System Administration Guide: Advanced Administration

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

Part No: 817–0403–10
January 2005

Adobe PostScript

040819@9495

Contents

Preface

System Administration Guide: Advanced Administration is part of a set that covers a significant part of the Solaris™ system administration information. This guide includes information for both SPARC® and x86 based systems.

This book assumes that you have installed the SunOS™ 5.10 Operating System. It also assumes that you have set up any networking software that you plan to use. The SunOS 5.10 Operating System is part of the Solaris 10 product family, which also includes many features, including the Solaris Common Desktop Environment (CDE). The SunOS 5.10 Operating System is compliant with AT&T's System V, Release 4 operating system.

For the Solaris 10 release, new features that are interesting to system administrators are covered in sections called *What's New in ... ?* in the appropriate chapters.

Note – This Solaris release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC®, SPARC64, IA-32, AMD64. The supported SPARC based systems are listed in the *Solaris Sun Hardware Platform Guide* at `http://docs.sun.com`. The supported x86 based systems appear in the *Solaris Hardware Compatibility List* at `http://www.sun.com/bigadmin/hcl`. This document cites any implementation differences between the platform types.

In this document the term "x86" refers to the Intel 32-bit family of microprocessors and compatible 64-bit and 32-bit microprocessors made by AMD. For supported systems, see the *Solaris Hardware Compatibility List*.

Who Should Use This Book

This book is intended for anyone responsible for administering one or more systems that are running the Solaris 10 release. To use this book, you should have 1-2 years of UNIX® system administration experience. Attending UNIX system administration training courses might be helpful.

How the System Administration Volumes Are Organized

Here is a list of the topics that are covered by the volumes of the System Administration Guides.

Book Title	Topics
System Administration Guide: Basic Administration	User accounts and groups, server and client support, shutting down and booting a system, and managing software (packages and patches)
System Administration Guide: Advanced Administration	Printing services, terminals and modems, system resources (disk quotas, accounting, and crontabs), system processes, and troubleshooting Solaris software problems
System Administration Guide: Devices and File Systems	Removable media, disks and devices, file systems, and backing up and restoring data
System Administration Guide: IP Services	TCP/IP networks, IPv4 and IPv6, DHCP, IP Security, Mobile IP, IP Network Multipathing, and IPQoS
System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)	DNS, NIS, and LDAP naming and directory services
System Administration Guide: Naming and Directory Services (NIS+)	NIS+ naming and directory services
System Administration Guide: Network Services	Web cache servers, time-related services, network file systems (NFS and Autofs), mail, SLP, and PPP
System Administration Guide: Security Services	Auditing, device management, file security, BART, PAM, Solaris cryptographic framework, privileges, RBAC, SASL, Solaris Secure Shell, and SEAM

Book Title	Topics
System Administration Guide: N1 Grid Containers, Resource Management, and Solaris Zones	Resource management topics such as projects and tasks, extended accounting, resource controls, and dynamic resource pools; also covers virtualization with Solaris Zones

To view license terms, attribution, and copyright statements for open source software included in this Solaris release, the default path is `/usr/share/src/`*freeware-name* or `/usr/sfw/share/src/`*freeware-name*. If the Solaris Operating System has been installed anywhere other than the default location, modify the given path to access the file at the installed location.

Related Third-Party Web Site References

Note – Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with the use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Accessing Sun Documentation Online

The docs.sun.com[SM] web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is http://docs.sun.com.

What Typographic Conventions Mean

The following table describes the typographic conventions that are used in this book.

TABLE P–1 Typographic Conventions

Typeface or Symbol	Meaning	Example
`AaBbCc123`	The names of commands, files, and directories; on-screen computer output	Edit your `.login` file. Use `ls -a` to list all files. `machine_name% you have mail.`
`AaBbCc123`	What you type, contrasted with on-screen computer output	`machine_name%` **`su`** `Password:`
AaBbCc123	Command-line placeholder: replace with a real name or value	To delete a file, type **`rm`** *filename*.
AaBbCc123	Book titles, new words or new terms, or words to be emphasized.	Read Chapter 6 in *User's Guide*. These options are called *class* options. Do *not* save changes yet.

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P–2 Shell Prompts

Shell	Prompt
C shell prompt	`machine_name%`
C shell superuser prompt	`machine_name#`
Bourne shell and Korn shell prompt	`$`
Bourne shell and Korn shell superuser prompt	`#`

General Conventions

Be aware of the following conventions that are used in this book.

- When following steps or using examples, be sure to type double-quotes (`"`), left single-quotes (`‘`), and right single-quotes (`’`) exactly as shown.

- The key referred to as Return is labeled Enter on some keyboards.
- It is assumed that the root path includes the `/sbin`, `/usr/sbin`, `/usr/bin`, and `/etc` directories, so the steps in this book show the commands in these directories without absolute path names. Steps that use commands in other, less common, directories show the absolute path in the example.
- The examples in this book are for a basic SunOS 5.10 software installation without the Binary Compatibility Package installed and without `/usr/ucb` in the path.

Caution – If `/usr/ucb` is included in a search path, it should always be at the end of the search path. Commands like `ps` or `df` are duplicated in `/usr/ucb` with different formats and different options from the SunOS 5.10 commands.

CHAPTER 1

Managing Printing Services (Overview)

This chapter provides information about managing printers in the Solaris Operating System (Solaris OS).

This is a list of the overview information in this chapter.

What's New in Printing?

This section describes new printing features in the current Solaris release.

Expanded Printer Support in the Solaris Operating System

In this Solaris release, modifications have been made to incorporate support for a wide array of printers. This functionality differs greatly from previous Solaris software releases. In previous releases, you could only print to printers that understood PostScript™ natively, or you could only print plain ASCII text. The list of supported printer types, and information about whether these printer types accepted PostScript or ASCII text, was limited. Now, through the use of additional transformation software, raster image processor (RIP), and PostScript Printer Description (PPD) files, you can print to a wider range of printers.

For more information on the expanded printer support features that are available in this release, see "Expanded Printer Support" on page 36.

Internet Printing Protocol Listener for the Solaris Operating System

The Internet Printing Protocol (IPP) listener is a new feature for printing on a network. IPP works in conjunction with the Apache web server that is shipped as part of the Solaris software, and IPP uses HTTP to transport requests. Using the Apache web server, the listener loads IPP print requests and processes them. The IPP listener listens on port 631 for HTTP requests and then communicates the requests to the printing system. For more information on the IPP listener, see "Internet Printing Protocol Listener" on page 41.

Where to Find Printer Tasks

Printer Task	For More Information
Set up printers with Solaris Print Manager and the LP print service	Chapter 3
Administer printers with the LP print service after the printers are set up	Chapter 4
Manage character sets, filters, forms, and fonts	Chapter 5
Customize the LP print service	Chapter 6

Printing in the Solaris Operating System

The Solaris printing software provides an environment for setting up and managing client access to printers on a network.

The Solaris printing software contains these tools:

- Solaris Print Manager – A graphical user interface (GUI) that provides the ability to manage printing configuration on a local system or in a name service.
- The LP print service commands – A command-line interface (CLI) that is used to set up and manage printers on a local system or in a name service. These commands also provide functionality that extend beyond the other print

management tools.

Even if you do use Solaris Print Manager to set up printing, you will have to use some of the `lp` print service commands to completely manage printing on the Solaris Operating System. For more information, see Chapter 4.

The limitations of Solaris print software include the following:

- No support for print servers defined as `s5` (the System V print protocol) in previous Solaris releases
- No print filtering on print clients

Solaris Print Manager

Solaris Print Manager is a Java technology-based GUI that enables you to manage local and remote printer configuration. This tool can be used in the following name service environments: LDAP, NIS, NIS+, and files. You must be logged in as superuser or assume an equivalent role to use this tool.

Solaris Printer Manager centralizes printer information when used in conjunction with a name service. Using a name service for storing printer configuration information is desirable because a name service makes printer information available to all systems on the network. This method provides easier printing administration.

Solaris Print Manager recognizes existing printer information on the printer servers, print clients, and in the name service databases. No conversion tasks are required to use Solaris Print Manager as long as the print clients are running either the Solaris 2.6, 7, 8, 9, or 10 releases.

The Solaris Print Manager package is `SUNWppm`.

For step-by-step instructions on setting up new printers with Solaris Print Manager, see Chapter 3.

Printing Support in the Name Service Switch

The `printers` database in `/etc/nsswitch.conf`, the name service switch file, provides centralized printer configuration information to print clients on the network.

By including the `printers` database and corresponding sources of information in the name service switch file, print clients automatically have access to printer configuration information without having to add it to their own systems.

The default printers entry in the `/etc/nsswitch.conf` file for files, LDAP, NIS, and NIS+ environments are described in the following table. The `nisplus` keyword represents the `printers.org_dir` table.

Name Service Type	Default `printers` Entry
`files`	`printers: user files`
`ldap`	`printers: user files ldap`
`nis`	`printers: user files nis`
`nis+`	`printers: user nisplus files`

For example, if your name service is NIS, printer configuration information on print clients is searched for in the following sources in this order:

- `user` – Represents the user's `$HOME/.printers` file
- `files` – Represents the `/etc/printers.conf` file
- `nis` – Represents the `printers.conf.byname` table

For more information, see the `nsswitch.conf`(4) and *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)* man pages.

LDAP Print Support Guidelines

Keep the following in mind when you manage printer information in the LDAP name service:

- Solaris Print Manager uses the `ldapclient` command to determine the default LDAP server name. If more than one server is specified, the first server is automatically selected. For more information, see `ldapclient`(1M).
- Solaris Print Manager always displays printer entries from the current LDAP server. If the current LDAP server is not the domain's LDAP master server, the list of printers displayed might not be the current list of printers. The LDAP replica server might not have been updated by the master server, and as a result, is out of sync with the master. Replica servers can have various update replication agreements. For example, a change is made on the master, and the replica servers can be updated immediately or once a day.
- If the selected LDAP server is an LDAP replica server, any updates are referred to the master server and are done there. This situation again means that the printer list could be out of sync with the master server. For example, a deleted printer might still appear in the displayed printer list until the replica is updated from the master server.
- Users can use the LDAP command-line utilities, `ldapadd` and `ldapmodify`, to update printer entries in the directory. However, this method is not recommended. If these utilities are used, the user must ensure that the *printer-name* attribute value is unique within the `ou=printers` container. If the *printer-name* attribute value is not unique, the result of modifications done by Solaris Print Manager, or by the `lpset` command, might not be predictable.

Choosing a Method to Manage Printers

Adding printer information to a name service makes access to printers available to all systems on the network. Generally, doing so makes printer administration easier because all the information about printers is centralized.

Name Service Configuration	Actions to Centralize Printer Information
Use a name service	Adding the printer to the LDAP, NIS, or NIS+ database makes the printer available to all systems on the network.
Do not use a name service	Adding the printer adds the printer information to the printer server's configuration files only. Print clients will not automatically know about the printer. You will have to add the printer information to every print client that needs to use the printer.

The following table describes the major printer-related tasks and the tools available to perform the printing tasks.

TABLE 1–1 Solaris Printing Tools and Their Features

Printing Tool	Availability	Graphical User Interface?	Configures Network Printers?	Manages Print Clients and Servers?	Uses LDAP, NIS, or NIS+?
Solaris Print Manager	Solaris 10 and compatible versions and Solaris Easy Access Server 3.0	Yes	Yes	Yes	Yes
`lp` commands	Solaris 10 and compatible versions	No	Yes	Yes	Yes

After using the preceding table to determine which printing tool is best for your network environment, see Chapter 3 for printer setup information.

Most printing configuration tasks can be accomplished with Solaris Print Manager. However, if you need to write interface scripts or add your own filters, you need to use the LP print service commands. These commands underlie Solaris Print Manager. Performing printing administration tasks with LP print service commands is described in Chapter 4.

The LP Print Service

The *LP print service* is a set of software utilities that allows users to print files while the users continue to work.

For background information about the LP print service, see Chapter 7.

Other LP print service topics are described in the following sections. Included are cross-references for further information.

Managing Network Printers

A *network printer* is a hardware device that is connected directly to the network. A network printer transfers data directly over the network to the output device. The printer or network connection hardware has its own system name and IP address.

Network printers often have software support provided by the printer vendor. If your printer has printer vendor-supplied software, then use the printer vendor software. If the network printer vendor does not provide software support, Sun supplied software is available. This software provides generic support for network-attached printers. However, this software is not capable of providing full access to all possible printer capabilities.

For step-by-step instructions on setting up a network printer, see Chapter 3.

Administering Printers

After you set up print servers and print clients, you might need to perform these administration tasks frequently:

- Delete a printer
- Delete remote printer access
- Check the status of printers
- Restart the print scheduler

For step-by-step instructions on how to perform these printer administration tasks, see Chapter 4.

Setting Definitions for Printers

Setting definitions for the printers on your network is an ongoing task that lets you provide a more effective print environment for users. For example, you can assign printer descriptions for all your site's printers to help users find where a printer is located. Or, you can define a class of printers to provide the fastest turnaround for print requests.

For information on setting up printer definitions, see Chapter 2.

Administering Character Sets, Filters, Forms, and Fonts

Depending on your site's requirements and the types of printers you have on the network, you might have to set up and administer printer-specific features of the LP print service. For example, you can assign different print wheels, filters, and forms to different printers. For background information and step-by-step instructions on how to set up character sets, print filters, forms, and fonts, see Chapter 5.

Customizing the LP Print Service

Although the LP print service is designed to be flexible enough to handle most printers and printing needs, the LP print service does not handle every possible situation. You might have a print request that is not accommodated by the standard features of the LP print service. Or, you might have a printer that does not quite fit into the way the LP print service handles printers.

You can customize the LP print service in the following ways:

- Adjust the printer port characteristics.
- Adjust the `terminfo` database.
- Customize the printer interface program.
- Create a print filter.
- Define a form.

For detailed descriptions and step-by-step instructions on customizing the LP print service, see Chapter 6.

The Solaris Print Client-Server Process

This section provides an overview of how Solaris printing works.

The Print Client Process

The following figure illustrates the path of a print request from the time the user initiates the request until the print request is printed.

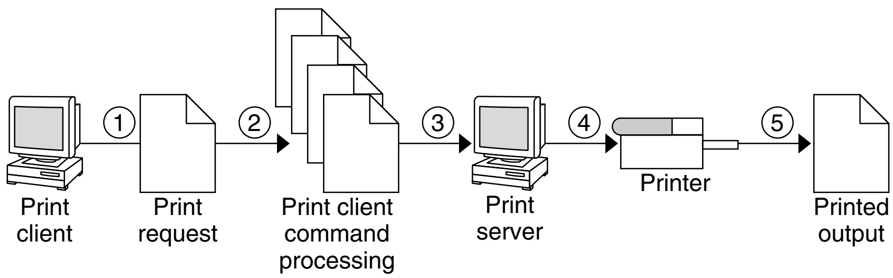

FIGURE 1–1 Overview of the Print Client Process

1. A user submits a print request from a print client.
2. The print command checks a hierarchy of print configuration resources to determine where to send the print request.
3. The print command sends the print request directly to the appropriate print server. A print server can be any server that accepts Berkley Software Distribution (BSD) printing protocol, including SVR4 (LP) print servers and BSD `lpd`-based print servers.
4. The print server sends the print request to the appropriate printer.
5. The print request is printed.

Using Print Clients

This section of the overview focuses on the *print client*, a system that can send print requests to a print server. This section also discusses print commands, which enable the print client to initiate print requests.

The following figure highlights the part of the print process in which the user submits a print request from a print client.

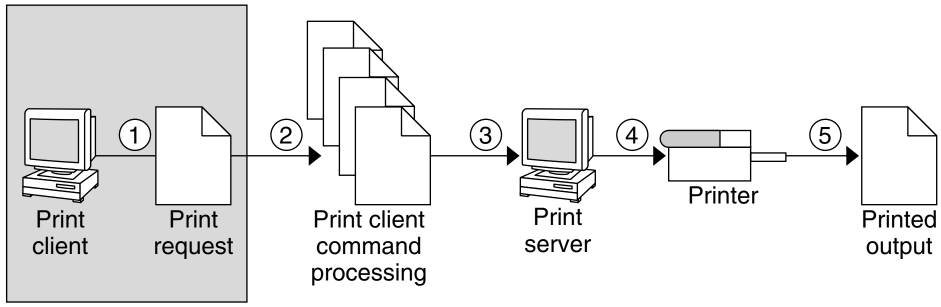

FIGURE 1–2 The User Submits a Print Request from a Print Client

What Is a Print Client?

A system becomes a print client when you install the Solaris print software and enable access to remote printers on the system.

The Solaris print software checks the following resources to locate printers and printer configuration information:

- The command-line interface by using the `lp -d` command (atomic or POSIX format)
- A user's `LPDEST` or `PRINTER` variables
- The `_default` variable in the sources configured for the `printers` database in the `/etc/nsswitch.conf` file
- The user's `$HOME/.printers` file
- The local `/etc/printers.conf` file
- The LDAP, NIS, or NIS+ name service database

The print client sends its requests to the print server's queue. The print client does not have a local queue. The client writes the print request to a temporary spooling area only if the print server is not available, or if an error occurs. This streamlined path to the server decreases the print client's use of resources, reduces the chances for printing problems, and improves performance.

Printer Configuration Resources

This section describes the resources that the print software uses to locate printer names and printer configuration information.

The print software can use a name service, which is a network (shared) resource for storing information that can be accessed on a network. The name service can be used for storing printer configuration information for all printers on the network. The name service, LDAP, NIS, or NIS+, simplifies printer configuration maintenance. When you add a printer in the name service, all print clients on the network can access the printer.

How the Print Software Locates Printers

The following figure highlights the part of the print process in which the print software checks a hierarchy of printer configuration resources to determine where to send the print request.

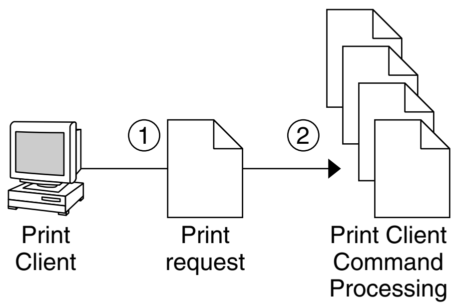

FIGURE 1–3 How the Print Client Software Locates Printers

1. A user submits a print request from a print client by using the `lp` or `lpr` command. The user can specify a destination printer name or class in any of three styles:
   - Atomic style, which is the `lp` command and option, followed by the printer name or class, as shown in this example:

     `% lp -d neptune` *filename*
   - POSIX style, which is the print command and option, followed by *server:printer*, as shown in this example:

     `% lpr -P galaxy:neptune` *filename*
   - Context-based style, as shown in this example:

     `% lpr -d thisdept/service/printer/`*printer-name filename*
2. The print command locates a printer and printer configuration information as follows:
   - The print command checks to see if the user specified a destination printer name or printer class in one of the three valid styles.
   - If the user didn't specify a printer name or class in a valid style, the command checks the user's `PRINTER` or `LPDEST` environment variable for a default printer name.

- If neither environment variable for the default printer is defined, the command checks the sources configured for the `printers` database in the `/etc/nsswitch.conf` file. The name service sources might be one of the following:
  - LDAP directory information tree in the domain's `ou=printers` container
  - NIS `printers.conf.byname` map
  - NIS+ `printers.conf_dir` map

Using Print Servers

The print server is a system that has a local printer connected to it and makes the printer available to other systems on the network. The following figure highlights the part of the print process in which the print server sends the print request to the printer.

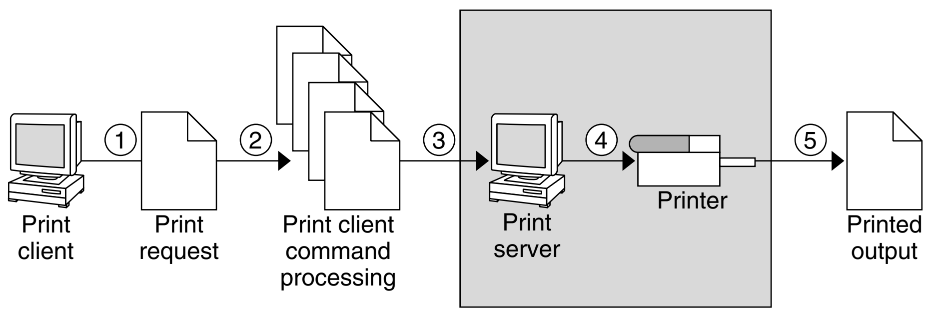

FIGURE 1–4 The Print Server Sends a Print Request to the Printer

1. The print client submits a print request.
2. The print request is processed on the print client.
3. The print request goes to the print server.
4. The print request goes to the printer.
5. The print output is produced on the printer.

