo, 429

RPC (continued)

sends one-way stream of packets to host —
spray, 568

server, autofs mount/unmount requests —
automountd, 64

server, NFS mount requests — mountd, 286

server which returns peer process
information — rpc.sprayd, 428

rpcbind — converts RPC program numbers to
universal addresses, 425

rpc.bootparamd — boot parameter server, 67,
427

rpc.getpeerinfod — Obtain peer process
information, 428

rpcinfo — report RPC information, 429

rpc.nisd — NIS+ service daemon, 355, 434
rpc.nisd_resolv, 355, 358, 434, 437

rpc.nispasswd — NIS+ password update
daemon, 359, 438

rpc.tbootparamd — Trusted Solaris boot
parameter server, 440

rpc.yppupdated — server for changing NIS
information, 443, 631

runpd

run a command for privilege
debugging, 447

rwall — write to all users over a network, 448

S

scheduler, process

administration — dispadmin, 89

sendmail — mail delivery system, 449

server for changing NIS information
— rpc.yppupdated, 443, 631
— ypupdated, 443, 631

servers

automountd — mount/unmount request
server, 64

inetd — Internet services daemon, 150

in.rexecd — remote execution server, 196,
406

mountd — mount request server, 286

RARP server — in.rarpd, 192, 395

yppasswdd — NIS password server, 441,
626

servers, NIS+
 location information — nis_cachemgr, 348
 setaudit — run a command with the audit mask set, 465
 setfsattr — tune up an existing file system, 339, 466
 setuname — changes machine information, 468
 share — make local resource available for mounting by remote systems, 469
 shareall — make multiple resources available for mounting, 471, 616
 share_nfs — make local NFS file system available for mounting by remote systems, 472
 shell
 remote shell server — in.rshd, 205, 444
 showmount — display remote mounts, 481
 smc — start the Solaris Management Console (SMC), 482
 smcron — manage jobs in the crontab database, 487
 smexec — manage entries in the exec_attr database, 493
 smgroup — manage group entries, 499
 smhost — manage entries in the hosts database, 503
 smmaillist — manage email alias entries, 508
 smmultiuser — manage bulk operations on user accounts, 512
 smnetidb — manage entries in the hosts database, 517
 smnetmpl — manage entries in the network template database, 521
 smnetwork — manage entries in the networks database, 526
 smprofile — manage profiles in the prof_attr and exec_attr databases, 530
 smrole — manage roles and users in role accounts, 536
 smuser — manage user entries, 545
 snoop — capture and inspect network packets, 556
 software package
 check installation accuracy — pkgchk, 382
 spray — sends one-way stream of packets to host, 568
 start the LP print service — lpsched, 263
 start the Solaris Management Console (SMC) — smc, 482
 statd — network status monitor, 569
 statistics
 audit — auditstat, 52
 NFS, display — nfsstat, 343
 stop the processor — halt, 134, 385
 stop the processor — poweroff
 poweroff, 134, 385
 stop the LP print service — lpshut, 265
 STREAMS
 automatically pushed modules — autopush, 65
 su — become a non-role user, 570
 super user command — su, 570
 swap — administer the system swap areas, 573
 sysdef — displays current system definition, 576
 sysh — system shell, 578
 system administration
 control for basic administrative functions — uadmin, 605
 install commands — install, 208
 system configuration
 print information — prtconf, 388
 system definition
 display current — sysdef, 576
 system parameters
 change value — setuname, 468
 system shutdown
 — halt, 134, 385

T

tbootparam — send a request to rpc.tbootparamd to inform it that a host is in normal (labeled) state now, 579
 TCP/IP
 File Transfer Protocol daemon on remote host — in.ftpd, 120, 153
 telinit — process control initialization, 163, 580
 tftpd — Internet Trivial File Transfer Protocol server, 210, 585

timed event services
 daemon for cron — cron, 77
 tmpfs
 mount — mount_tmpfs, 302
 tnchddb — check file syntax of trusted network
 databases, 586
 tnctl — configure Trusted Solaris
 network-daemon control parameters, 588
 tnd — Trusted network daemon, 590
 tninfo — print information and statistics about
 kernel-level network, 592
 tokmapctl — configure token-mapping
 daemon, 594
 tokmapd — token-mapping daemon, 596
 traceroute — print the route packets take to
 network host, 598

U

uadmin — administrative control, 605
 ufs
 mount — mount_ufs, 305
 ufs file system debugger — fsdb_ufs, 111
 umount — unmount file systems and remote
 resources, 277, 606
 umountall — unmount multiple file
 systems, 284, 613
 unshare — make local resource unavailable for
 mounting by remote systems, 615
 unshareall — make multiple resources
 unavailable for mounting, 471, 616
 unshare_nfs — make local NFS filesystem
 unavailable for mounting by remote
 systems, 617
 updatehome — update the home directory copy
 and link files for the current label, 618
 user IDs
 become a non-role user — su, 570

V

Volume Management
 removable media mouter —
 rmmount, 410

W

writeaudit — write an audit record, 620

Y

ypbind — NIS binder process, 623
 ypserv — NIS server and binder
 processes, 628, 640
 ypserv, 629, 641
 ypupdated — server for changing NIS
 information, 443, 631
 ypxfr — transfer NIS map from a NIS server to
 host, 632, 634, 636, 638
 ypxfrd — NIS server and binder
 processes, 628, 640
 yppasswdd — NIS password server, 441, 626