The BSD Printing Protocol

The print commands use the BSD printing protocol. One of the big advantages of this protocol is that it can communicate with a variety of print servers:

- `lpd`-based print servers
- SunOS 5.10 and compatible SVR4 (LP) print servers
- Any other print server or printer that accepts the BSD printing protocol

The BSD printing protocol is an industry standard that is widely used because it provides compatibility between different types of systems from various manufacturers. Sun has chosen to support the BSD printing protocol to provide interoperability in the future.

Expanded Printer Support

In this Solaris release, printer support has been expanded to include support for a wide array of printers. Specific modifications have been made to incorporate the RIP feature and PPD files into the Solaris print subsystem. Two new interface scripts, `standard_foomatic`, and `netstandard_foomatic`, were created. These interface scripts provide the generic Solaris interface between the Solaris spooler and the back-end process of the Solaris print server.

The following are examples of the types of printers that are now supported:

- Lexmark Optra E312
- Epson Stylus Photo 1280
- Canon BJC-55
- QMS magicolor 2+

Features of Expanded Printer Support

The following new or modified features are available in this release:

- "Raster Image Processor Support" on page 36
- "Support for PostScript Printer Description Files" on page 36
- "Modifications to Solaris Print Manager Printer Definitions" on page 37
- "New `-n` Option to the `lpadmin` Command" on page 39
- "Changes to `lpstat` Command Output When Using PPD Files" on page 40

Raster Image Processor Support

The RIP enables you to print to printers that do not have resident PostScript processing capabilities. The Solaris printing software now provides the print server RIP and supporting technologies. The RIP occurs behind the scenes. However, to use the appropriate driver you need to configure each printer, by using either Solaris Print Manager or the `lpadmin -n` command. For step-by-step instructions on using the `lpadmin -n` command, see "How to Add a New Attached Printer by Using LP Print Service Commands" on page 65.

Support for PostScript Printer Description Files

The `lpadmin` and `lpstat` commands, as well as the Solaris Print Manager printer definition screens, have been modified to support the use of PPD files.

The following new software packages are associated with this feature:

- `SUNWa2psr`
- `SUNWa2psu`
- `SUNWespgs`
- `SUNWffiltersr`
- `SUNWffiltersu`
- `SUNWfppd`
- `SUNWgimpprint`
- `SUNWhpijs`
- `SUNWimagick`
- `SUNWpsutils`

For more information about supported printers, see "Where to Find Information About Supported Printers" on page 40.

Modifications to Solaris Print Manager Printer Definitions

The following changes have been made to printer definitions in Solaris Print Manager:

- A new Use PPD files option has been added to the Print Manager drop-down menu in Solaris Print Manager. This option is set as the default when creating a new print queue. You can choose to not use PPD files by removing the checkmark from this option. To switch back to using PPD files, recheck the Use PPD files option.

  Note – If Solaris Print Manager cannot locate the `ppdcache` file to use, the Use PPD files option in the Print Manager drop-down menu is greyed out and not available for use.

- The printer type and file content type printer definitions cannot be assigned with Solaris Print Manager if you create a new print queue with PPD files or modify an existing print queue that was created with PPD files. However, If you choose to not use PPD files to create a print queue, you can assign these printer definitions by deselecting the Use PPD files option in the Print Manager drop-down menu in Solaris Print Manager.
- When creating a new print queue or modifying an existing print queue with PPD files, the following Solaris Print Manager screens have been modified:
  - New Attached Printer
  - Modify Printer Properties (for an attached printer)
  - New Network Printer
  - Modify Printer Properties (for a network printer)

  For these screens, when creating a print queue with PPD files, the printer type and file content type printer definitions have been replaced with the following:

  - Printer Make

- Printer Model
- Printer Driver

This table describes the printer definitions that are used by Solaris Print Manager and their availability, depending on which Solaris software release you are running.

Note – The printer make, model, and driver printer definitions are only available when you create a print queue with the Use PPD files default option in the Print Manager drop-down menu selected. Note that this option is the default. Conversely, the printer type and file content type printer definitions are only available when you have deselected the Use PPD files option.

TABLE 1–2 Printer Definitions Used by Solaris Print Manager

Printer Definition	Available in releases prior to Solaris 10	Available in Solaris 10 and later releases
`Printer Name`	Yes	Yes
`Printer Server`	Yes	Yes
`Description`	Yes	Yes
`Printer Port`	Yes, for attached printers	Yes, for attached printers
`Printer Type`	Yes Not by default for the Solaris 9 9/04 release	Yes, by deselecting the Use PPD files option in Print Manager drop-down menu
`File Content Type`	Yes Not by default in the Solaris 9 9/04 release	Yes, by deselecting the Use PPD files option in Print Manager drop-down menu
`Printer Make`	No Yes, available in Solaris the 9 9/04 release only	Yes
`Printer Model`	No Yes, available in the Solaris 9 9/04 release	Yes
`Printer Driver`	No Yes, available in the Solaris 9 9/04 release	Yes
`Fault Notification`	Yes	Yes
`Destination`	Yes, for a network printer	Yes, for a network printer

TABLE 1–2 Printer Definitions Used by Solaris Print Manager *(Continued)*

Printer Definition	Available in releases prior to Solaris 10	Available in Solaris 10 and later releases
`Protocol`	Yes, for a network printer	Yes, for a network printer
`Options`	Yes	Yes
`User Access List`	Yes	Yes
`Default Printer`	Yes	Yes
`Always Print Banner`	Yes	Yes

For a detailed description of the printer make, printer model, and printer driver printer definitions, see "Setting Definitions for Printers" on page 47.

Note – Using PPD files is a per-queue configuration option, and is not a requirement. PPD files are not required for printers that already contain the RIP feature. You can continue to use Solaris Print Manager to configure these printers the same way that you did for previous Solaris releases. To do so, you must remove the checkmark from the Use PPD files default option in the Print Manager drop-down menu in Solaris Print Manager before adding the new printer.

New `-n` Option to the `lpadmin` Command

The `lpadmin` command has a new `-n` option. This option enables you to specify the PPD file that you want to use to create a new print queue or modify an existing print queue.

For example, if you wanted to add the queue, `foobar`, to a Lexmark printer, you would type the following command:

```
# lpadmin -p foobar -v device -I postscript -m standard_foomatic \
-n /path/Lexmark-Optra_E312-Postscript.ppd.gz
```

When using the `-n` option to the `lpadmin` command, specify the full path and file name to the PPD file, as shown in the previous example.

For more information, see the `lpadmin`(1M) man page.

Where to Find Information About Supported Printers

Caution – The location of the PPD files and the `ppdcache` are private interfaces, as is the contents of the `ppdcache` file. The placement of these files and the contents of the `ppdcache` are subject to change in future releases. You should not build software that relies on these files being in their current location or that relies on the data being in the format that is currently used.

If the file that is required by your printer is not available, you can add your own PPD file. If you use the `lpadmin -n` command to create a new print queue, you can store your own PPD files anywhere that you choose. However, if you use Solaris Print Manager to create the print queue, the PPD file must have an entry in the `ppdcache` file.

Directories of PPD files are located in the `/usr/lib/lp/model/ppd/` directory. The output of the `ls` command lists all the PPD files for a particular printer manufacturer. The names of the files reflect the printer model that is supported.

Specifying Your Own PPD Files

When you use the `lpadmin` command with the `-n` option to specify a PPD file, you are required to provide the full path to the PPD file. If you choose to use a PPD file that is located in a different directory, you must specify the full path to that file.

For example, to specify the `xyz.ppd` file that is located in the user's home directory, you would type the following command:

```
# lpadmin -p foo -n /home/user/stuff/xyz.ppd -m \
standard_foomatic -I postscript -v /device
```

Changes to `lpstat` Command Output When Using PPD Files

If a PPD file was used to create a print queue, the PPD file that the specified queue is configured to use is now listed in the `lpstat` command output. If you do not specify a PPD file when creating a new print queue, the `lpstat` command continues to display the same output as it did previously. For an example of the `lpstat` command output when using PPD files, see "How to Check the Status of Printers" on page 87.

You can use Solaris Print Manager to modify queues that were created without the use of PPD files. In such cases, Solaris Print Manager uses the previous printer definitions, which do not include the printer make, model, and driver printer definitions.

For tasks that are associated with expanded printer support, see "Setting Up Printing (Task Map)" on page 57.

Internet Printing Protocol Listener

The Internet Printing Protocol (IPP) listener implements a standardized printing protocol on a network. The IPP listener provides a single, standard interface for querying printing systems attached to a network. The IPP listener for the Solaris OS performs by listening for HTTP requests on port 631. The listener receives print client requests and communicates those requests to the printing system.

To implement the IPP listener on your network, set up a print server on your local system. For more information on setting up a print server, see "Setting Up a Print Server" on page 63. Once the print server has been configured, the IPP listening service automatically starts at system reboot. If you want to stop the service or start the service without rebooting, you can use a run control script. For information on how to use run control scripts, see "How to Use a Run Control Script to Stop or Start a Service" in *System Administration Guide: Basic Administration*.

You also need to configure a print client in order to access the printer system. For example, on a Microsoft Windows system, configure a network printer with the network path set to `http://`*server-name*`:631/printers/`*printer-name*.

The Solaris OS implementation of the IPP listener is in the form of an Apache web server module (`mod-ipp.so`). The default configuration in the Solaris OS enables all of the implemented operations for the IPP listener. To modify the configuration information, edit the `/etc/apache/httpd-standalone-ipp.conf` file. Then stop and restart the IPP service. For information on how to modify the configuration, see `mod-ipp(1).`

CHAPTER 2

Planning Printers on Your Network (Tasks)

This chapter provides information about distributing printers across your network to gain the best efficiency and about planning for printer setup.

This is a list of overview information in this chapter.

Distributing Printers on the Network

The goal of setting up printers on a network is to give users access to one or more printers. As an administrator, you must determine whether each printer would be best used if it is dedicated to one system or available to many systems. In a network environment, distribute your printers on several print servers. The advantage of setting up several print servers is that when one print server has a problem, you can route print requests to other print servers.

If you use a centralized print configuration, you can still connect printers to user systems for convenience or for improved response. A printer that is connected to a user system is still available to other systems on the network.

The following figure shows an example of how you can have a centralized print configuration and still connect printers to users' systems.

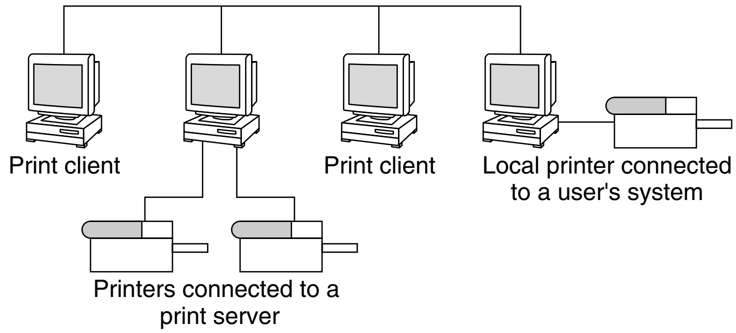

FIGURE 2–1 How to Distribute Printers on a Network

Assigning Print Servers and Print Clients

You must decide which systems will have local printers physically attached to them. You must also decide which will systems use printers on other systems. A system that has a local printer attached to it and makes the printer available to other systems on the network is called a *print server*. A system that sends its print requests to a print server is called a *print client*.

The LP print service software provides printing services in the Solaris Operating System. Besides physically connecting a printer to a system, you must define the printer server characteristics to the LP print service. Once you have print servers set up, you can set up other systems as print clients.

Print servers and print clients can run different versions of the Solaris release and different versions of the UNIX operating system. Print clients running the Solaris 9 release and compatible versions can print to print servers that are running an `lpd`-based print service, such as the SunOS 4.1, BSD UNIX, and Linux releases. In addition, print clients running an `lpd`-based print service can print to print servers running the Solaris 9 release and compatible versions.

The following figure shows a print configuration on a network with systems running the Solaris release and an `lpd`-based print service. For details on how Solaris printing works in heterogeneous environments, see "How Remote Printing Works" on page 189.

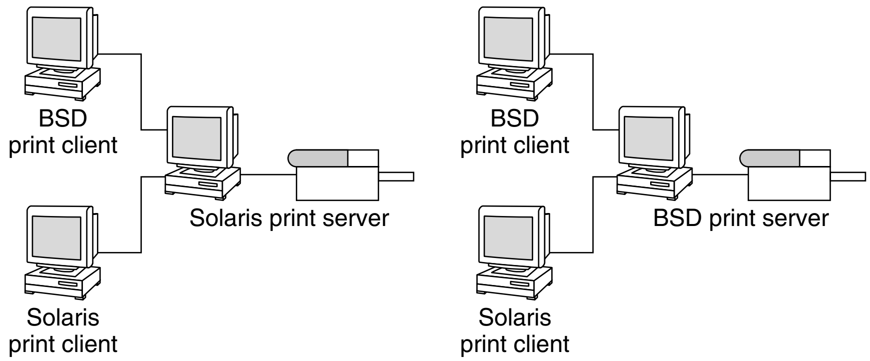

FIGURE 2–2 Solaris and lpd-Based Systems Print Configuration

Print Server Requirements and Recommendations

You can attach a printer to a standalone system or to any system on the network. Any networked system with a printer can be a print server, as long as the system has adequate resources to manage the printing load.

Spooling Space

Spooling space is the amount of disk space that is used to store and process requests in the print queue. Spooling space is the single most important factor to consider when deciding which systems to designate as print servers. When users submit files for printing, the files are stored in the `/var/spool/lp` directory until the files are printed. The size of the `/var` directory depends on the size of the disk and how the disk is partitioned. Spooling space can be allocated in the `/var` directory on the print server, or mounted from a file server and accessed over the network.

Note – If `/var` is not created as a separate file system, the `/var` directory uses space in the root (/) file system. The root (/) file system is likely to be insufficient on a print server.

Disk Space

When evaluating systems as possible print servers, consider their available disk space. A large spool directory can consume 600 Mbytes of disk space. Evaluate the size and division of disk space on systems that can be designated as print servers.

Also, carefully evaluate the printing needs and use patterns of print client systems. For example, assume that users in a small group typically print only short email messages, which are simple ASCII files without sophisticated formatting requirements. In this example, a print server with 20 to 25 Mbytes of disk space allocated to the `/var` directory is probably sufficient. If, however, many print client users are printing large documents, bit mapped images, or raster images, the users will likely fill the spooling space quite frequently. When users cannot queue their jobs for printing, work flow is interrupted. Requests for more spooling space can force you to either add disk space for spooling or designate a different system as the print server.

If the print server has a `/var` directory that is too small, and a larger file system is available, you can mount the larger file system on the `/var` directory for additional spooling space. For information about mounting file systems and editing the `vfstab` file, see "Mounting File Systems" in *System Administration Guide: Devices and File Systems*.

Memory

The Solaris Operating System requires a minimum of 64 Mbytes of memory to run. A print server does not require additional memory. However, you might find that more memory improves performance in managing print requests.

Swap Space

The swap space allocation on the print server should be sufficient to handle LP print service requirements. For information about how to increase swap space, see Chapter 20, "Configuring Additional Swap Space (Tasks)," in *System Administration Guide: Devices and File Systems*.

Hard Disk

For optimal performance, the print server should have a hard disk and a local `/var` directory. You should mount spooling space for a print server on a local hard disk. If a print server has its own hard disk and a local `/var` directory, printing is much faster. And, you can more accurately predict the time needed to process print requests.

Planning for Printer Setup

This section provides an overview of planning for printing that includes the following:

- Setting definitions for printers such as printer name, printer description, printer port
- Selecting a printer make, model, and driver (for setting up printers with PPD files)
- Selecting a printer type and file content type (for setting up printers without PPD files)
- Determining which PPD file to use (for setting up printers with PPD files)
- Setting up a default printer destination
- Determining whether you want to print banner pages
- Limiting user access to printer
- Setting up printer classes
- Setting up fault notification
- Setting up fault recovery

Setting Definitions for Printers

Setting definitions for the printers on your network is an ongoing task that lets you provide a more effective print environment for users. For example, you can assign parameters for all your site's printers to help users find where a printer is located. Or, you can define a class of printers to provide the fastest turnaround for print requests.

The `lpadmin` command lets you set all of the printer definitions. However, Solaris Print Manager lets you set only some printer definitions when you install or modify a printer. The following table lists the printer definitions and shows whether you can assign the definition with Solaris Print Manager.

TABLE 2–1 Printer Definitions Set With Solaris Print Manager

Printer Definition	Can You Set Printer Definitions With Solaris Print Manager?
Printer name	Yes
Printer description	Yes
Printer port	Yes, for adding an attached printer

TABLE 2–1 Printer Definitions Set With Solaris Print Manager *(Continued)*

Printer Definition	Can You Set Printer Definitions With Solaris Print Manager?
Printer type	Yes, if the Use PPD files default option is deselected in the Print Manager drop-down menu
File contents	Yes, if the Use PPD files default option is deselected in the Print Manager drop-down menu
Printer Make	Yes, if the Use PPD files default option in the Print Manager drop-down menu is selected
Printer Model	Yes, if the Use PPD files default option in the Print Manager drop-down menu is selected
Printer Driver	Yes, if the Use PPD files default option in the Print Manager drop-down menu is selected
Fault notification	Yes, but with less functionality than the `lpadmin` command
Default printer	Yes
Destination	Yes, for adding a network printer
Protocol	Yes, for adding a network printer
Printing banner pages	Yes, but with less functionality than the `lpadmin` command
Limiting user access to a printer	Yes, but with less functionality than the `lpadmin` command
Printer class	No
Fault recovery	No

Printer Name

When adding a printer to a system, you specify a *printer name* for the printer. A printer name must adhere to these guidelines:

- The printer name must be unique among all printers within the bounds of an administrative domain.
- The printer name can have a maximum of 14 alphanumeric characters, which can include dashes and underscores.
- The printer name should be easy to remember and might identify the type of printer, its location, or the print server name.

Establish a naming convention that works for your site. For example, if you have different types of printers on the network, including the printer type as part of the printer name can help users choose an appropriate printer. For instance, you could identify PostScript printers with the letters `PS`. However, if all of the printers at your site are PostScript printers, you would not need to include the initials `PS` as part of the printer name.

Printer Description

You can assign a description to a printer by using the `lpadmin -D` command or Solaris Print Manager. The printer's description should contain information to help users identify the printer. You might include the room number where the printer is located, the type of printer, the manufacturer, or the name of the person to call if printing problems occur.

Users can view a printer description by using the following command:

```
$ lpstat -D -p printer-name
```

Printer Port

When you install a printer or later change its setup, you can specify the device, or the *printer port*, to which the printer is connected. You can use either Solaris Print Manager or the `lpadmin -p` *printer-name* `-v` *device-name* command.

Most systems have two serial ports, plus a parallel port or USB ports. Unless you add ports, you cannot directly connect more than two serial printers and a parallel printer, or two USB printers, to one system.

You can select the following printer port types with Solaris Print Manager. These options give you as much flexibility as the `lpadmin` command.

Printer Port Type	Corresponding Device Name Options
Serial	`/dev/term/a`
Serial	`/dev/term/b`
Parallel	`/dev/printers/0` —> `/dev/ecpp0`
USB	`/dev/printers/[1-9]`
Specify any port name that the print server recognizes	`Other`

The LP print service initializes the printer port by using the settings from the standard printer interface program. For more information about printer interface programs, see "Managing Print Filters" on page 126. If you have a parallel printer or a serial printer for which the default settings do not work, see "Adjusting Printer Port Characteristics" on page 148 for information about customizing the port settings.

---

x86 only – If you use multiple ports on an x86 based system, only the first port is enabled by default. The second port, and any subsequent ports, are disabled by default. To use more than one port, you must manually edit the device driver port configuration file for each additional `asy` (serial) port or `ecpp` (parallel) port. The path names for the x86 port configuration files are the following:

`/kernel/drv/asy.conf`

`/kernel/drv/ecpp.conf`

Printer Make

When you set up a printer with PPD files, the *printer make* is the name of the manufacturer of the printer. The printer make is found on the printer itself, and on the packaging materials and documentation that is shipped with the printer.

The following are examples of printer makes that are available in this release:

- Lexmark
- Epson
- Canon
- QMS
- Xerox

Printer Model

The printer manufacturers produce several types and models of printers. When you set up a printer with PPD files, the *printer model* defines the printer precisely. The model is usually stamped on the front or top of the printer. This information is also displayed on the packaging materials and documentation that is shipped with the printer.

The following are examples of printer models that are available in this release:

- Lexmark Optra E312
- Lexmark Z32
- Lexmark 1000

Printer Driver

The *printer driver* is determined when you select the PPD file for the specified printer make and printer model

Printer Type

The *printer type* is a generic name for a type of printer. Printer type identifies the `terminfo` database entry that contains various control sequences for the printer. By convention, printer type is usually derived from the manufacturer's model name. For example, the printer type name for the DECwriter printer is `decwriter`. However, the common printer type `PS` does not follow this convention. `PS` is used as the printer type for many models of PostScript printers, such as the Apple LaserWriterI and Apple LaserWriterII printers.

You can specify the printer type by using the `lpadmin -T` command or Solaris Print Manager.

Printer Names in the `terminfo` Database

Information about each printer type is stored in the `terminfo` database (`/usr/share/lib/terminfo`). This information includes the printer capabilities and initialization control data. The printer you install must correspond to an entry in the `terminfo` database.

```
$ pwd
/usr/share/lib/terminfo
$ ls
1  3  5  7  9  B  H  P  a  c  e  g  i  k  m  o  q  s  u  w  y
2  4  6  8  A  G  M  S  b  d  f  h  j  l  n  p  r  t  v  x  z
$
```

Each subdirectory contains compiled database entries for terminals or printers. The entries are organized by the first letter of the printer or terminal type. For example, if you have an Epson printer, look in the `/usr/share/lib/terminfo/e` directory to find your particular model of Epson printer.

```
$ cd /usr/share/lib/terminfo/e
$ ls
emots          ep2500+high   ep48            ergo4000      exidy2500
env230         ep2500+low    epson2500       esprit
envision230    ep40          epson2500-80    ethernet
ep2500+basic   ep4000        epson2500-hi    ex3000
ep2500+color   ep4080        epson2500-hi80  exidy
$
```

The entries for Epson printers begin with `epson`.

If you have an NEC printer, look in the `/usr/share/lib/terminfo/n` directory for your NEC printer model.

```
$ cd /usr/share/lib/terminfo/n
$ ls
ncr7900         ncr7900iv       netronics       network         nuc
ncr7900-na      ncr7901         netty           netx            nucterm
ncr7900i        nec             netty-Tabs      newhp
ncr7900i-na     net             netty-vi        newhpkeyboard
$
```

The nec entry in this directory is for the NEC printer.

Selecting a Printer Type

In this Solaris release, you can assign a printer type with Solaris Print Manager if you deselect the Use PPD files default option in the Print Manager drop-down menu in Solaris Print Manager. This printer definition is not available when the Use PPD files option is selected. When you add a printer with PPD files, the printer make, model, and driver printer definitions replace the printer type definition.

For a local PostScript printer, use a printer type of either PostScript (PS) or Reverse PostScript (PSR). If your printer supports PostScript, choose PS or PSR, even if the specific printer type is listed in the terminfo database.

If your PostScript printer prints pages face up, documents appear to be printed backwards. As such, the first page is at the bottom of the stack and the last page is on the top. If you specify the printer's type as PSR, the LP print service reverses the order of the pages before sending the pages to the printer. The last page is printed first, and the pages are stacked in forward order. However, the LP print service can reliably change the page order only for PostScript files that conform to the Adobe® Technical Note #5001, *PostScript Language Document Structuring Conventions Specification*. This document is available on the Adobe Developer Relations web site.

If a printer can emulate more than one kind of printer, you can assign several printer types by using the lpadmin -T command. If you specify more than one printer type, the LP print service uses the printer type that is appropriate for each print request.

You might not find the printer type in the appropriate terminfo directory. The type of printer is not necessarily linked to the manufacturer's name on the printer. For example, for any type of PostScript printer, you can use the PS or PSR entry, found in the /usr/share/lib/terminfo/P directory, instead of an entry specific to the manufacturer or product name.

If you have an unusual type of printer, you might need to try different entries before you can determine whether a particular terminfo entry works for your type of printer. If possible, find an entry in the terminfo database that works for your printer. Trying to use an existing entry is much easier than trying to create an entry. If you must create your own entry, "Adding a terminfo Entry for an Unsupported Printer" on page 150 contains some useful tips.

Selecting a File Content Type

In this Solaris release, you can assign the file content type printer definition with Solaris Print Manager if you deselect the Use PPD files default option in the Print Manager drop-down menu in Solaris Print Manager. This printer definition is not available if you use PPD files. The file content type for printers added with PPD files is PostScript. This file content type is automatically set by the tool. You can also specify file content type when adding a new printer or modifying an existing printer by using the `lpadmin -I` command.

Print filters convert the content type of a file to a content type that is acceptable to the destination printer. The *file content type* tells the LP print service the type of file contents that can be printed directly, without filtering. To print without filtering, the necessary fonts must also be available in the printer. You must set up and use filtering for other types of files.

Most printers can print directly the following types of files:

- The same type as the printer type. For example, `PS` for a PostScript printer.
- The `simple` type. For example, an ASCII text file.

When submitting a file for printing, the user can indicate the content type of the file by using the `lp -T` *content-type* command. If no file content type is supplied when the request is submitted, the LP print service checks the first file in the request to determine the content type. If the file begins with `^D%!` or `%!`, the request is considered to contain PostScript data. Otherwise, the request is assumed to contain `simple` (ASCII) text. The LP print service uses the file content type to determine which filters to use to convert the file contents into a type the printer can handle.

When you are not specifying PPD files, Solaris Print Manager provides a list of file content types from which you can choose when you install or modify an attached or network printer. The choices are translated to the names that the LP print service uses. The following table describes the file content types that you can choose with Solaris Print Manager.

TABLE 2–2 Choosing File Content Type With Solaris Print Manager When Not Using PPD Files

File Contents Choice	LP Print Service Name	Description
PostScript	`postscript`	PostScript files do not require filtering.
ASCII	`simple`	ASCII files do not require filtering.
Both PostScript and ASCII	`simple, postscript`	PostScript files and ASCII files do not require filtering.
None	`""`	All files require filtering, except those files that match the printer's type.

TABLE 2–2 Choosing File Content Type With Solaris Print Manager When Not Using PPD Files *(Continued)*

File Contents Choice	LP Print Service Name	Description
Any	`any`	No filtering is required. If the printer cannot handle a file content type directly, the file will not be printed.

Choose the file content type that best matches the printer's capabilities. PostScript is the default choice in Solaris Print Manager and is probably correct most of the time. PostScript means that filtering is not needed for PostScript files.

Frequently Used Printers

This section provides the printer type and file content type for the printers most commonly used with Solaris software. Although not shown, many of these printers can also directly print files with `simple` file content type.

If you have a PostScript printer, use a printer type of `PS` or `PSR` and a content type of `postscript`. `PSR` reverses the pagination and prints the banner page last.

The following table lists additional non-PostScript printers and shows the printer type to use for configuring each printer. For all these printers, the file content type is `simple`.

Note – Sun Microsystems does not supply filtering software for the printers listed in the following table, among others. However, you can use unsupported printers if you supply filtering or if the printer can directly print the file content type. If you have questions about any printer for which Sun Microsystems does not supply filters, contact the printer manufacturer.

TABLE 2–3 Some Non-PostScript Printers for Which Sun Does Not Supply Filters

Printer	Printer Type
Daisy	`daisy`
Datagraphix	`datagraphix`
DEC LA100	`la100`
DEC LN03	`ln03`
DECwriter	`decwriter`
Diablo	`diablo`

TABLE 2–3 Some Non-PostScript Printers for Which Sun Does Not Supply Filters *(Continued)*

Printer	Printer Type
	`diablo-m8`
Epson 2500 variations	`epson2500`
	`epson2500-80`
	`epson2500-hi`
	`epson2500-hi80`
Hewlett-Packard HPCL printer	`hplaser`
IBM Proprinter	`ibmproprinter`

If you want to set up a printer that is not in the `terminfo` database, see "How to Add a `terminfo` Entry for an Unsupported Printer" on page 152.

CHAPTER 3

Setting Up Printers (Tasks)

This chapter describes how to set up a printer and how to make the printer accessible to systems on the network by using Solaris Print Manager and the LP print service.

For information on the procedures associated with setting up printers by using Solaris Print Manager, see "Setting Up Printing (Task Map)" on page 57.

For overview information about printing, see Chapter 1.

Setting Up Printing (Task Map)

Task	Description	For Instructions
Add a new attached printer by using Solaris Print Manager.	Use Solaris Print Manager to make the printer available for printing after you physically attach the printer to a system.	"How to Add a New Attached Printer by Using Solaris Print Manager" on page 64
Add a new attached printer by using LP print service commands.	Use LP print service commands to make the printer available for printing after you physically attach the printer to a system.	"How to Add a New Attached Printer by Using LP Print Service Commands" on page 65
Add access to a printer.	Use Solaris Print Manager to add printer access on the print client.	"How to Add Printer Access by Using Solaris Print Manager" on page 69

Task	Description	For Instructions
(Optional) Set up a `.printers` file.	Use a `$HOME/.printers` file so that users can establish their own custom printer aliases.	"How to Set Up a `.printers` File" on page 70
Add a new network printer by using vendor-supplied tools.	Use Printer Vendor Supplied Tools to configure the network printer after you physically connect the printer to the network. After you physically connect the printer to the network, configure the network printer.	"How to Add a Network Printer by Using Printer Vendor-Supplied Tools" on page 76
Add a new network printer by using Solaris Print Manager.	Use Solaris Print Manager to make a printer available for printing over the network.	"How to Add a New Attached Printer by Using Solaris Print Manager" on page 64
Add a new network printer by using LP print service commands.	Use LP print service commands to make a printer available for printing over the network.	"How to Add a New Network Printer by Using LP Print Service Commands" on page 78
Turn off banner pages.	You can optionally turn off banner pages so that the banner pages are never printed.	"How to Turn Off Banner Pages" on page 95
Set up fault alerts.	You can optionally set up more specific fault alerts for the printer than the fault alerts provided by Solaris Print Manager.	"How to Set Fault Alerts for a Printer" on page 98
Set up fault recovery.	You can optionally set up how a printer should recover after it faults.	"How to Set Printer Fault Recovery" on page 100
Limit access to a printer.	Solaris Print Manager enables you to optionally set up an allow list. If you want to limit the access of a few users to the printer, you might want to set up a deny list.	"How to Limit User Access to a Printer" on page 102

Setting Up Printing by Using Solaris Print Manager

Solaris Print Manager is a Java-based graphical user interface that enables you to manage local and remote printer configuration. This tool can be used in the following name service environments: LDAP, NIS, NIS+, and files. You must be logged in as superuser to use this tool.

You can use Solaris Print Manager to set up print servers (Add New Attached Printer or Add a Network Printer) and print clients (Add Access to Printer). An attached printer is a printer that is physically cabled to the print server. A network printer is physically attached to the network. You can use Solaris Print Manager to add an attached printer or a network printer. Adding access to a printer, or adding remote access, is the process of giving print clients (all those machines which are not the server) access to the printer.

Printer Setup With Solaris Print Manager

The following table describes each printer attribute to help you determine the information that is needed to set up a printer with Solaris Print Manager.

Printer Definition	Description	Example	Default Setting	Required or Optional?
`Printer Name`	Name of printer	`laser1`	N/A	Required to install an attached or network printer and to add access to a printer
`Printer server`	Name of print server	`venus`	N/A	Required to install an attached or network printer and to add access to a printer — This field is filled in by the tool
`Description`	User defined string	`laser printer near breakroom`	N/A	Optional
`Printer Port`	Device that the printer is attached to	`/dev/term/a`	`/dev/term/a`	Required to install an attached printer
`Printer Type`	Type of printer	`unknown`	`PostScript`	Required to install an attached printer or a network printer without PPD files

Printer Definition	Description	Example	Default Setting	Required or Optional?
File Content Type	Content to be printed	any	PostScript	Required to install an attached printer or a network printer without PPD files
Printer Make	Make of printer	Lexmark	N/A	Required to install an attached printer or a network printer with PPD files
Printer Model	Model of printer	Lexmark Optra E312	N/A	Required to install an attached printer or a network printer with PPD files
Printer Driver	Driver that is used by specifying PPD file	Foomatic/PostScript	N/A	Required to install an attached printer or a network printer with PPD files
Destination	Destination name for the network printer	For examples, see "Selecting the Destination Name (Also Called the Network Printer Access Name)" on page 74	N/A	Required to install a network printer
Protocol	Protocol used to communicate with the printer	TCP	BSD	Required to install a network printer
Fault Notification	Specifies how to notify user of errors	Mail to superuser	Write to superuser	Optional
Default Printer	Identifies the default printer	N/A	N/A	Optional
Always Print Banner	Print banner with print job?	N/A	Banner is printed	Optional
User Access List	List of users allowed to print on the print server	rimmer,lister	All users can print	Optional

Solaris Print Manager and Special Characters

Solaris Print Manager checks user input for the various text fields in the input screens. There are two types of checking: general illegal input and input that is illegal for specific fields.

Solaris Print Manager does not accept the following characters as input, except for the help screens:

- Shell metacharacters, such as “\$^&*(){}‘‘|;:?<>, except for the destination field on the network printer screen, which accepts colons (:)
- Multibyte characters
- Pound signs (#), spaces, or tabs, except the description field, which accepts tabs

Starting Solaris Print Manager

To start Solaris Print Manager, either select Printer Administrator from the CDE Workspace menu or start the Print Manager from the command line. See the following procedure for details.

▼ How to Start Solaris Print Manager

Steps

1. **Verify that the following prerequisites are met. To use Solaris Print Manager, you must meet the following requirements:**
   - Have a bitmapped display monitor, such as the standard display monitor that comes with a Sun workstation. Solaris Print Manager can be used only on a system with a console that is a bitmapped screen.
   - Be running an X Window System, such as the CDE environment, or be using the remote display feature on a system running an `xhost` environment.
   - Be logged in as superuser or assume an equivalent role on the print server to install an attached or network printer, or on the print client to add access to a printer.
   - Have the required access privileges for managing the LDAP, NIS, or NIS+ database:
     - If your name service is LDAP, you must have the following:
       - The distinguished name (DN) of a printer administrator and password in the directory.
       - The name or IP address of an LDAP server hosting the administered domain. This information is usually displayed and selected automatically.
     - If your name service is NIS, you must have the root password for the NIS master.
     - If you name service is NIS+, you might need to do the following:

a. Log in to the NIS+ master as superuser.

b. Identify the group that owns the printers table:

```
# niscat -o printers.org_dir.domain-name.com
          .
          .
          .
         Group          : "admin.domain-name.com"
```

c. If necessary, add the system that runs Solaris Print Manager to the NIS+ admin group authorized to update the `printers.org_dir.`<*domain*> file.

```
# nisgrpadm -a admin.domain-name.com hostname
```

d. As superuser, log in to the system that runs Solaris Print Manager. Your NIS+ configuration might make it necessary to run the `/usr/bin/keylogin` command. For more information, see the `keylogin`(1) man page.

- Have the `SUNWppm` package installed.

```
# pkginfo | grep SUNWppm
system      SUNWppm        Solaris Print Manager
```

2. **Start Solaris Print Manager by using one of the following methods:**

- Select Printer Administrator from the Tools option of the CDE Workspace menu.
- Select the Applications menu from the CDE front panel. Then, click the Printer Administrator icon in the Application Manager's System Admin window.
- Type the following command:

```
# /usr/sadm/admin/bin/printmgr &
```

The Select Naming Service window overlays the Solaris Print Manager main window.

- If you want to use Solaris Print Manager from a remote system, do the following:

Use the `xhost` command on the local system to give the remote system display access:

```
# xhost +remote-system
```

Then log in to the remote system, set the `DISPLAY` environment variable, and start Solaris Print Manager:

```
# DISPLAY=local-system:display-number
# export DISPLAY
# /usr/sadm/admin/bin/printmgr &
```

Note – If Solaris Print Manager fails to start from the CDE menu or from the command line, check the following:

- Superuser (`root`) might not have permission to connect to the X-server process on the local system or the remote system. In this case, type the following:

  ```
  $ xhost +hostname
  $ su
  (Enter root's password)
  # /usr/sadm/admin/bin/printmgr &
  ```

  Replace *hostname* with either the local system name or the remote system name before restarting Solaris Print Manager.

- Verify that the `SUNWppm` package is installed on the local system or the remote system.

  ```
  $ pkginfo | grep SUNWppm
  ```

3. **Select the name service that is used in your network from the Select Naming Service window. The choices are: LDAP, NIS, NIS+, or files.**
4. **Check that the domain name is correct.**

   The Solaris Print Manager main menu is displayed after the name service is loaded successfully.

Setting Up a Print Server

When you add an attached printer or a network printer to a system, the printer is made accessible to the local system. The system on which you install the printer becomes the *print server*.

The following procedure describes how to use Solaris Print Manager to add a new attached printer with the expanded printer support that is the default in this release.

For overview information about the new or changed Solaris Print Manager features in this release, see "What's New in Printing?" on page 25.

▼ How to Add a New Attached Printer by Using Solaris Print Manager

In this release, the printer definitions that you can assign with Solaris Print Manager have been expanded to enable you to use PPD files. A new Use PPD files option has been added to the Print Manager menu drop-down menu in Solaris Print Manager. This default option enables you to select the printer make, model, and driver when adding new printer.

This procedure shows you how to add a new attached printer by using Solaris Print Manager with the Use PPD files default option selected. To add a new attached printer without PPD files, you must first deselect this option in the Print Manager drop-down menu. For overview information about new or modified Solaris Print Manager features in this release, see "What's New in Printing?" on page 25.

Steps

1. **Select the system to be the print server.**

2. **Connect the printer to the print server. Turn on the power to the printer.**

   Consult the printer vendor's installation documentation for information about the hardware switches and cabling requirements.

3. **Start Solaris Print Manager on the print server where you connected the printer.**

   For instructions, see "How to Start Solaris Print Manager" on page 61.

4. **Select New Attached Printer from the Printer menu.**

   The New Attached Printer window is displayed.

5. **Fill in the required fields with the appropriate information.**

   If you need information to complete a field, click the Help button.

   For example, to add a new attached printer with PPD files, complete the following fields:

   ```
   Printer Name:
   Description:
   Printer Port:
   Printer Make:
   Printer Model:
   Printer Driver:
   Fault Notification:
   Options:
   Default Printer
   Always Print Banner
   User Access List:
   ```

   The Printer Server field is filled in by Solaris Print Manager.

   To add a new attached printer without PPD files, complete the following fields:

   ```
   Printer Name:
   Description:
   ```

```
Printer Port:
Printer Type:
File Contents:
Fault Notification:
Options:
Default Printer
Always Print Banner
User Access List:
```

The Printer Server field is filled in by Solaris Print Manager.

6. **Click OK.**

7. **Verify that the printer has been installed. Check for the new printer entry in the Solaris Print Manager main window.**

8. **Verify that the printer can print requests.**

   $ **lp -d** *printer-name filename*

9. **Exit Solaris Print Manager.**

   Choose Exit from the Print Manager Menu.

More Information

Verifying That the New Solaris Print Manager Feature Is Working

When you are using Solaris Print Manager to create or modify a print queue, choose the `files` option as the name service. Select the Printer → New Attached Printer option from the menu. If the drop-down menu displays printer make and printer model, the new feature is working.

Verifying That the RIP Feature Is Working

The RIP enables you to print to printers that do not have resident PostScript processing capabilities.

To ensure that the RIP feature is working you need to create a new print queue by selecting one of the make/model combinations that is available in the New Attached Printer and New Network Printer screens used by Solaris Print Manager. Then, try printing to the new printer. If the output of the print job is recognizable, the RIP feature is working. If the output of the print job is unrecognizable, then the RIP feature is not working.

▼ How to Add a New Attached Printer by Using LP Print Service Commands

This procedure describes how to add a new attached printer with the expanded printer support that is available in this release. The `-n` option to the `lpadmin` command enables you to specify a PPD file when adding a new attached printer or modifying an existing attached printer. Two examples follow this procedure. The first example shows you how to add a new attached printer with PPD files. The second example shows you how to add a new attached printer without PPD files.

Steps

1. **Connect the printer to the system, and turn on the power to the printer.**

   Consult the printer vendor's installation documentation for information about the hardware switches and cabling requirements.

2. **Collect the information that is required to configure an attached printer with PPD files.**

   - Printer name and port device
   - File content type
   - PPD file

     To determine the PPD file that the printer will use, first define the printer make, model, and driver.

     For overview information about using PPD files when adding a printer or modifying an existing printer by using LP print service commands, see "What's New in Printing?" on page 25.

3. **Define the printer name, port device, file content type, and PPD file that the printer will use.**

   When using PPD files, the file content type is typically PostScript.

   a. **Specify the printer name and the port device that the printer will use.**

   ```
   # lpadmin -p printer-name -v /dev/printers/0
   ```

   The device to use is `/dev/printers/0`.

   b. **Specify the interface script that the printer will use.**

   ```
   # lpadmin -p printer-name -m standard_foomatic
   ```

   c. **Specify the file content type and the PPD file that the printer will use.**

   ```
   # lpadmin -p printer-name -I content-type -n /path/ppdfile
   ```

   The PPD file that you supply is located in the /*path* directory.

   d. **Specify the printer description.**

   ```
   # lpadmin -p printer-name -D "printer-description"
   ```

   For more information, see the `lpadmin`(1M) man page.

4. **(Optional) Add filters to the print server.**

   ```
   # cd /etc/lp/fd
   # for filter in *.fd;do
       > name=`basename $filter .fd`
       > lpfilter -f $ name-F $ filter
       > done
   ```

5. **Enable the printer to accept print requests and to print those requests.**

   ```
   # accept printer-name
   # enable printer-name
   ```

6. **Verify that the printer is correctly configured.**

   ```
   # lpstat -p printer-name -l
   ```

   The output of the `lpstat` command will list the PPD file that you used.

Example 3–1 Adding a New Attached Printer With PPD Files by Using LP Print Service Commands

This example shows how to add a new attached printer with PPD files. The `-n` option to the `lpadmin` command enables you to add a new print queue by specifying PPD files. The following information is used as an example. The information that you provide will vary:

- Printer name: `paper`
- Port device: `/dev/printers/0`
- File content type: `postscript`
- PPD file: `/path/Lexmark-Optra_E312-Postscript.ppd.gz`

```
# lpadmin -p paper -v /dev/printers/0     (1)
# lpadmin -p paper -I postscript     (2)
# lpadmin -p paper -m standard_footmatic     (3)
# lpadmin -p paper -n /path/Lexmark-Optra_E312-Postscript.ppd.gzmake     (4)
# lpadmin -p paper -D "Color printer on third floor, rm 3003"     (5)
# cd /etc/lp/fd
# for filter in *.fd;do
   > name=`basename $ filter .fd`
   > lpfilter -f $ name -F $ filter
   > done     (6)
# accept paper
    destination "paper" now accepting requests
# enable paper        (7)
printer "paper" now enabled
# lpstat -p paper     (8)
printer paper is idle. enabled since Feb 28 11:21 2004...
```

1. Defines the printer name and the port device that the printer will use
2. Sets the file content types to which the printer can print to directly
3. Specifies the interface script for utilizing PPD files
4. Specifies the path to the PPD file and the PPD file name
5. Adds a description for the printer
6. Adds print filters to the print server
7. Accepts print requests for the printer and enables the printer
8. Verifies that the printer is ready for printing

Example 3–2 Adding a New Attached Printer Without PPD Files by Using LP Print Service Commands

This example shows how to add a new attached printer without PPD files. The commands in this example must be executed on the print server where the printer is connected. The following information is used as an example. The information that you provide will vary:

- Printer name: `luna`
- Port device: `/dev/term/b`
- Printer type: `PS`
- File content type: `postscript`

```
# chown lp /dev/term/b
# chmod 600 /dev/term/b     (1)
# lpadmin -p luna -v /dev/term/b     (2)
# lpadmin -p luna -T PS     (3)
# lpadmin -p luna -I postscript     (4)
# lpadmin -p luna -D "Room 1954 ps"     (5)
# cd /etc/lp/fd
# for filter in *.fd;do
   > name =`basename $ filter.fd`
   > lpfilter -f $ name -F $ filter
   > done     (6)
# accept luna
destination "luna" now accepting requests
# enable luna     (7)
printer "luna" now enabled
# lpstat -p luna     (8)
printer luna is idle.enabled since Feb 28 11:21 2003.available
```

1. Gives `lp` ownership and sole access to a port device
2. Defines the printer name and the port device that the printer will use
3. Sets the printer type of the printer
4. Specifies the file content types to which the printer can print directly
5. Adds a description for the printer
6. Adds print filters to the print server
7. Accepts print requests for the printer and enables the printer
8. Verifies that the printer is ready for printing

See Also

- "How to Add Printer Access by Using Solaris Print Manager" on page 69 to add printer access on a print client.
- "How to Set Up a `.printers` File" on page 70 to allow users to customize printer aliases.

Setting Up a Print Client

A *print client* is a system that is not the server for the printer. Yet, this system has access to the printer. A print client uses the services of the print server to spool, schedule, and filter the print jobs. Note that one system can be a print server for one printer and be a print client for another printer.

Access to a printer can be configured on a domain-wide basis or on a per-machine basis. If you add the printer information to the name service database, access is configured on a domain-wide basis.

The following procedure describes how to use the Solaris Print Manager to add access to a printer on a print client. The example that follows describes how to add printer access by using the `lp` print service commands.

▼ How to Add Printer Access by Using Solaris Print Manager

Steps

1. **Start Solaris Print Manager on the system where you want to add access to a remote printer.**

   For instructions, see "How to Start Solaris Print Manager" on page 61.

2. **Select Add Access to Printer from the Printer menu.**

   The Add Access to Printer window is displayed.

3. **Type the appropriate data in the entry fields.**

   If you need information to complete a field, click the Help button.

4. **Click OK.**

5. **Verify that access to the printer is added by checking for the new printer entry in the Solaris Print Manager main window.**

6. **Verify that the printer can print requests.**

   ```
   $ lp -d printer-name filename
   ```

7. **Exit Solaris Print Manager.**

   Choose Exit from the Print Manager Menu.

Example 3–3 Adding Printer Access by Using `lp` Print Service Commands

If you want to print to a remote printer, you must add access to the remote printer. This example shows how to configure access to a printer named `luna`, whose print server is `saturn`. The system `saturn` becomes a print client of the printer `luna`.

```
# lpadmin -p luna -s saturn     (1)
# lpadmin -p luna -D "Room 1954 ps"     (2)
# lpadmin -d luna     (3)
# lpstat -p luna     (4)
```

1. Identifies the printer and the print server
2. Adds a description for the printer
3. Sets the printer as the system's default printer destination
4. Verifies that the printer is ready for printing

Setting Up a `.printers` File

There is no need to set up a `.printers` file in your users' home directories if the users don't need customized printer information. However, the `.printers` file enables users to establish their own printer aliases. You can use the `_default` alias to make a printer the default printer. You can also set up a special `_all` alias to define a list of printers affected when you cancel a print request or to check the status of printers.

The use of the `.printers` file by the LP print service is controlled by the name service switch (`/etc/nsswitch.conf`). The default configuration specifies that the print service checks a user's home directory to locate printer configuration information before it checks the other name services. So, you can tailor a user's printer configuration file to use custom printer information rather than the shared information in the name service.

For more information about the `.printers` file, see the`printers`(4) man page. For more information about the name service switch, see the`nsswitch.conf`(4) man page.

▼ How to Set Up a `.printers` File

Steps

1. **Log in to the user's system as superuser or assume an equivalent role.**
2. **Start the file editor that you want to use to create a `.printers` file in the user's home directory.**
3. **(Optional) Set up the `_default` alias to make a specific printer the user's default printer. Use an entry that is similar to this entry:**

_default *printer-name*

4. **(Optional) Set up the `_all` alias to define the printers affected when you cancel a print request or check the status of printers. Use an entry similar to this entry:**

   _all *printer1,printer2,printer3*

5. **Save the file as `.printers`.**

Introduction to Adding a Network Printer

A *network printer* is a hardware device that is directly connected to the network. This printer can be accessed from a print server without actually connecting the printer to the print server with a cable. The network printer has its own system name and IP address.

Caution – Configuring a network printer without utilizing queue configuration on a print server is an unsupported method that results in printer misconfiguration.

You can correctly configure a network printer with Solaris Print Manager by selecting the New Network Printer option. Or, you can correctly configure a network printer by using the `lpadmin` command. For more information on how to configure a network printer by using the `lpadmin` command, see "How to Add a New Network Printer by Using LP Print Service Commands" on page 78.

Note that you should not use the `lpadmin -p` *queue-name* `-s` *printer-name* command or the Add Access to Printer option of Solaris Print Manager to add a network printer. These methods should be used exclusively to point to an already existing queue on a remote print server.

The print server provides the following:

- Queueing capabilities
- Filtering
- Printing administration for a network printer

Printing directly to a network printer results in a printer misconfiguration that is likely to cause a number of problems. Additionally, job options, such as copies, the use or disuse of burst page, and filtering, are lost.

Network printers might use one or more special protocols that require a vendor-supplied printing program. The procedures used to set up the vendor-supplied printing program can vary. If the printer does not come with vendor–supplied support, Solaris network printer support can be used with most devices. Use the printer vendor-supplied software, when possible.

The vendor might supply an SVR4 printer interface script to replace the standard printer interface script. If so, that SVR4 interface script calls the vendor-supplied printing program to send the job to the printer. If not, you need to modify the standard interface script to call the vendor-supplied printing program. To do so, edit the per-printer copy of the standard interface script to call the vendor-supplied printing program.

The following terms are used in network printer configuration:

- **Print server** – The system that spools and schedules the jobs for a printer. The print server is the system on which the printer is configured.
- **Printer-host device** – The software and hardware supplied by a vendor that provides network printer support for a non-network capable printer. The combination of the printer-host device with one or more printers attached to the device creates a *network printer*.
- **Printer node** – Is either the physical printer or the printer-host device. The printer node is the physical printer when the network support resides in the physical printer. The printer node is the printer-host device when an external box is used to provide the network interface. The printer node name is the system name given with the IP address. You select this name, which has no default or vendor requirement. The printer node name, as with all nodes, must be unique.
- **Printer name** – The name typed on the command line when you use any of the printer commands. You select the printer name at the time of printer configuration. Any one physical printer can have several printer name or queue names. Each name provides access to the printer.
- **Destination or network printer access name** – The internal name of the printer node port that is used by the printer subsystem to access the printer. The access name is the name of the printer node, or the name of the printer node with a printer vendor port designation. Any printer vendor port designation is explicitly defined in the printer vendor documentation. The designation is printer specific. If the printer is a printer-host device and a printer, the port designation is documented in the printer-host device documentation. The format is:

  printer-node-name

  or

  printer-node-ame:*port-designation*
- **Protocol** – The over-the-wire protocol used to communicate with the printer. The documentation from the printer vendor supplies the information regarding the protocol to select. The network printer support supplies both the BSD print protocol and raw TCP. Due to implementation variations, you might want to try both protocols.

- **Timeout, or retry interval timeout** – Is a seed number that represents the number of seconds to wait between attempting connections to the printer. This seed number is the smallest amount of time to wait between attempted connections. The seed number increases when there is an increase in failed connections. After repeated failures to connect to the printer, a message is returned to the user requesting possible human intervention. Attempts to reconnect continue until the connection is successful or the job is cancelled by the job owner.

Printer Vendor-Supplied Software for Network Printers

Network printers often have software support provided by the printer vendor. If your printer has printer vendor- supplied software, use the printer vendor software, when possible. The software is designed to support the attributes of the printer and can take full advantage of the printer capabilities. Read the printer vendor documentation to install and configure the printer, so that it can operate with the LP print service.

Sun Support for Network Printers

If the network printer vendor does not provide software support, Sun supplied software is available. The software provides generic support for network printers and is not capable of providing full access to all possible printer attributes.

A general discussion of how to add a network printer is provided in Chapter 3. The following is a discussion of printer management when you use the Sun supplied software.

Invoking Network Printer Support

The software support for network printers is called through the network interface script `netstandard`. Configuring a network printer with this script causes the network printer support module to be called. Here is the command to configure the printer with the network support.

```
lpadmin -p printer-name -m netstandard
```

Selecting the Print Protocol

The print subsystem uses the BSD print protocol and raw TCP to communicate with the printer. The printer vendor documentation provides information about which protocol to use. In general, the TCP protocol is more generic across printers.

Type the following command to select the protocol:

```
lpadmin -p printer-name -o protocol=bsd
```

or

```
lpadmin -p printer-name -o protocol=tcp
```

If the protocol that is selected is the BSD print protocol, you can further select the order of sending the control file to the printer. Some printers expect the control file, then the data file. Other printers expect the reverse. For this information, see the printer vendor documentation. The default is to send the control file first.

Type the following command to select the order:

```
lpadmin -p printer-name -o bsdctrl=first
```

or

```
lpadmin -p printer-name -o bsdctrl=last
```

Selecting the Printer Node Name

You select the printer node name. This name must be unique, as with any node on the network. The printer node name is associated with the IP address of the printer.

Selecting the Destination Name (Also Called the Network Printer Access Name)

The print subsystem requires access information for the printer. The subsystem uses the destination name when making the network connection to the printer. You supply this name to the print subsystem by using the `lpadmin` command. This name then becomes part of the printer configuration database. The printer access name is the name of the printer node, sometimes qualified by a port name. Port designation varies across printer vendors. You can find information about port designation in the printer documentation. Here is the format of printer access name:

printer-node-name[:*port-designation*]

EXAMPLE 3–4 Destination Name (or Network Printer Access Name) With Port Designation (Number)

A common port designation with TCP is 9100. If the printer node name is `pn1`, and the printer vendor defines the port as 9100, then the printer access name is `pn1:9100`. To configure a printer in this case, use the following command:

```
lpadmin -p printer_name -o dest=pn1:9100
```

EXAMPLE 3–5 Destination Name (or Network Printer Access Name) With BSD Port Designation (Name)

When you use the BSD protocol, the port designation might not be a number, but some name defined by the printer vendor, for example: *xxx*`_parallel_1`. If the printer node name is *cardboard*, then the printer access name is `cardboard:xxx_parallel_1`. To configure a printer in this case, use the following command:

```
lpadmin -p printer-name -o dest=cardboard:xxx_parallel_1
```

EXAMPLE 3–6 Destination Name (or Network Printer Access Name) With No Port Designation

If there is no port designation, and the printer node name is `newspaper`, the printer access name is the printer node name: `newspaper`. To configure a printer in this case, use the following command:

```
lpadmin -p printer-name -o dest=newspaper
```

Setting the Timeout Value

The timeout option allows the user to select the amount of time (in seconds) to wait between successive attempts to connect to the printer. Some printers have a long warm up time, and a longer timeout value is advised. The default is 10 seconds.

The timeout value does not impact the success or failure of the print process. The value is a seed value that the software uses as the initial timeout count. On repeated failures, this count is increased. A message is sent to the spooler when repeated attempts to connect to the printer fail. This message alerts the user that intervention might be required. The failure could be anything from the printer being turned off to the printer being out of paper. If these messages are produced too often, for example when the printer is warming up, increasing the timeout value will eliminate spurious messages.

You ncan experiment to find the optimal timeout value. Type the following command to set the timeout value:

```
lpadmin -p printer-name -o timeout=n
```

Managing Network Printer Access

Each network printer should have only one server that provides access to the printer. This restriction enables the server to manage access to the printer and keep jobs coherent.

The default device for the network printer is `/dev/null`. This device is sufficient when the printer has only one queue. If more queues are required, set the device to a file. This setting enables the print system to restrict access to the printer across queues. The following commands create a device file and configure it as the network printer device.

```
touch /path/filename
chmod 600 /path/filename
lpadmin -p printer-name -v /path/filename
```

The following is an example of how to create a device file called `devtreedown`.

```
# touch /var/tmp/devtreedown
# chmod 600 /var/tmp/devtreedown
# lpadmin -p treedown -v /var/tmp/devtreedown
```

▼ How to Add a Network Printer by Using Printer Vendor-Supplied Tools

Steps

1. **Connect the printer to the network. Turn on the power to the printer.**

   Consult the printer vendor's installation documentation for information about the hardware switches and cabling requirements. Get an IP address and select a name for the printer node. This method is equivalent to adding any node to the network.

2. **Follow the printer vendor instructions to add the network printer.**

   Use the printer vendor instructions to configure the network printer. These instructions are specific to the vendor and printer.

3. **Add client access to the new printer.**

   Now that the printer has been added, create access to the printer for the print clients. For more information, see "Setting Up a Print Client" on page 69.

4. **Complete any optional tasks.**

   There are several optional tasks that you might want to complete when setting up a network printer. For pointers to these tasks, see "Setting Up Printing (Task Map)" on page 57.

▼ How to Add a New Network Printer by Using Solaris Print Manager

In this release, the printer definitions that you can assign with Solaris Print Manager have been expanded to enable you to use PPD files. A new Use PPD files option has been added to the Print Manager menu drop-down menu in Solaris Print Manager.

This default option enables you to select the printer make, model, and driver when adding new printer. For overview information about new or modified Solaris Print Manager features, see "What's New in Printing?" on page 25.

This procedure shows you how to add a new network printer by using Solaris Print Manager with the Use PPD files default option selected. To add a new network printer without PPD files, you must first deselect this option in the Print Manager drop-menu.

Steps

1. **Select the system to be the print server.**
2. **Connect the printer to the network. Turn on the power to the printer.**

   Consult the printer vendor's installation documentation for information about the hardware switches and cabling requirements.
3. **Start Solaris Print Manager on the print server.**

   For instructions, see "How to Start Solaris Print Manager" on page 61.
4. **Select New Network Printer from the Printer menu.**

   The New Network Printer window is displayed.
5. **Fill in the required fields with the appropriate information.**

   If you need information to complete a field, click the Help button.

   For example, to add a new network printer with PPD files, complete the following fields:

   ```
   Printer Name:
   Description:
   Printer Make:
   Printer Model:
   Printer Driver:
   Fault Notification:
   Destination:
   Protocol:
   Options:
   Default Printer
   Always Use Banner
   User Access List:
   ```

   The Printer Server field is filled in by Solaris Print Manager.

   To add a new network printer without PPD files, complete the following fields:

   ```
   Printer Name:
   Description:
   Printer Type:
   File Contents:
   Fault Notification:
   destination:
   Protocol:
   Options:
   Default Printer
   Always Use Banner
   ```

```
User Access List:
```

The Printer Server field is filled in by Solaris Print Manager.

6. **Click OK.**
7. **Verify that the printer has been installed. Check for the new printer entry in the Solaris Print Manager main window.**
8. **Verify that the printer can print requests.**

   `$ lp -d` *printer-name filename*

9. **Exit Solaris Print Manager.**

   Choose Exit from the Print Manager Menu.

▼ How to Add a New Network Printer by Using LP Print Service Commands

This procedure describes how to add a new network printer with the expanded printer support that is available in this release. There are two examples that follow this procedure. The first example shows you how to add a new network printer with PPD files by using `lp` commands. The second example shows you how to add a new network printer without PPD files by using `lp` commands.

The following steps are necessary to set up a network printer by using the network printer support software. The use of this software is intended for printers that do not come with vendor-supplied software.

For overview information about adding a new network printer or modifying an existing network printer with PPD files, see "What's New in Printing?" on page 25.

Steps

1. **Connect the printer to the network, and turn on the power to the printer.**

   Consult the printer vendor's installation documentation for information about the hardware switches and cabling requirements. Get an IP address and select a name for the printer node. This step is equivalent to adding any node to the network.

2. **Collect the information that is required to configure a new network printer with PPD files.**
   - Printer name and port device
   - Network printer access name
   - Protocol
   - Timeout
   - File content type
   - PPD file

To determine the PPD file, first define the printer make, model, and driver that the printer will use.

For more information, see the terms that are described in "Introduction to Adding a Network Printer" on page 71.

3. **Define the printer name, port device, file content type, and PPD file that the printer will use. When using PPD files, the file content type is usually PostScript.**

   a. **Specify the printer name and the port device that the printer will use.**

   ```
   # lpadmin -p printer-name -v /dev/null
   ```

   The device to use is `/dev/null`.

   b. **Specify the interface script that the printer will use.**

   ```
   # lpadmin -p printer-name -m netstandard_foomatic
   ```

   The interface script that is supplied with the support software for network printer is `/usr/lib/lp/model/netstandard_foomatic`.

   c. **Specify the file content type and PPD file that the printer will use.**

   When using PPD files, the file content type is usually PostScript.

   ```
   # lpadmin -p printer-name -I content-type -n
   /path/ppdfile
   ```

   The PPD file that you specify is located in the */path* directory

   For more information, see `lpadmin`(1M).

   d. **Specify the printer description.**

   ```
   # lpadmin -p printer-name -D "printer-description"
   ```

4. **Set the printer destination, protocol, and timeout values.**

   ```
   # lpadmin -p printer-name -o dest=access-name:port -o protocol=protocol
   -o timeout=value
   ```

`-p` *printer-name*	Specifies the network printer name.
`-o dest=`*access-name:port*	Sets the printer destination to the network printer access name and a designated printer vendor port, if the port is defined in the printer vendor documentation.
`-o protocol=`*protocol*	Sets the over-the-wire protocol used to communicate with the printer. Both BSD and raw TCP are supported.

-o timeout=*value*	Sets a retry timeout value that represents a number of seconds to wait between attempting connections to the printer.

5. **(Optional) Add filters to the print server.**

```
# cd /etc/lp/fd
# for filter in *.fd;do
    > name=`basename $ filter .fd`
    > lpfilter -f $name -F $ filter
    > done
```

6. **Enable the printer to accept print requests and to print those requests.**

```
# accept printer-name
# enable printer-name
```

7. **Verify that the printer is correctly configured.**

```
# lpstat -p printer-name
```

8. **Add client access to the new printer.**

   Now that the printer has been added, create access to the printer for the clients. For more information, see "Setting Up a Print Client" on page 69. For information about how to define a print client by using a name service, see "Printing Support in the Name Service Switch" on page 27.

9. **Complete any optional tasks.**

   There are several optional tasks that you might want to complete when setting up a printer. For pointers to the remaining tasks, see "Setting Up Printing (Task Map)" on page 57. For information about how to define print clients by using the name service, see "Printing Support in the Name Service Switch" on page 27.

Example 3–7 Adding a New Network Printer With PPD Files by Using LP Print Service Commands

This example shows how to add a new network printer with PPD files. The `-n` option to the `lpadmin` command enables you to add a new print queue by specifying PPD files. The commands used in this example must be executed on the print server. The following information is used as an example only. The information that you provide will vary.

- Printer name : `luna1`
- Device: `/dev/null`
- Interface: `/usr/lib/lp/model/netstandard_foomatic`
- Network printer access name: `nimquat:9100`
- Protocol: `tcp`
- Timeout: 5

- File content type: postscript
- PPD file: */path/ppdfile*

```
# lpadmin -p luna1 -v /dev/null     (1)
# lpadmin -p luna1 -m netstandard_foomatic     (2)
# lpadmin -p luna1 -o dest=nimquat:9100 -o protocol=tcp
-o timeout=5     (3)
# lpadmin -p luna1 -I postscript     (4)
# lpadmin -p luna1 -n /path/ppdfile     (5)
# lpadmin -p luna1 -D "Second floor color printer"     (6)
# cd /etc/lp/fd
# for filterin *.fd;do
       > name =`basename $ filter .fd`
       > lpfilter -f $ name -F $ filter
       > done     (7)
# accept luna1
    destination "luna1" now accepting requests
# enable luna1     (8)
printer "luna1" now enabled
# lpstat -p luna1     (9)
 printer luna1 is idle. enabled since Mon Apr 19 15:31:15 2004. available.
```

1. Defines printer name and sets the device to /dev/null
2. Defines the interface script for network printers
3. Sets the destination, protocol, and timeout
4. Specifies the file content types to which the printer can print to directly
5. Specifies the PPD file to use
6. Adds a description for the printer
7. Adds print filters to the print server
8. Accepts print requests for the printer and enables the printer
9. Verifies that the printer is ready for printing

Example 3–8 Adding a New Network Printer Without PPD Files by Using LP Print Service Commands

This example shows how to add a new network printer without PPD files. The commands must be executed on the print server. The following information is used as an example. The information that you provide will vary.

- Printer name: luna1
- Device: /dev/null
- Interface: /usr/lib/lp/model/netstandard
- Network printer access name: nimquat:9100
- Protocol: tcp
- Timeout: 5
- File content type: postscript
- Printer type: PS

```
# lpadmin -p luna1 -v /dev/null     (1)
# lpadmin -p luna1 -m netstandard     (2)
# lpadmin -p luna1 -o dest=nimquat:9100 -o protocol=tcp
-o timeout=5     (3)
```

```
# lpadmin -p luna1 -I postscript -T PS      (4)
# lpadmin -p luna1 -D "Room 1954 ps"     (5)
# cd /etc/lp/fd
# for filter in *.fd;do
       > name ='basename $ filter .fd'
       > lpfilter -f $ name -F $ filter
       > done     (6)
# accept luna1
    destination "luna1" now accepting requests
# enable luna1     (7)
printer "luna1" now enabled
# lpstat -p luna1     (8)
 printer luna1 is idle. enabled since Feb 28 11:21 2003.  available.
```

1. Defines printer name and sets the device to `/dev/null`
2. Defines the interface script for network printers
3. Sets the destination, protocol, and timeout
4. Specifies the file content type to which the printer can print directly, and the printer type
5. Adds a description for the printer
6. Adds print filters to the print server
7. Accepts print requests for the printer and enables the printer
8. Verifies that the printer is ready for printing

CHAPTER 4

Administering Printers (Tasks)

This chapter describes the procedures for administer printers with the `lp` commands.

For information on the step-by-step procedures associated with administering printers, see the following:

- "Managing Printers and the Print Scheduler (Task Map)" on page 83
- "Setting or Resetting Miscellaneous Printer Definitions (Task Map)" on page 90
- "Managing Print Requests (Task Map)" on page 103

For overview information about printing and the LP print service, see Chapter 1.

Managing Printers and the Print Scheduler (Task Map)

Task	Description	For Instructions
Delete a printer and remove printer access.	If a printer needs to be replaced or moved to a different location, you can delete information about the printer from the *print-client*.	"How to Delete a Printer and Remove Printer Access" on page 85

Task	Description	For Instructions
Check the status of printers.	Use the `lpstat` command to check the status of all printers or a specific printer. This command allows you to determine which printers are available for use and enables you to examine the characteristics of those printers.	"How to Check the Status of Printers" on page 87
Stop the Print Scheduler.	If the print scheduler stops accepting print requests, you might need to stop and restart the print scheduler.	"How to Stop the Print Scheduler" on page 89
Restart the Print Scheduler.	After you stop the print scheduler, you will need to restart the print scheduler, so the printer can begin accepting requests.	"How to Restart the Print Scheduler" on page 90

Managing Printers and the Print Scheduler

This section provides instructions for day-to-day tasks you perform to manage printers and the print scheduler.

Deleting Printers and Printer Access

If a printer needs to be replaced or you want to move the printer to a different location, you must delete the printer information from the LP print service before you physically remove it from the print server. You should also make sure that all the current print requests on the printer are printed or moved to another printer to be printed.

Besides deleting the printer information from the print server, you need to delete the information from the print clients or network name service. If you delete a local printer from a print server, you should delete the remote printer entry from the print clients or network name service. If you move a printer to another print server, you need to delete the old remote printer entry from the print clients or network name service. Then, add access to the remote printer in its new location.

For detailed information on how to delete a local printer and a remote printer, see "How to Delete a Printer and Remove Printer Access" on page 85. You can use Solaris Print Manager to delete either type of printer. However, Solaris Print Manager does not enable you to move queued print requests to another printer.

▼ How to Delete a Printer and Remove Printer Access

Use this procedure when you remove a printer from service, and you want to remove the printer access from the print clients. Then, remove the printer information from the print server.

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on a print client that has access to the printer you want to delete.**

2. **Delete information about the printer from the print client.**

   print-client# **`lpadmin -x`** *printer-name*

`-x`	Deletes the specified printer.
printer-name	Specifies the name of the printer you want to delete.

   Information for the specified printer is deleted from the print client's `/etc/lp/printers` directory.

3. **If the print client does not use another printer on the same print server, delete information about the print server from the print client.**

   print-client# **`lpsystem -r`** *print-server*

`-r`	Removes the specified print server.
print-server	Specifies the name of the print server you want to delete.

   The print server is deleted from the print client's `/etc/lp/Systems` file.

4. **Repeat Step 2 through Step 3 on each print client that has access to the printer.**

5. **Log in as superuser `lp`, or assume an equivalent role on the print server.**

6. **Stop accepting print requests on the printer.**

   print-server# **`reject`** *printer-name*

   `reject` *printer-name* rejects print requests for the specified printer.

   This step prevents any new requests from entering the printer's queue while you are in the process of removing the printer. For a detailed description, see "How to Accept or Reject Print Requests for a Printer" on page 110.

7. **Stop the printer.**

   print-server# **`disable`** *printer-name*

This step stops print requests from printing. For a detailed description on how to stop printing, see "How to Enable or Disable a Printer" on page 108.

8. **Move any print requests that are still in the queue to another printer.**

   For a detailed description on how to move print requests to another printer, see "How to Move Print Requests to Another Printer" on page 113.

9. **Delete the printer from the print server.**

   print-server# **`lpadmin -x`** *printer-name*

   Configuration information for the printer is deleted from the print server's `/etc/lp/printers` directory.

10. **Delete information about the print clients that were using the printer you just deleted, unless they are still using another printer on the print server.**

    print-server# **`lpsystem -r`** *print-client1* `[`,*print-client2* `...]`

`-r`	Removes the specified print client.
print-client	Specifies the name of the print client you want to delete from the print server. You can specify multiple print clients in this command. Use a space or a comma to separate print client names. If you use spaces, enclose the list of print clients in quotation marks.

    The specified print clients are deleted from the print server's `/etc/lp/Systems` file.

11. **Verify that the printer information has been deleted:**

    a. **Confirm that the printer information has been deleted on the print client.**

       print-client$ **`lpstat -p`** *printer-name* **`-l`**

       In the command output, you should receive an error indicating that the printer does not exist.

    b. **Confirm that the printer information has been deleted on the print server.**

       print-server$ **`lpstat -p`** *printer-name* **`-l`**

       In the command output, you should receive an error indicating that the printer does not exist.

Example 4–1 Deleting a Printer and Remote Printer Access

The following example shows how to delete the printer `luna` from the print client `terra` and from the print server `jupiter`. This example also shows how to delete the print client `terra` from the print server.

```
terra# lpadmin -x luna
Removed "luna".
terra# lpstat -p luna -l
```

```
jupiter# lpadmin -x luna
jupiter# lpsystem -r terra
Removed "terra".
jupiter# lpstat -p luna -l
```

Checking Printer Status

Many routine printer administration tasks require information about the status of the LP print service or a specific printer. For example, you can determine which printers are available for use. You can also examine the characteristics of those printers. You can use the `lpstat` command to find status information about the LP print service or a specific printer.

In this release, if a PPD file was used to create or modify a print queue, the specific PPD file is listed in the `lpstat` output. If a PPD file was not specified when the print queue was created, the PPD entry in the `lpstat` output is `none`.

▼ How to Check the Status of Printers

Steps

1. **Log in to any system on the network.**
2. **Check the status of printers.**

   Only the most commonly used options are shown here. For other options, see the`lpstat`(1) man page.

   `$ lpstat [-d] [-p] printer-name [-D] [-l] [-t]`

   `-d` Shows the system's default printer.

   `-p` *printer-name* Shows if a printer is active or idle, when the printer was enabled or disabled, and whether the printer is accepting print requests.

   You can specify multiple printer names with this command. Use a space or a comma to separate printer names. If you use spaces, enclose the list of printer names in quotation marks. If you don't specify *printer-name*, the status of all printers is displayed.

   `-D` Shows the description of the specified *printer-name*.

   `-l` Shows the characteristics of the specified *printer-name*.

   `-t` Shows status information about the LP print service, including the status of all printers, such as whether they are active and whether they are accepting print requests.

Example 4–2 Checking the Status of Printers

The following example shows how to display the system's default printer.

```
$ lpstat -d
system default destination: luna
```

The following example shows how to display the status of the printer `luna`.

```
$ lpstat -p luna
printer luna is idle. enabled since Jul 12 11:17 2001. available.
```

The following example shows how to display the description of the printers `asteroid` and `luna`.

```
$ lpstat -p "asteroid luna" -D
printer asteroid faulted. enabled since Jan 5 11:35 2004. available.
unable to print: paper misfeed jam

Description: Printer by break room
printer luna is idle. enabled since Jan 5 11:36 2004. available.
Description: Printer by server room.
```

The following example shows how to display the characteristics of the printer `luna`.

```
$ lpstat -p luna -l
printer luna is idle. enabled since Thu Jul 12 15:02:32 PM PDT
        Form mounted:
        Content types: postscript
        Printer types: PS
        Description:
        Connection: direct
        Interface: /usr/lib/lp/model/standard
                PPD: none
        After fault: continue
        Users allowed:
                (all)
        Forms allowed:
                (none)
        Banner not required
        Character sets:
                             (none)
                Default pitch:
        Default page size: 80 wide 66 long
        Default port settings:
```

Example 4–3 `lpstat` Command Output When Using PPD Files

In this example, the `lpstat` command output indicates that the print queue was configured by using the PPD file, `Mitsubishi-CP50_Color_Printer-cp50.ppd.gz`.

```
# lpstat -l -p paper
printer paper is idle. enabled since Tue 30 Mar 2004 01:48:38 PM PST
            available.
        Form mounted:
        Content types: any
        Printer types: unknown
        Description:
        Connection: direct
        Interface: /usr/lib/lp/model/standard_foomatic
        PPD:
/path/Mitsubishi-CP50_Color_Printer-cp50.ppd.gz
        After fault: continue
        Users allowed:
                (all)
        Forms allowed:
                (none)
        Banner required
        Character sets:
                (none)
        Default pitch:
        Default page size:
        Default port settings# lpstat l p <queue>
```

Stopping and Restarting the Print Scheduler

The print scheduler, lpsched, handles print requests on print servers. However, the print scheduler might sometimes stop running on a system, so print requests stop being accepted or printed.

In this Solaris release, the lpadmin command automatically enables the lpsched service when local printers are added to the system and disables it when the last local printer is removed. The following procedures describe how to stop and restart the print scheduler. If a print request was printing when the print scheduler stopped running, the print request will be printed in its entirety when you restart the print scheduler.

▼ How to Stop the Print Scheduler

Steps

1. **Log in as superuser, lp, or assume an equivalent role on the print server.**
2. **Determine if the print scheduler is running.**

   ```
   # svcs application/print/server
   ```

   If the print scheduler is not running, the message scheduler is not running is displayed.

3. **If the print scheduler is running, stop it.**

```
# svcadm disable application/print/server
```

How to Restart the Print Scheduler

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Determine if the print scheduler is running.**

```
# svcs application/print/server
```

If the print scheduler is not running, the message `scheduler is not running` is displayed.

3. **If the print scheduler is not running, start it.**

```
# svcadm enable application/print/server
```

Setting or Resetting Miscellaneous Printer Definitions (Task Map)

Task	Description	For Instructions
Add a printer description.	User the `lp` command with the `-p` option to add a printer description.	"How to Add a Printer Description" on page 92
Set a default printer for a system.	You can specify a default printer destination for a user so the user does not need to type the printer name.	"How to Set a System's Default Printer Destination" on page 93
Make banner pages optional.	You can give users a choice to turn off printing of a banner page when they submit a print request.	"How to Make Banner Pages Optional" on page 95

Task	Description	For Instructions
Turn off banner pages.	By setting the `-o banner=never` variable for a printer you can turn off banner pages under all circumstances.	"How to Turn Off Banner Pages" on page 95
Define a class of printers.	You can group printers into a class by using the `lpadmin -c` command. Printer could be grouped into the following classes: ■ Printer type ■ Location ■ Work group	"How to Define a Class of Printers" on page 97
Set fault alters for a printer.	If you choose, the print service can notify you when it detects a printer fault. Use the `-A` option with the `lpadmin` command to set fault alerts for a printer.	"How to Set Fault Alerts for a Printer" on page 98
Set printer fault recovery.	You can choose to not receive any fault notification. You can find out about printing faults so you can correct the problem. Use the `-F` option with the `lpadmin` command to define fault recovery options.	"How to Set Printer Fault Recovery" on page 100
Limit user access to a printer.	Create allow and deny lists on the print serve to control user access to printers	"How to Limit User Access to a Printer" on page 102

Setting or Resetting Miscellaneous Printer Definitions

The following procedures provide step-by-step instructions on setting or resetting printer definitions. Some of the following printer definitions can be set by using Solaris Print Manager. The following procedures show how to use the LP print service commands to quickly set or reset printer definitions.

▼ How to Add a Printer Description

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**
2. **Add a printer description.**

   ```
   # lpadmin -p printer-name -D "comment"
   ```

   `-p` *printer-name* — Specifies the name of the printer for which you are adding a description.

   `-D` *"comment"* — Specifies the characteristics of the printer, such as the location or administrative contact. Enclose characters that the shell might interpret, such as `*`, `?`, `\`, `!`, `^`, in single quotation marks.

   For more information, see the `lpadmin`(1M) man page.

   The printer description is added in the print server's `/etc/lp/printers/`*printer-name*`/comment` file.
3. **Verify that the `Description` information is correct.**

   ```
   # lpstat -p printer-name -l
   ```

Example 4–4 Adding a Printer Description

The following example shows how to add a printer description for the printer `luna`.

```
# lpadmin -p luna -D "Nathans office"
```

Setting Up a Default Printer Destination

You can specify a default printer destination for a user so that the user does not need to type the printer name when using the print commands. Before you can designate a printer as the default, the printer must be known to the print service on the system. You can set a user's default printer destination by setting any of the following:

- `LPDEST` environment variable
- `PRINTER` environment variable
- The `_default` variable in the user's `.PRINTERS` file
- The system's default printer by using the `lpadmin -d` command or Solaris Print Manager

When an application provides a printer destination, that destination is used by the print service, regardless of whether you have set a system's default printer destination. If an application doesn't provide a printer destination or if you don't provide a printer name when using a print command, the print command searches for the default printer in a specific order. The following table shows the search order for a system's default printer destination.

TABLE 4–1 Search Order for Default Printer Destinations

Search Order	Using `/usr/bin/lp` Command	Using `lpd`-Based Compatibility Commands (`lpr`, `lpq`, and `lprm`)
First	`LPDEST` variable	`PRINTER` variable
Second	`PRINTER` variable	`LPDEST` variable
Third	System's default printer	System's default printer

▼ How to Set a System's Default Printer Destination

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the system where you want to set a default printer.**

2. **Set the system's default printer.**

   ```
   # lpadmin -d [printer-name]
   ```

   `-d` *printer-name* specifies the name of the printer you are assigning as the system's default printer. If you don't specify *printer-name*, the system is set up with no default printer.

   The default printer name is specifed in the system's `/etc/lp/default` file.

3. **Check the system's default printer.**

   ```
   # lpstat -d
   ```

Example 4–5 Setting a System's Default Printer Destination

The following example shows how to set the printer `luna` as the system's default printer. The printer `luna` is used as the system's default printer if the `LPDEST` or the `PRINTER` environment variables are not set.

```
# lpadmin -d luna
# lpstat -d
system default destination: luna
```

Printing Banner Pages

A banner page identifies who submitted the print request, the print request ID, and when the request was printed. A banner page also has a modifiable title to help users identify their printouts.

Banner pages make identifying the owner of a print job easy, which is especially helpful when many users submit jobs to the same printer. However, printing banner pages uses more paper and might be unnecessary if a printer has only a few users. In some cases, printing banner pages is undesirable. For example, if a printer has special paper or forms mounted, such as paycheck forms, printing banner pages might cause problems.

By default, the print service forces banner pages to be printed. However, you can give users a choice to turn off printing of a banner page when they submit a print request. You can set this choice through the `lpadmin` command or through Solaris Print Manager. If you give the users a choice, they have to use the `-o banner` option to turn off banner page printing.

Also, if you don't need or want banner pages, you can turn off banner pages so that they are never printed. You can turn off banner page printing by using the `lpadmin` command.

TABLE 4–2 Banner Page Printing

Command Used	Banner Page Printing	Override?
`lpadmin -p` *printer* `-o banner` OR `lpadmin -p` *printer* `-o banner=always`	Required and printed	If you are a regular user and use `p -o nobanner`, the request is printed. However, the `nobanner` argument is ignored. If you are superuser (`root`), or another privileged user, the `nobanner` argument is honored.

`lpadmin -p` *printer* `-o nobanner`
 	
`lpadmin -p` *printer* `-o banner=optional`
	On by default, but can be disabled on a per request basis with the `lp -o nobanner` command	Not applicable.

`lpadmin -p` *printer* `-o banner=never`	Disabled	No.

For step-by-step command-line instructions, see "How to Turn Off Banner Pages" on page 95 and "How to Make Banner Pages Optional" on page 95.

▼ How to Make Banner Pages Optional

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Make banner pages optional.**

   ```
   # lpadmin -p printer-name -o banner=optional
   ```

`-p` *printer-name*	Specifies the name of the printer for which you are making banner pages optional.
`-o banner=optional`	Enables users to specify no banner page when they submit a print request.

   If you want to force a banner page to print with every print request, specify the `-o banner=always` option. The banner page setting is specified in the print server's `/etc/lp/printers/`*printer-name*`/configuration` file.

3. **Verify that the banner page is optional. The command output contains the line `Banner not required`.**

   ```
   # lpstat -p printer-name -l
   ```

Example 4–6 Making Banner Pages Optional

The following example shows how to make the banner page optional on the printer `luna`.

```
# lpadmin -p luna -o banner=optional
```

▼ How to Turn Off Banner Pages

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Turn off banner printing.**

   ```
   # lpadmin -p printer-name -o banner=never
   ```

`-p` *printer-name*	Specifies the name of the printer for which you are turning off banner pages.
`-o banner=never`	Disables banner page printing under all circumstances.

   The banner page setting is specified in the print server's `/etc/lp/printers/`*printer-name*`/configuration` file.

3. **Verify that banner printing is turned off.**

   ```
   # lpstat -p printer-name -l
   ```

The command output contains the line `Banner not printed`.

4. **Submit a print request to the printer to ensure that a banner page does not print.**

Example 4–7 Turning Off Printing Banner Pages

The following example shows how to disable the printing of banner pages on the printer `luna`.

```
# lpadmin -p luna -o banner=never
```

Setting Up Printer Classes

The LP print service enables you to group several locally attached printers into one class. You can perform this task only by using the `lpadmin -c` command.

After you set up a printer class, users can then specify that class, rather than individual printers, as the destination for a print request. The first printer in the class that is available to print is used. The result is faster turnaround because printers are kept as busy as possible.

No default printer classes are known to the print service. Printer classes exist only if you define them. Here are some ways you could define printer classes:

- **By printer type** – For example, a PostScript printer.
- **By location** – For example, a printer located on the 5th floor.
- **By work group or department** – For example, the accounting department.

Alternatively, a class might contain several printers that are used in a particular order. The LP print service always checks for an available printer in the order in which printers were added to a class. Therefore, if you want a high-speed printer to be accessed first, add the high-speed printer to the class before you add a low-speed printer. As a result, the high-speed printer handles as many print requests as possible. The low-speed printer is reserved as a backup printer when the high-speed printer is in use.

Note – Print requests are balanced between printers in a class only for local printers.

Class names, similar to printer names, must be unique and can contain a maximum of 14 alphanumeric characters and underscores.

You are not obligated to define printer classes. You should add printer classesonly if you determine that using them would benefit users on the network.

▼ How to Define a Class of Printers

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Define a class of printers.**

   ```
   # lpadmin -p printer-name -c printer-class
   ```

`-p` *printer-name*	Specifies the name of the printer you are adding to a class of printers.
`-c` *printer-class*	Specifies the name of a class of printers.

   The specified printer is added to the end of the class list in the print server's `/etc/lp/classes/`*printer-class* file. If the printer class does not exist, it is created.

3. **Verify that the printers are in a printer class.**

   ```
   # lpstat -c printer-class
   ```

Example 4–8 Defining a Class of Printers

The following example shows how to add the printer `luna` in the printer class `roughdrafts`.

```
# lpadmin -p luna -c roughdrafts
```

Setting Up Printer Fault Alerts

If you choose, the LP print service can notify you when it detects a printer fault. With the `lpadmin -A` command or with Solaris Print Manager, you can select any of the following methods to receive printer fault notification:

- A message to the terminal on which `root` is logged in
- Electronic mail to `root`
- No notification

However, the `lpadmin -A` command offers you an additional option of receiving a message specified by the program of your choice. The `lpadmin -A` command also enables you to selectively turn off notification for an error that you already know about.

Unless you specify a program to deliver fault notification, the content of the fault alert is a predefined message that states that the printer has stopped printing and needs to be fixed.

The following table lists the alert values that you can set for a printer with the `lpadmin -A` command. These alert values can also be set for print wheels, font cartridges, and forms.

TABLE 4–3 Values for Printer Fault Alerts

Value for `-A alert`	Description
'`mail` [*username*]'	Sends the alert message by email to `root` or `lp` on the print server, or the specified *username*, which is a name of a user.
'`write` [*user-name*]'	Sends the alert message to the `root` or `lp` console window on the print server, or to the console window of the specified *username*, which is a name of a user. The specified user must be logged in to the print server to get the alert message.
'*command*'	Runs the specified *command* file for each alert. The environment variables and current directory are saved and restored when the file is executed.
`quiet`	Stops alerts until the fault is fixed. Use this valve when you, `root` or a specified user, receive repeated alerts.
`none`	Does not send any alerts. This value is the default if you don't specify fault alerts for a printer.

▼ How to Set Fault Alerts for a Printer

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Set fault alerts for a printer.**

   ```
   # lpadmin -p printer-name -A alert [-W minutes]
   ```

   `-p` *printer-name* — Specifies the name of the printer for which you are specifying an alert for printer faults.

   `-A` *alert* — Specifies what kind of alert occurs when the printer faults. For detailed information about the valid values for *alert*, see Table 4–3. Some valid values are `mail`, `write`, and `quiet`.

   `-W` *minutes* — Specifies how often, in minutes, the fault alert occurs. If you don't specify this option, the alert is sent one time.

   The fault alert setting is specified in the print server's `/etc/lp/printers/`*printer-name*`/alert.sh` file.

3. **Verify that the fault alert has been sent properly.**

   ```
   # lpstat -p printer-name -l
   ```

Example 4–9 Setting Fault Alerts for a Printer

The following example shows how to set up the printer mars to send fault alerts by email to a user named joe. A reminder is sent every 5 minutes.

```
# lpadmin -p mars -A 'mail joe' -W 5
```

The following example shows how to set up the printer venus to send fault alerts to the console window. A reminder i sent every 10 minutes.

```
# lpadmin -p venus -A write -W 10
```

The following example shows how to stop fault alerts for the printer mercury.

```
# lpadmin -p mercury -A none
```

The following example shows how to stop fault alerts until the printer venus has been fixed.

```
# lpadmin -p venus -A quiet
```

Setting Up Printer Fault Recovery

If you choose not to send any fault notification, you can still find out about printer faults so that you can correct the problem. The LP print service will not continue to use a printer that has a fault. In addition to alerts for printer faults, you can also provide alerts that instruct you to mount print wheels, font cartridges, and forms when print requests require them.

You can define the fault recovery options for a printer only by using the lpadmin -F command. This task is not available in Solaris Print Manager.

Printer faults can be as simple as running out of paper or needing to replace a toner cartridge. Other more serious problems can include complete printer failure or power failure. After you fix a printer fault, the print request that was active when the fault occurred begins printing in one of three ways:

- Starts printing from the beginning
- Continues printing from the top of the page where printing stopped
- After you enable the printer, continues printing from the top of the page where the printing stopped

The LP print service requires a separate print filter to continue printing from the top of a page where the printing stopped. This print filter records the control sequences set by the default print filters. The printer uses these control sequences to track page boundaries. You will be notified by the LP print service if recovery cannot proceed with the specified print filter. For information about writing filters, see "How to Create a New Print Filter" on page 166.

If you want printing to resume immediately after a printer fault is fixed, enable the printer by using the `enable` command.

The following table lists the fault recovery values that you can set for a printer by using the `lpadmin -F` command.

TABLE 4–4 Values for Printer Fault Recovery

Value for `-F` *recover-options*	Description
`beginning`	After a fault recovery, printing restarts from the beginning of the file.
`continue`	After a fault recovery, printing starts at the top of the page where the printing stopped. This recovery option requires a print filter.
`wait`	After a fault recovery, printing stops until you enable the printer. After you enable the printer by using the `enable` command, printing starts at the top of the page where printing stopped. This recovery option requires a print filter.

▼ How to Set Printer Fault Recovery

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Set up fault recovery for the printer.**

   ```
   # lpadmin -p printer-name -F recovery-options
   ```

   `-p` *printer-name* — Specifies the name of the printer for which you are specifying fault recovery.

   `-F` *recovery-options* — Specifies one of the three valid recovery options: `beginning`, `continue`, or `wait`.

   For detailed information about the valid values for *recovery-options*, see Table 4–4.

   For more information, see the `lpadmin`(1M) man page.

   The fault recovery setting is specified in the print server's `/etc/lp/printers/`*printer-name*`/configuration` file.

3. **Verify that printer fault recovery has been set up properly.**

   ```
   # lpstat -p printer-name -l
   ```

Example 4–10 Setting Printer Fault Recovery

The following example shows how to set up the printer `luna` to continue printing at the top of the page where printing stopped.

```
# lpadmin -p luna -F continue
```

Limiting User Access to a Printer

You can control which users can access some printers or all of the available printers. For example, you can prevent some users from printing on a high-quality printer to minimize expense. To restrict user access to printers, you create `allow` and `deny` lists by using the `lpadmin -u` command on the print server. Solaris Print Manager enables you to create only `allow` lists. If you create neither list, a printer is available to all users who can access the printer.

An `allow` list contains the names of users who are allowed access to the specified printer. A `deny` list contains the names of users denied access to the specified printer.

The rules for `allow` and `deny` lists are explained in the following table:

Allow and Deny List Rule	User Consequence
Do not create `allow` and `deny` lists. Or, if you leave both lists empty.	All users can access the printer.
Specify `all` in the `allow` list.	All users can access the printer.
Specify `all` in the `deny` list	All users, except `root` and `lp` user on the server, are denied access to the printer.
Make any entry in the `allow` list.	The `deny` list is ignored. Only users who are listed can access the printer.
Create a deny list, but you do not create an `allow` list or you leave the `allow` list empty.	Users who are listed in the `deny` list are denied access to the printer.

Because the print server is responsible for controlling access to the printer, `allow` and `deny` lists can only be created on the print server. If you create `allow` and `deny` lists, the print server will exclusively control user access to printers.

The following table lists the values that you can add to an `allow` list or a `deny` list to limit user access to a printer.

TABLE 4–5 Values for Allow and Deny Lists

Value for *user-list*	Description
user	*User* on any system
`all`	All users on all systems
`none`	No user on any system
system`!`*user*	*User* on *system* only
`!`*user*	*User* on local system only
`all!`*user*	*User* on any system
`all!all`	All users on all systems
system`!all`	All users on *system*
`!all`	All users on local system

▼ How to Limit User Access to a Printer

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Allow or deny users access to a printer.**

   `# `**`lpadmin -p`** *printer-name* **`-u allow`**`:`*user-list* `[`**`deny`**`:`*user-list*`]`

   `-p` *printer-name* — Specifies the name of the printer to which the `allow` or `deny` access list applies.

   `-u allow:`*user-list* — Specifies the user names to be added to the `allow` access list. You can specify multiple user names with this command. Use a space or a comma to separate names. If you use spaces, enclose the list of names in quotation marks.

   Table 4–5 provides the valid values for *user-list*.

   `-u deny:`*user-list* — Specifies user names to be added to the deny user access list. You can specify multiple user names with this command. Use a space or a comma to separate names. If you use spaces, enclose the list of names in quotation marks.

   Table 4–5 provides the valid values for *user-list*.

   The specified users are added to the `allow` or `deny` list for the printer in one of the following files on the print server:

   `/etc/lp/printers/`*printer-name*`/users.allow`

/etc/lp/printers/*printer-name*/users.deny

Note – If you specify none as the value for *user-list* in the allow list, the following files are not created on the print server:

/etc/lp/printers/*printer-name*/alert.sh

/etc/lp/printers/*printer-name*/alert.var

/etc/lp/printers/*printer-name*/users.allow

/etc/lp/printers/*printer-name*/users.deny

---

3. **Verify that information is correct under the `Users allowed` or the `Users denied` heading in the following command output:**

   ```
   # lpstat -p printer-name -l
   ```

Example 4–11 Limiting User Access to a Printer

The following example shows how to allow only the users nathan and george access to the printer luna.

```
# lpadmin -p luna -u allow:nathan,george
```

The following example shows how to deny the users nathan and george access to the printer asteroid.

```
# lpadmin -p asteroid -u deny:"nathan george"
```

Managing Print Requests (Task Map)

Task	Description	For Instructions
Check the status of print requests.	Use the lpstat command to check the status of print requests.	"How to Check the Status of Print Requests" on page 106

Task	Description	For Instructions
Enable or disable a printer.	Use the `enable` and `disable` commands to control whether a printer prints requests or stops printing requests that are in the print queue.	"How to Enable or Disable a Printer" on page 108
Accept or reject print requests for a printer.	Use the `accept` and `reject` commands to turn on or turn off a print queue that stores requests to be printed.	"How to Accept or Reject Print Requests for a Printer" on page 110
Cancel a print request.	You can use the can use the cancel command to cancel print requests in one of the following three ways: ■ Cancel requests by request identification number (request ID) ■ Cancel requests from a specific user on all printers or specific printers ■ Cancel the job currently printing	"How to Cancel a Print Request" on page 111
Cancel a print request from a specific user.	You can cancel the request for a user by using the `cancel -u` command and specifying that user.	"How to Cancel a Print Request From a Specific User" on page 112
Move print requests to another printer.	You can move or cancel any requests that are currently queued a printer. You can use the `lpmove` command to move individual or all print requests to another local printer.	"How to Move Print Requests to Another Printer" on page 113

Task	Description	For Instructions
Change the priority of a print request.	You can change the priority of a print request in one of the following ways: ■ Put a print request on hold if it has not finished printing. ■ Move a print request to the head of the queue. ■ Change the priority of a job still waiting to be printed.	"How to Change the Priority of a Print Request" on page 115

Managing Print Requests

When a user submits a print request from a print client, the print request is added to a queue on the print server before it is sent to the printer. While a print request is in the queue, you can cancel or gain status information on the request from a client system. You must log in to the print server to move, hold, resume, or change the priorities of print requests with LP print service commands. These actions can help you keep printing services operating smoothly.

The following table lists the values for changing the priority of a print request by using the `lp -H` command.

TABLE 4–6 Values for Changing the Priority of a Print Request

Value for `-H` *change-priority*	Description
`hold`	Places the print request on hold until you cancel the request or instruct the LP print service to resume printing the request.
`resume`	Places a print request that has been on hold back in the queue. The request is printed according to its priority and placement in the queue. If you place on hold a print job that is already printing, `resume` places the print request at the head of the queue so that it becomes the next request printed.
`immediate`	Places a print request at the head of the queue. If a request is already printing, you can put the request on hold to allow the next request to print immediately.

▼ How to Check the Status of Print Requests

Steps

1. **Log in on any system on the network.**
2. **Check the status of printers and print requests.**

   Only the most commonly used options are shown here. For other valid options, see the `lpstat`(1).man page.

   ```
   $ lpstat -o [list] | -u [user-list]
   ```

   `-o` *list* — Shows the status of print requests on a specific printer. *list* can be one or more printer names, printer class names, or print request IDs.

   To specify multiple printer names, class names, and IDs for *list*, use a space or a comma to separate values. If you use spaces, enclose the list of values in quotation marks.

   If you don't specify *list*, the status of print requests sent to all printers is displayed.

   `-u` *user-list* — Shows the status of print requests for a specific user. *user-list* can be one or more user names.

   To specify multiple users with this command, use a space or a comma to separate user names. If you use spaces, enclose the list of names in quotation marks.

   If you don't specify *user-list*, the status of print requests for all users is displayed.

   When used to check the status of print requests, the `lpstat` command displays one line for each print request. From left to right, the line shows the following information:

   - Request ID
   - User
   - Output size in bytes
   - Date and time of the request,
   - Information about the request, such as "being filtered."

Example 4–12 Checking the Status of Print Requests

The following example shows that user `fred` has one print request queued on the printer `luna`.

```
$ lpstat
luna-1    fred     1261    Jul 12 17:34
```

The following example shows that the user `paul` currently has no print requests in queue.

```
$ lpstat -u paul
```

The following example shows that two print requests are queued on the printer `moon`.

```
$ lpstat -o moon
moon-78    root     1024   Jul 14 09:07
moon-79    root     1024   Jul 14 09:08
```

Enabling or Disabling a Printer From Printing

The `enable` and `disable` commands control whether a printer prints or stops printing requests that are in the print queue. When you disable a printer, the printer stops printing requests in the queue. However, requests are still added to the queue. You must set the printer to reject print requests so that requests are not added to the queue. For information about rejecting print requests, see "Accepting or Rejecting Print Requests" on page 109.

A printer is enabled to print and accepts print requests when the printer is added by using Solaris Print Manager. Solaris Print Manager doesn't provide any additional printer processing management.

You must enable the printer whenever it has been disabled, which can happen when a printer fault occurs. When you enable a printer, it prints requests from the print queue until the queue is empty, even if the print service rejects additional requests for the print queue.

The following figure shows the point at which theprocessing of print requests is interrupted when a printer is disabled.

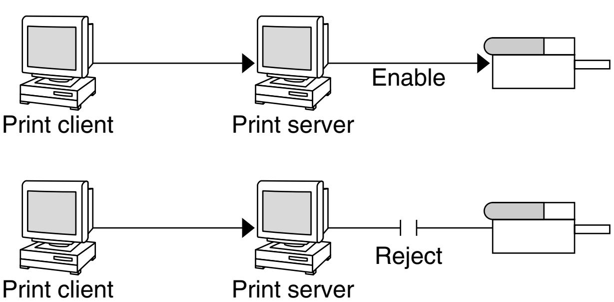

FIGURE 4–1 What Happens When a Printer Is Enabled or Disabled

▼ How to Enable or Disable a Printer

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Stop printing print requests.**

   ```
   # disable [-c | -W] [-r "reason"] printer-name
   ```

`disable`	Cancels the current job, then disables the printer. The current job is saved to reprint when the printer is enabled.
`-c`	Cancels the current job, then disables the printer. The current job is not printed later.
`-W`	Waits until the current job is finished before disabling the printer.
`-r` *"reason"*	Provides users with a reason why the printer is disabled. The reason is stored and displayed whenever a user checks on the status of the printer by using the `lpstat -p` command.
printer-name	Specifies the name of the printer that will stop printing print requests.

   Note – You cannot enable or disable classes of printers. Only individual printers can be enabled or disabled.

   ---

3. **Start printing print requests.**

   ```
   # enable printer-name
   ```

4. **Verify that the printer is enabled.**

   ```
   # lpstat -p printer-name
   ```

Example 4–13 Enabling or Disabling a Printer

The following example shows how to stop the current job on the printer `luna`, save the job to print later, and provide a reason why the printer has stopped printing print requests.

```
# disable -r "changing the form" luna
```

The following example shows how to start printing print requests on the printer `luna`.

```
# enable luna
printer "luna" enabled
```

Accepting or Rejecting Print Requests

The `accept` and `reject` commands enable you to turn on or turn off a print queue that stores requests to be printed.

When you use the `reject` command, the print queue for a specified printer is turned off. No new print requests can enter the queue on the print server. However, all print requests that are in the queue are still printed. You must disable the printer if you want it to stop printing requests that are already in the queue. The following table compares the functions of the `accept`, `reject`, `enable`, and `disable` commands.

TABLE 4–7 Functions of `accept`, `reject`, `enable`, and `disable` Print Commands

Command	Function
`accept`	Accepts print requests that are sent to the print queue.
`disable`	Stops printing requests that are currently in the print queue.
`enable`	Prints the requests that are in the print queue.
`reject`	Rejects print requests that are sent to the print queue.

If a print request is rejected, the LP print service writes or mails a message to the user who submitted the request. the message states that print requests are not being accepted for the specified printer.

You can also specify a reason for not accepting requests through the command line. The reason is displayed on a user's system when a user tries to check the printer's queue. The following figure shows the point at which the processing of print requests is interrupted when a print queue rejects print requests.

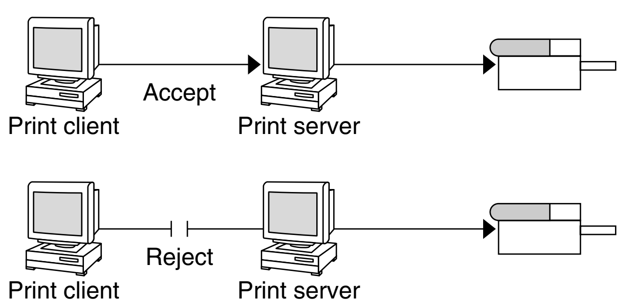

FIGURE 4–2 What Happens When a Print Queue Accepts or Rejects Requests

▼ How to Accept or Reject Print Requests for a Printer

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Stop accepting print requests for the printer.**

   ```
   # reject [-r "reason"] printer-name
   ```

   -r *"reason"* — Provides users with a reason why the printer is rejecting print requests. The reason is stored and displayed whenever a user checks on the status of the printer by using the `lpstat -p` command.

   printer-name — Specifies the name of the printer that will stop accepting print requests.

   The queued requests continue printing as long as the printer is enabled. For instructions on disabling a printer so that it stops printing, see "How to Enable or Disable a Printer" on page 108.

3. **Start accepting print requests for the printer.**

   ```
   # accept printer-name
   ```

4. **Check that the status of the printer to see whether it is accepting or rejecting print requests.**

   ```
   $ lpstat -p printer-name
   ```

Example 4–14 Accepting or Rejecting Print Requests for a Printer

The following example shows how to stop the printer `luna` from accepting print requests.

```
# reject -r "luna is down for repairs" luna
destination "luna" will no longer accept requests
```

The following example shows how to set the printer `luna` to accept print requests.

```
# accept luna
destination "luna" now accepting requests
```

Canceling a Print Request

You can use the `cancel` command to cancel print requests in printer queues or to cancel jobs that are printing. Three ways to use the `cancel` command are as follows:

- Cancel requests by request identification number (request ID)

- Cancel requests from a specific user on all printers or specific printers
- Cancel the job currently printing

When you use the `cancel` command, a message tells you the request or requests are canceled, and the next request in queue is being printed. You can cancel a print request only under the following conditions:

- You are the user who submitted the request and you are logged in on the system from which you submitted the request
- You are the user who submitted the request on any client system and the print server has the "user-equivalence" option configured for the printer in its `/etc/printers.conf` file.
- You are logged in as superuser, `lp`, or have assumed an equivalent role on the print server.

To cancel a specific request, you need to know its request ID. The request ID contains the name of the printer, a dash, and the number of the print request. For example, `luna-185`.

When you submit the print request, the request ID is displayed. If you do not remember the print request ID, you can find it by using the `lpstat` command with the `-o printer` option.

▼ How to Cancel a Print Request

Steps

1. **If you are going to cancel the print requests of other users, then become superuser, `lp`, or assume an equivalent role.**
2. **Determine the request IDs of the print requests to cancel.**

   ```
   # lpstat
   ```

   For more details, see "How to Check the Status of Print Requests" on page 106.
3. **Cancel the print request.**

   ```
   # cancel request-id | printer-name
   ```

   request-id Specifies the request ID of a print request to be canceled. You can specify multiple request IDs with this command. Use a space or a comma to separate request IDs. If you use spaces, enclose the list of request IDs in quotation marks.

   printer-name Specifies the printer for which you want to cancel the currently printing print request.

You can specify multiple printer names with this command. Use a space or a comma to separate printer names. If you use spaces, enclose the list of printer names in quotation marks.

4. **Verify that the print requests are canceled.**

   ```
   $ lpstat -o printer-name
   ```

Example 4–15 Canceling a Print Request

The following example shows how to cancel the `luna-3` and `luna-4` print requests.

```
$ cancel luna-3 luna-4
request "luna-3" cancelled
request "luna-4" cancelled
```

The following example shows how to cancel the print request that is currently printing on the printer `luna`.

```
# cancel luna
request "luna-9" cancelled
```

▼ How to Cancel a Print Request From a Specific User

Steps

1. **(Optional) Become superuser, `lp`, or assume an equivalent role if you are going to cancel the print requests of other users.**

2. **Cancel a print request from a specific user.**

   ```
   $ cancel -u user-list [printer-name]
   ```

   `-u` *user-list* — Cancels the print request for a specified user.

   user-list can be one or more user names. Use a space or a comma to separate user names. If you use spaces, enclose the list of names in quotation marks

   printer-name — Specifies the printer for which you want to cancel the specified user's print requests.

   printer-name can be one or more printer names. Use a space or a comma to separate printer names. If you use spaces, enclose the list of printer names in quotation marks.

   If you don't specify the *printer-name*, the user's print requests will be canceled on all printers.

Example 4–16 Canceling a Print Request From a Specific User

The following example shows how to cancel all the print requests submitted by the user `george` on the printer `luna`.

```
# cancel -u george luna
request "luna-23" cancelled
```

The following example shows how to cancel all the print requests submitted by the user `george` on all printers.

```
# cancel -u george
request "asteroid-3" cancelled
request "luna-8" cancelled
```

Moving a Print Request

If you plan to change the way a printer is used or decide to take a printer out of service, you should set up the LP print service to reject additional print requests. Then, move or cancel any requests that are currently queued to the printer. You can use the `lpmove` command to move individual or all print requests to another local printer.

Request IDs are not changed when you move print requests, so users can still find their requests. Print requests that have requirements, such as a specific file content type or form, that cannot be met by the newly specified printer cannot be moved. These print requests must be canceled.

▼ How to Move Print Requests to Another Printer

To move all print requests from one printer to another printer, you do not need to know the request IDs. However, first determine how many print requests are affected before you move them.

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**
2. **(Optional) Determine if the request IDs of the print requests on the original printer.**

   ```
   # lpstat -o printer-name1
   ```

3. **(Optional) Check to see if the destination printer is accepting print requests.**

   ```
   # lpstat -p printer-name2
   ```

   `-p` *printer-name2* specifies the name of the printer to which you are moving the print requests.

4. **Move all the print requests from the original printer to the destination printer.**

   ```
   # lpmove printer-name1 printer-name2
   ```

printer-name1	Specifies the name of the printer from which all print requests will be moved.
printer-name2	Specifies the name of the printer to which all print requests will be moved.

   For more information, see the `lpmove`(1M) man page.

   If some requests cannot be printed on the destination printer, the requests are left in the original printer's queue. By using request IDs, you can also move specific print requests to another printer by using the `lpmove` command.

5. **(Optional) Start accepting print requests on the original printer.**

   If you move all the print requests to another printer, the `lpmove` command automatically stops accepting print requests for the printer. This step is necessary if you want to begin accepting new print requests for the printer.

   ```
   # accept printer-name1
   ```

6. **Check for any remaining print requests in the original printer's queue.**

   ```
   # lpstat -o printer-name1
   ```

   Ensure that all specified print requests were moved to the destination printer's queue.

   ```
   # lpstat -o printer-name2
   ```

Example 4–17 Moving Print Requests to Another Printer

The following example shows how to move print requests from the printer `luna` to the printer `terra`. Then, the original printer, `luna`, is instructed to resume accepting print requests.

```
# lpmove luna terra
# accept luna
```

Changing the Priority of Print Requests

After a user has submitted a print request, you can change its priority in the print server's queue by doing any of the following:

- Putting any print request on hold if it has not finished printing. Putting a request on hold stops the request, if it is currently printing, and keeps it from printing until you resume printing it. Other print requests go ahead of the on-hold request.

- Moving any print request to the head of the queue, where it will become the next job eligible for printing. If you want a job to start printing immediately, you can interrupt the job that is currently printing by putting it on hold.
- Changing the priority of a job still waiting to be printed or moving the job in the queue. Doing so will move the job ahead of lower priority requests and behind other requests that are the same priority or higher priority.

▼ How to Change the Priority of a Print Request

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server that is holding the print request.**

2. **Determine the request IDs of the print requests whose priority you want to change.**

   ```
   # lpstat
   ```

   For more information, see "How to Check the Status of Print Requests" on page 106.

3. **Change the priority of a print request.**

   ```
   # lp -i request-id -H change-priority
   ```

   `-i` *request-id* — Specifies the request ID of a print request you want to change.

   You can specify multiple request IDs with this command. Use a space or a comma to separate request IDs. If you use spaces, enclose the list of request IDs in quotation marks.

   `-H` *change-priority* — Specifies one of the three ways to change the priority of a print request: `hold`, `resume`, `immediate`.

   For detailed information about valid values for *change-priority*, see Table 4–6.

   You can also use the `lp -q` command to change the priority level of a specified print request. You can change the priority level from 0, the highest priority, to 39, the lowest priority.

Example 4–18 Changing the Priority of a Print Request

The following example shows how to change a print request with the request ID `asteroid-79`, to priority level 1.

```
# lp -i asteroid-79 -q 1
```

CHAPTER 5

Managing Character Sets, Filters, Forms, and Fonts (Tasks)

This chapter provides overview information and step-by-step instructions for setting up and administering character sets, print filters, forms, and fonts.

For information on the step-by-step procedures associated with managing character sets, filters, forms, and fonts, see the following:

- "Managing Character Sets (Task Map)" on page 117
- "Managing Print Filters (Task Map)" on page 126
- "Managing Forms (Task Map)" on page 131
- "Managing Fonts (Task Map)" on page 141

For more information about printing, see the following:

- Chapter 1
- Chapter 2
- Chapter 3
- Chapter 4

Managing Character Sets (Task Map)

Task	Description	For Instructions
Define a print wheel or font cartridge.	Define a print wheel or font cartridge that can be used with a printer.	"How to Set an Alert to Mount a Print Wheel or Font Cartridge" on page 123

Task	Description	For Instructions
Unmount and mount a print wheel or font cartridge.	Delete a current print wheel or font cartridge from the print server's configuration file. Adds a new print wheel or font cartridge to the print server's configuration file.	"How to Unmount and Mount a Print Wheel or Font Cartridge" on page 122
Set up alerts to mount a print wheel or font cartridge.	Set up email alerts for print requests.	"How to Set an Alert to Mount a Print Wheel or Font Cartridge" on page 123
Set up an alias for a selectable character set.	Set up aliases when `terminfo` names for the selectable character sets are inadequate.	"How to Set Up an Alias for a Selectable Character Set" on page 124

Managing Character Sets

Printers differ in the method that they use to print text in various font styles. For example, PostScript printers treat text as graphics. These printers can generate text in different fonts. They can also place the text in any position, size, or orientation on the page. Other types of printers support a more limited number of font styles and sizes. These printers use either print wheels, font cartridges, or preprogrammed selectable character sets. Usually, only one printing method applies to a given printer type.

From the perspective of the LP print service, print wheels and font cartridges are similar. In both cases, someone must intervene and mount the hardware on the printer, when needed. Character sets that require you to physically mount a print wheel or font cartridge are referred to as *hardware character sets*. Character sets that do not require hardware mounting, come preprogrammed with the printer. These character sets can be selected by a print request. They are referred to as *software character sets*.

When you set up a non-PostScript printer, you need to tell the LP print service which print wheels or selectable character sets are available to users. When users submit print requests, the `lp -S` command enables them to specify a print wheel or selectable character set to use for the print job. Users do not have to know which type of character set applies. They just refer to the font style by the name you have defined. For example, you can define a print wheel as `gothic`. To request the `gothic` print wheel, the user would use the following command:

```
% lp -S gothic filename
```

Selectable Character Sets

The selectable character sets supported by a printer are listed in the `terminfo` entry for that printer. For example, the entry for the `ln03` printer is `/usr/share/lib/terminfo/l/ln03`. You can find the names of selectable character sets for any printer type in the `terminfo` database by using the `tput` command. The syntax for the `tput` command is as follows:

`tput -T` *printer-type* `csn`

The `cs`*n* option is an abbreviation for character set number. The number starts with 0, which is always the default character set number after the printer is initialized. To display the names of the other character sets, you can repeat the command by using `-1`, `-2`, `-3`, and so on, in place of the `-0`. For each selectable character set, a `terminfo` name, for example, `usascii`, `english`, `finnish`, and so forth, is returned.

In general, the `terminfo` character set names should closely match the character set names that are used in the manufacturer's documentation for the printer. Because all manufacturers do not use the same character set names, the `terminfo` character set names can differ from one printer type to the next printer type.

You do not have to register the selectable character set names with the LP print service. However, you can give them more meaningful names or aliases.

Note – If you do not specify the selectable character sets that can be used with a printer, the LP print service assumes that the printer can accept any character set name, such as `cs0`, `cs1`, or `cs2`, or the `terminfo` name known for the printer.

Character sets for PostScript printers are not listed when you use the `lpstat -p -l` command. The PostScript fonts are controlled by PostScript filters, not by entries in the `terminfo` database. For information about how to administer PostScript fonts, see "Managing Fonts" on page 141.

Users can use the `lpstat -p -l` command to display the names of the selectable character sets that you have defined for each printer on a print server.

Hardware Character Sets

Another method to obtain alternative character sets is to use removable print wheels or font cartridges that you physically attach, or mount, in a printer.

To administer hardware-mounted character sets, you inform the LP print service of the following:

- The names you want to use for the available print wheels
- How you want to be alerted when a printer needs a different print wheel

Then, when a user requests a particular character set by using the `lp -S` command, the print scheduler sends an alert to mount the print wheel. In addition, the print request is placed in the print queue. When you mount the correct print wheel and tell the LP print service that the print wheel is mounted, the job is printed. For more information, see "How to Unmount and Mount a Print Wheel or Font Cartridge" on page 122.

If you do not specify multiple print wheels or font cartridges for a printer, the LP print service assumes that the printer has a single, fixed print wheel or font cartridge. As such, users cannot specify a special print wheel or font cartridge when using that printer.

Unlike selectable character sets, the names you use for print wheels or font cartridges are not tied to entries in the `terminfo` database. Print wheel names or font cartridge names are used only for the purpose of communicating with the LP print service and its users.

However, the names you choose for print wheels or font cartridges should have meaning to the users. The names should refer to font styles. In addition, the names should be the same across printers that have similar print wheels, font cartridges, or selectable character sets. That way, users can specify a font style (character set), without regard to which printer will be used, or whether a print wheel, font cartridges, or a selectable character set will be used.

You and the printer users should agree on the meanings of print wheel names or font cartridge names. Otherwise, what a user asks for and what you mount might not be the same character set.

Tracking Print Wheels

The procedure for tracking print wheels is similar to the procedure for tracking forms. Some printers, usually letter-quality printers, have removable print heads that provide a particular font or character set. Removable print heads include print wheels and print cartridges. A user can request a named character set. If that character set is not available, the LP print service notifies `root` of the request. The job is stored in the print queue until the print wheel is changed.

Alerts for Mounting Print Wheels or Font Cartridges

You request alerts for mounting print wheels or font cartridges in the same way that you request other alerts from the LP print service. For general information about alerts, see "Setting Up Printer Fault Alerts" on page 97.

▼ How to Define a Print Wheel or Font Cartridge

These procedure shows how to define a print wheel or font cartridge that can be used with the printer.

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Define a print wheel or font cartridge.**

   ```
   # print-server lpadmin -p printer-name -S hard-charset1[,hard-charset2...]
   ```

   `-p` *printer-name* — Specifies the name of the printer for which you are defining a print wheel or font cartridge.

   `-s` *hard-charset* — Specifies the hardware character set name of the print wheel or font cartridge.

   You can specify multiple hardware character sets with this command. Use commas or spaces to separate character set names. If you use spaces, enclose the list of character set names in quotation marks.

   Define names that are meaningful to users. Inform the users of these names.

   The print wheel or font cartridge definition is specified in the print server's `/etc/lp/printers/`*printer-name*`/configuration` file.

3. **Log in as superuser, `lp`, or assume an equivalent role on a print client of the print server.**

4. **Define the same print wheel or font cartridge for the print client.**

   ```
   # print-client lpadmin -p printer-name -S hard-charset1, hard-charset2...
   ```

   In this command, the variables are the same as the variables in Step 2.

   The print wheel or font cartridge definition is added to the print client's `/etc/lp/printers/`*printer-name*`/configuration` file.

5. **Repeat Step 3 and Step 4 for each print client that might need to use the print wheel or font cartridge.**

6. **Verify that the print wheel or font cartridge is defined.**

   You should see the name of the character set and the notation `(mounted)`.

   ```
   # lpstat -p printer-name -l
   ```

Example 5–1 Defining a Print Wheel

The following example shows how to define the `pica` print wheel on the printer `luna`. The print client is named `asteroid`.

```
asteroid# lpadmin -p luna -S pica
```

▼ How to Unmount and Mount a Print Wheel or Font Cartridge

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Unmount the print wheel or font cartridge that is mounted in the printer.**

   ```
   # lpadmin -p printer-name -M -S none
   ```

   `-p` *printer-name* — Specifies the printer on which you are unmounting a print wheel or font cartridge.

   `-M -S none` — Indicates that you want to unmount the current print wheel or font cartridge.

   For more information, see the `lpadmin`(1M) man page.

   The current print wheel or font cartridge is deleted from the print server's `/etc/lp/printers/`*printer-name*`/configuration` file.

3. **Remove the print wheel or font cartridge from the printer.**

4. **Put the new print wheel or font cartridge in the printer.**

5. **Mount the new print wheel or font cartridge.**

   ```
   # lpadmin -p printer-name -M -S hard-charset...
   ```

   `-p` *printer-name* — Specifies the printer on which you are mounting a print wheel or font cartridge.

   `-M -S` *hard-charset* — Specifies the hardware character set name of the print wheel or font cartridge you want to mount.

   The print wheel or font cartridge is added to the print server's `/etc/lp/printers/`*printer-name*`/configuration` file. The mounted print wheel or font cartridge remains active until it is unmounted or until a new print wheel or font cartridge is mounted.

6. **Check the information under the `Print wheels` or `Character set` heading in the output of the following command.**

   You should see the name of the print wheel or character set and the notation `(mounted).`

```
# lpstat -p printer-name -l
```

Example 5–2 Unmounting and Mounting a Print Wheel

The following example shows how to unmount the current print wheel on the printer `luna` and mount the `pica` print wheel.

```
# lpadmin -p luna -M -S none
# lpadmin -p luna -M -S pica
```

▼ How to Set an Alert to Mount a Print Wheel or Font Cartridge

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**
2. **Set an alert to mount a print wheel or font cartridge.**

   ```
   # lpadmin -S hard-charset -A alert [-Q requests] [-W minutes]
   ```

   `-S` *hard-charset* — Specifies the hardware character set name of the print wheel or font cartridge for which you want to set an alert.

   `-A` *alert* — Specifies what kind of alert will occur when a print wheel or font cartridge is requested. For detailed information about the valid values for *alert*, see Table 4–3. Some valid values are `mail`, `write`, and `quiet`.

   If you specify `mail` or `write`, a predefined alert message says to mount the specified print wheel or font cartridge. The message includes the names of one or more printers that have been set up to use such a print wheel or font cartridge.

   `-Q` *requests* — Specifies the number of print requests that require the print wheel or font cartridge that must be in the queue before an alert occurs. If you don't specify this option, only one print request in the queue triggers an alert.

   `-W` *minutes* — Specifies how often (in minutes) the alert occurs If you don't specify this option, the alert is sent one time.

   For more information, see the `lpadmin`(1M) man page.

   The alert is added to the print server's `/etc/lp/pwheels/`*charset-name*`/alert.sh` file.
3. **Verify that the alert has been added for the print wheel or font cartridge by checking the output of the following command.**

   ```
   # lpadmin -S hard-charset -A list
   ```

Alternately, if you have set a low number of print requests to trigger the alert, submit enough print requests to meet the minimum requirements. Ensure that you receive an alert to mount the print wheel or font cartridge.

Example 5–3 Setting an Alert to Mount a Print Wheel

The following example shows how to set email alerts to occur every 5 minutes for the `elite` print wheel when the print queue contains 10 print requests for the `elite`.

```
# lpadmin -S elite -A mail -Q 10 -W 5
```

Example 5–4 Setting an Alert to Mount a Font Cartridge

The following example shows how to set email alerts to occur every minute for the `finnish` font cartridge when the print queue contains 5 print requests for the `finnish`.

```
# lpadmin -S finnish -A mail -Q 5 -W 1
```

Example 5–5 Setting Up No Alerts for a Print Wheel

The following example shows how to set up no alerts for the `elite` print wheel.

```
# lpadmin -S elite -A none
```

▼ How to Set Up an Alias for a Selectable Character Set

You do not need to perform this procedure if the `terminfo` names for the selectable character sets are adequate. For more information on using the `terminfo` database, see “Adding a `terminfo` Entry for an Unsupported Printer” on page 150.

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**
2. **Display the names of the selectable character sets for the specified printer type.**

   ```
   # tput -T printer-type csn
   ```

`-T` *printer-type*	Specifies the printer type found in the `terminfo` database. For information on entries in the `terminfo` database, see “Printer Type” on page 51.
n	Specifies the number (0, 1, 2, 3, 4, 5, and so on) that represents a selectable character set for the specified printer type. The system displays the selectable character set name followed by the prompt symbol. For example, `cs1` could cause the system to

display `english#`.

For more information, see the `tput`(1) man page.

3. **Set up an alias for a selectable character set.**

```
# lpadmin -p printer-name -S select-charset1=alias1,select-charset2=alias2...
```

`-p` *printer-name*	Specifies the printer on which you are setting up aliases for selectable character sets.
`-S`*select-charset*	Specifies the electable character set name. The name can be found in Step 2.
alias	Is the alias for the specified selectable character set. This alias can be used in addition to the selectable character set name. You can set up more than one alias with this command. Use commas or spaces to separate the aliases. If you use spaces, enclose the list of aliases in quotation marks.

The alias is added to the print server's `/etc/lp/printers/`*printer-name*`/configuration` file.

4. **Log in as superuser, `lp`, or assume an equivalent role on a system that is a print client of the print server.**

5. **Set up an alias for the selectable character set.**

```
# lpadmin -p printer-name -S select-charset1=alias1,select-charset2=alias2...
```

In this command, the variables are the same as the variables in Step 3.

The alias is added to the print client's `/etc/lp/printers/`*printer-name*`/configuration` file.

6. **Repeat Step 4 and Step 5 for each print client that might need to use the alias.**

7. **On the print server and print clients, verify that the selectable character set alias is listed in the output of the following command.**

```
# lpstat -p printer-name -l
```

Alternately, submit a print request that uses the alias for the selectable character set and check for output.

Example 5–6 Setting Up an Alias for a Selectable Character Set

The following example shows how to display the names of selectable character sets for the `usascii` selectable character set on the printer `luna`. The printer `luna` is a `ln03` printer type. This example also shows how to specify text as an alias for the `usascii` selectable character set.

```
# tput -T ln03 cs0
usascii# tput -T ln03 cs1
english# tput -T ln03 csn2
finnish# tput -T ln03 csn3
japanese# tput -T ln03 cs4
norwegian#
# lpadmin -p luna -S usascii=text
```

Managing Print Filters (Task Map)

Task	Description	For Instructions
Add a print filter.	Add information about print filters, which are programs that convert the content type of a file to a content type that is acceptable to the destination printer.	"How to Add a Print Filter" on page 128
Delete a print filter.	Delete information about print filters, which are programs that convert the content type of a file to a content type that is acceptable to the destination printer.	"How to Delete a Print Filter" on page 129
View information about a print filter.	Delete information about print filters, which are programs that convert the content type of a file to a content type that is acceptable to the destination printer.	"How to View Information About a Print Filter" on page 130

Managing Print Filters

Print filters are programs that convert the content type of a file to a content type that is acceptable to the destination printer. The LP print service uses filters to complete the following tasks:

- Convert a file from one data format to another format so that it can be printed properly on a specific type of printer

- Handle the special modes of printing, such as two-sided printing, landscape printing, draft-quality, and letter-quality printing
- Detect printer faults and notify the LP print service, so that the print service can alert users and system administrators

Not every print filter can perform all these tasks. Because each task is printer-specific, the tasks can be implemented separately.

The LP print service provides the PostScript filters listed in Table 5–1. The filter programs are located in the `/usr/lib/lp/postscript` directory. For PostScript printing, you usually only need to install the filter programs when setting up a print server. Solaris Print Manager automatically enables the supplied filters. However, if you administer other printers, you might need to administer print filters for them.

Creating Print Filters

To create a new print filter, you must write a print filter program and create a print filter definition. Filters contain input types, output types, and complex options that provide a language to process command-line arguments within the filter. For more information and step-by-step instructions, see "Creating a New Print Filter" on page 157.

Adding, Changing, Deleting, and Restoring Print Filters

Print filters are added, changed, deleted, or restored on the print server only.

You can use the `lpfilter` command to manage the list of available filters. System information about filters is stored in the `/etc/lp/filter.table` file. The `lpfilter` command gets the information about filters to write to the table from filter descriptor files. The filter descriptor files that are supplied (PostScript only) are located in the `/etc/lp/fd` directory. The actual filter programs are located in the `/usr/lib/lp` directory.

The LP print service imposes no fixed limit on the number of print filters you can define. You can remove filters that are no longer used to avoid extra processing by the LP print service. (LP examines all filters to find a filter that works for a specific print request.) If in doubt, do not remove a filter.

As you add, change, or delete filters, you can overwrite or remove some of the original filters provided by the LP print service. You can restore the original set of filters, if necessary. You can also remove any filters you have added.

SunOS software provides a default set of PostScript filters. Solaris Print Manager automatically adds these filters to a print server. Some of the TranScript filters used with `lpd`-based print services have SunOS equivalents, but other TranScript filters do not. The following table lists the default PostScript filters and identifies the equivalent TranScript filters, where applicable.

TABLE 5–1 Default PostScript Filters

PostScript Filter	Action	TranScript Equivalent
`download`	Used to download fonts	
`dpost`	Changes `ditroff` to PostScript	`psdit`
`postdaisy`	Changes `daisy` to PostScript	
`postdmd`	Changes `dmd` to PostScript	
`postio`	Used for serial interface for PostScript printer	`pscomm`
`postior`	Used to communicate with printer	
`postmd`	Changes matrix gray scales to PostScript	
`postplot`	Changes `plot` to PostScript	`psplot`
`postprint`	Changes `simple` to PostScript	`enscript`
`postreverse`	Used to reverse or select pages	`psrev`
`posttek`	Changes `TEK4014` to PostScript	`ps4014`

The SunOS software does *not* provide the following filters:

- `TEX`
- `oscat` (NeWSprint `opost`)
- Enscript

The `postreverse`, `postprint`, `postio`, and `dpost` filters are provided in place of Enscript.

Solaris Print Manager adds the default PostScript filters to a print server. If you have printing needs that are not met by these filters, see "How to Create a New Print Filter" on page 166. This section provides information about writing a custom print filter.

▼ How to Add a Print Filter

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Add a print filter that is based on a print filter definition.**

```
# lpfilter -f filter-name -F filter-definition
```

`-f` *filter-name* Specifies the name you choose for the print filter.

`-F` *filter-definition* Specifies the name of the print filter definition.

For more information, see the `lpfilter`(1M) man page.

The print filter is added in the print server's `/etc/lp/filter.table` file.

3. **Verify that the print filter was added.**

```
# lpfilter -f filter-name -l
```

Example 5–7 Adding a Print Filter

The following example shows how to add the `daisytroff` print filter that has the `daisytroff.fd` print filter definition.

```
# lpfilter -f daisytroff -F /etc/lp/fd/daisytroff.fd
```

▼ How to Delete a Print Filter

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**
2. **Delete the print filter.**

```
# lpfilter -f filter-name -x
```

`-f` *filter-name*

Specifies the name of the print filter to be deleted.

`-x` Deletes the specified filter.

The print filter is deleted from the print server's `/etc/lp/filter.table` file.

3. **Verify that the filter was deleted.**

```
# lpfilter -f filter-name -l
```

You should receive an error indicating that no filter by the specified name exists.

Example 5–8 Deleting a Print Filter

The following example shows how to delete the `daisytroff` print filter.

```
# lpfilter -f daisytroff -x
```

▼ How to View Information About a Print Filter

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**
2. **Request information about a print filter.**

   ```
   # lpfilter -f filter-name -l
   ```

   `-f` Specifies the print filter for which you want to view information. Specify `all` for *filter-name* to view information about all available print filters.

   `-l` Displays information about the specified filter.

   Information about the specified print filter or filters is displayed.

Example 5–9 Viewing Information About a Print Filter

This example shows how to request information for the `postdaisy` print filter. Sample command output follows.

```
# lpfilter -f postdaisy -l
Input types: daisy
Output types: postscript
Printer types: any
Printers: any
Filter type: slow
Command: /usr/lib/lp/postscript/postdaisy
Options: PAGES * = -o*
Options: COPIES * = -c*
Options: MODES group = -n2
Options: MODES group\=\([2-9]\) = -n\1
Options: MODES portrait = -pp
Options: MODES landscape = -pl
Options: MODES x\=\(\-*[\.0-9]*\) = -x\1
Options: MODES y\=\(\-*[\.0-9]*\) = -y\1
Options: MODES magnify\=\([\.0-9]*\) = -m\1
```

Example 5–10 Viewing Information About All Print Filters on a System

This example shows how to display all the print filters that have been added to the system. Sample command output follows.

```
# lpfilter -f all -l | grep Filter
(Filter "download")
Filter type: fast
(Filter "postio")
Filter type: fast
(Filter "postior")
Filter type: fast
(Filter "postreverse")
Filter type: slow
```

Example 5–11 Redirecting Information About a Print Filter

This example shows how to redirect information about the `daisytroff` filter to a file. This task is useful if a filter definition is removed unintentionally.

```
# lpfilter -f daisytroff -l > daisytroff.fd
```

Managing Forms (Task Map)

Task	Description	For Instructions
Limit user access to a form.	Control user access to printers and forms that are available on the network.	"How to Limit User Access to a Form" on page 139
Add a form.	Add the definition of a form to the list of available forms, so that the LP print service can recognize the form.	"How to Add a Form" on page 132
Delete a form.	Remove the definition of a form from the list of available forms that are recognized by the LP print service.	"How to Delete a Form" on page 133
Mount a form.	Use the `mount` command to notify the LP print service to accept print requests for a form with a specific definition.	"How to Unmount and Mount a Form" on page 134
Set an alert to mount a form.	Set alerts for forms that are mounted to a printer.	"How to Set an Alert to Mount a Form" on page 136
View information about a form.	View form names and other form information by listing the contents of the `/etc/lp/forms` directory.	"How to View Information About a Form" on page 138
View the current status of a form.	Use the `lpforms` command to display the current status of a form.	"How to View the Current Status of a Form" on page 138
Limit printer access to a form.	Control what forms printers have access to.	"How to Limit Printer Access to a Form" on page 140

Managing Forms

A *form* is a sheet of paper on which information is printed in a predetermined format. Unlike plain paper stock, forms usually have text or graphics preprinted on them. Common examples of forms are company letterhead, invoices, blank checks, receipts, and labels.

The term *form* has two meanings: the physical medium (the paper) and the software that defines a form to the LP print service.

The LP print service allows you to control the use of forms. This section provides information about adding, changing, deleting, mounting, and limiting access to forms.

Adding, Changing, or Deleting Forms

When you add a form, you tell the LP print service to include the form in its list of available forms. You also have to supply the information required to describe or define the form. Although you can specify such definitions when you add the form, it helps to create the definitions first and save them in files. You can then change the form definition by editing the file.

Note – No form definitions are supplied with the LP print service.

To change a form, you must re-add the form with a different definition.

The LP print service imposes no limit on the number of forms you can define. However, you should delete forms that are no longer appropriate. Obsolete forms can result in unnecessary processing by the print service.

▼ How to Add a Form

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**
2. **Add a form that is based on a form definition.**

   `# `**`lpforms -f`** *form-name* **`-F /etc/lp/forms/`***form*

`-f`	Specifies the name you choose for the form
`-F /etc/lp/forms/` *form*	Specifies the name of the form definition.

   For more information, see the `lpforms`(1M) man page.

The form is added to the print server's /etc/lp/forms/*form-name*/describe file.

3. **Verify that the form was added by checking the output of the following command:**

```
# lpforms -f form-name -l
```

Example 5–12 Adding a Form

The following example shows how to add the medical form that uses the medical.fmd form definition.

```
# lpforms -f medical -F /etc/lp/forms/medical.fmd
```

See Also Before the form can be used, one or more printers must be given access to the form. For more information, see "How to Limit Printer Access to a Form" on page 140.

▼ How to Delete a Form

Steps

1. **Log in as superuser, lp, or assume an equivalent role on the print server.**

2. **Delete the form.**

```
# lpforms -f form-name -x
```

-f Is the form to be deleted.

-x Deletes the specified form

The form is deleted from /etc/lp/forms/*form-name* file.

For more information, see the lpforms(1M) man page.

3. **Verify that the form was deleted.**

```
# lpforms -f form-name -l
```

You should receive an error indicating that a form by the specified name does not exist.

Example 5–13 Deleting a Form

The following example shows how to delete the medical form.

```
# lpforms -f medical -x
```

Mounting Forms

To print a form, you must load the paper in the printer and use a command to *mount* the form. This command notifies the LP print service that print requests submitted to the printer are to be printed using the form definition. If you use one printer for different types of printing, including forms, you should do the following:

- Disable the printer before you load the paper and mount the form.
- Re-enable the printer when the form is ready. Otherwise, the LP print service will continue to print files that do not need the form on the printer.

When you mount a form, ensure that it is aligned properly. If an alignment pattern has been defined for the form, you can request that the pattern print repeatedly after you have mounted the form. Let the pattern print repeatedly until you have adjusted the printer so the alignment is correct.

When you want to change or discontinue using a form on a printer, you must notify the LP print service by unmounting the form.

▼ How to Unmount and Mount a Form

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Stop accepting print requests on the printer on which you are unmounting the current form.**

   ```
   # reject printer-name
   ```

   printer-name is the name of the printer on which you are unmounting a form.

   New print requests (which might not require the form) are not allowed to enter the printer's queue.

3. **Unmount the current form.**

   ```
   # lpadmin -p printer-name -M -f none
   ```

   In this command, the variable *printer-name* is the same variable that is used in Step 2.

   The current form is deleted from the print server's `/etc/lp/printers/`*printer-name*`/configuration` file.

4. **Remove the form paper from the printer.**

5. **Load the new form paper for the next print request.**

6. **Mount the form.**

   ```
   # lpadmin -p printer-name -M -f form-name[-a -o filebreak]
   ```

`-p`	Specifies the printer on which you are mounting a form.
`-M`	Specifies the name of the form to be mounted.
`-a -o filebreak`	Optionally enables you to print a copy of the alignment pattern defined for the form, if this definition exists.

`-p` *printer-name*	Printer on which you are mounting a form.
`-M -f` *form-name*	Name of the form to be mounted.

The specified form is added in the print server's `/etc/lp/printers/`*printer-name*`/configuration` file.

7. **Start accepting print requests on the printer.**

   `#` **`accept`** *printer-name*

   The printer is ready to print the form you just mounted.

8. **Verify that the form has been mounted.**

   `#` **`lpstat -p`** *printer-name* **`-l`**

   Otherwise, submit a print request that requires the new form and check the printer for output.

Example 5–14 Unmounting a Form

The following example shows the process of unmounting the currently mounted form on the printer `luna`.

```
# reject luna
destination "luna" will no longer accept requests
# lpadmin -p luna -M f none
# accept luna
destination "luna" now accepting requests
```

Example 5–15 Mounting a Form

The following example shows the process of mounting the `medical` form on the printer `luna`.

```
# reject luna
destination "luna" will no longer accept requests
# lpadmin -p luna -M f medical -a -o filebreak
# accept luna
destination "luna" now accepting requests
```

Tracking Forms

The LP print service helps you track which forms are mounted on each printer. The print service also notifies you when it cannot find a description it needs to print a form. You are responsible for creating form descriptions, and mounting and unmounting form paper in each printer. You complete this task either as part of setting up a printer or in response to alerts from the LP print service.

Users can specify the form on which they want a job to print. As superuser, you can mount a specific form, then tell the LP print service that the form is available and on which printer it is mounted. Users can submit print requests specifying a particular form. When the LP print service receives the request, it sends an alert message to `root` requesting that you mount the form.

Defining Alerts for Mounting Forms

You request alerts for mounting forms in the same way that you request other alerts from the LP print service. For general information about alerts, see "Setting Up Printer Fault Alerts" on page 97.

▼ How to Set an Alert to Mount a Form

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Set a request alert for mounting a form.**

   `# lpforms -f` *form-name* `-A` *alert* `[-Q` *requests*`] [-W` *minutes*`]`

`-f`	Specifies the form for which you want to set a request alert.
`-A` *alert*	Specifies what kind of alert will occur when a form is requested. See Table 4–3 for detailed information about the valid values for *alert*. Some valid values are `mail`, `write`, and `quiet`. If you choose `mail` or `write`, a predefined alert message says to mount the specified form. The message includes names of one or more printers that have been set up to use the form.
`-Q` *requests*	Specifies how many print requests that require the form must be in the queue to trigger an alert. If you don't specify this option, an alert occurs with just one print request in the queue.
`-W` *minutes*	Specifies how often (in minutes) the alert will occur. If you don't specify this option, the alert is sent one time.

   The request alert is added to the print server's `/etc/lp/forms/`*form-name*`/alert.sh` file.

3. **Verify that the alert has been added for the form.**

   ```
   # lpforms -f form-name -A list
   ```

   Alternately, if you have set a low number of print requests to trigger the alert, submit print requests to meet the minimum requirement. Ensure that you receive an alert to mount the form.

Example 5–16 Setting an Alert to Mount a Form

This example shows how to set email alerts to occur every 5 minutes for the `letterhead` form when 10 print requests for `letterhead` reside in the print queue.

```
# lpforms -f letterhead -A mail -Q 10 -W 5
```

Example 5–17 Setting a Console Window Alert

This example shows how to set console window alerts to occur every 10 minutes for the `letterhead` form when 5 requests for `letterhead`reside in the print queue.

```
# lpforms -f letterhead -A write -Q 5 -W 10
```

Example 5–18 Setting up No Request Alerts

This example shows how to set no request alerts for the `invoice` form.

```
# lpforms -f invoice -A none
```

Checking Forms

When you have defined a form for the LP print service, you can check the form with either of two commands. Which command you choose depends on the type of information you want to check.

- To show the attributes of the form, use the `lpforms` command. You can also redirect the output of the command into a file to save it for future reference.
- To display the current status of the form, use the `lpstat` command. To protect potentially sensitive content, the alignment pattern is not shown.

If you are not sure about the name of an existing form, you can view the form names by listing the contents of the `/etc/lp/forms` directory.

▼ How to View Information About a Form

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Request information about a form.**

   ```
   # lpforms -f form-name -l
   ```

   `-f` Specifies the form for which you want to view information. Specify `all` for *form-name* to view information about all the available forms.

   `-l` Lists the specified form.

   Information about the specified form or forms is displayed.

Example 5–19 Viewing Information About a Form

The following example shows how to display information about the `medical` form.

```
# lpforms -f medical -l
Page length: 62
Page width: 72
Number of pages: 2
Line pitch: 6
Character pitch: 12
Character set choice: pica
Ribbon color: black
Comment:
Medical claim form
```

The following example shows how to redirect the information about the `medical` form to a file. This command creates the form definition for the form. This command is useful if a form definition is removed unintentionally.

```
# lpforms -f medical -l > medical.fmd
```

▼ How to View the Current Status of a Form

Steps

1. **Log in on the print server.**

2. **Request information about the current status of a form.**

   ```
   $ lpstat -f form-name
   ```

   The `-f` option specifies the form for which you want to view the current status. Specify `all` for *form-name* to view the current status of all the forms.

   Information about the current status of the specified form or forms is displayed.

Example 5–20 Viewing the Current Status of a Form

This example shows how to display the status of the `medical` form.

```
$ lpstat -f medical
form medical is available to you
```

Limiting Access to Forms

You can control which printers and which users have access to some or all of the forms available on the network. For example, you might want only the people in the payroll or accounts payable department to be able to print check forms. In addition, you might want the check forms to be available only on certain printers.

For more information, see "How to Limit User Access to a Form" on page 139. To limit printer access to a form, see "How to Limit Printer Access to a Form" on page 140.

▼ How to Limit User Access to a Form

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Allow or deny users access to a form.**

   ```
   # lpforms -f form-name -u allow:user-list | deny:user-list
   ```

`-f`	Specifies the name of the form for which the allow user access lists or the deny user access list is being created.
`-u allow:`*user-list*	Represents users to be added to the allow user access list. Use a comma or a space to separate user login IDs. If you use spaces, enclose the list of IDs in quotation marks. Table 4–5 provides the valid values for *user-list*.
`deny:`*user-list*	Represents users to be added to the deny user access list. Use a comma or a space to separate user login IDs. If you use spaces, enclose the list of IDs in quotation marks. Table 4–5 provides the valid values for *user-list*.

   The specified users are added to the appropriate user access list for the specified form in one of the following files on the print server:

   - `/etc/lp/forms/`*form-name*`/allow`
   - `/etc/lp/forms/`*form-name*`/deny`

3. **Verify that the allow user access list and deny user access list are set up correctly.**

```
# lpforms -f form-name -l
```

Example 5–21 Allowing User Access to a Form

This example shows how to allow only the users `nathan` and `marcia` access to the `check` form.

```
# lpforms -f check -u allow:nathan,marcia
```

Example 5–22 Denying User Access to a Form

This example shows how to deny users `jones` and `smith` access to the `dental` form.

```
# lpforms -f dental -u deny:jones,smith
```

▼ How to Limit Printer Access to a Form

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**
2. **Allow or deny use of forms on a printer.**

   ```
   # lpadmin -p printer-name -f allow:form-list | deny:form-list
   ```

`-p`	Specifies the name of the printer for which the allow forms list or deny forms list is being created.
`-f` allow: *form-list*	Specifies the form name to be added to the allow list. Use a space or a comma to separate multiple form names. If you use spaces to separate form names, enclose the list of form names in quotation marks.
`-f` deny: *form-list*	Is the form name to be added to the deny list. Use a space or a comma to separate multiple form names. If you use spaces to separate form names, enclose the list of form names in quotation marks.

   The specified form or forms are added to the allow forms list or the deny forms list in one of the following files on the print server:

   - `/etc/lp/printers/`*printer-name*`/form.allow`
   - `/etc/lp/printers/`*printer-name*`/form.deny`

3. **Verify that the allow forms list and the deny forms list are set up correctly.**

   ```
   # lpstat -p printer-name -l
   ```

Example 5–23 Allowing Printer Access to a Form

This example shows how to allow the printer luna to access only the medical, dental, and check forms.

```
# lpadmin -p luna -f allow:medical,dental,check
```

Example 5–24 Limiting Printer Access to a Form

This example shows how to deny the printer luna from accessing the medical, dental, and check forms.

```
# lpadmin -p luna -f deny:"medical dental check"
```

Managing Fonts (Task Map)

Task	Description	For Instructions
Install downloaded PostScript fonts.	Permanently download a font if most print requests serviced by the printer use that font. To avoid transmitting fonts across the network, add all the printer-resident fonts.	"How to Install Downloaded PostScript Fonts" on page 144
Install host-resident PostScript fonts.	Install host-resident fonts on a system that is shared by many users. Use this method when numerous fonts are available, or when these fonts are not used by all print requests.	"How to Install Host-Resident PostScript Fonts" on page 144

Managing Fonts

If you have a laser printer, you might need to install and maintain PostScript fonts. For many printers, the fonts are set up as part of the printer installation process.

PostScript fonts are stored in outline form, either on the printer or on a system that communicates with the printer. When a document is printed, the PostScript interpreter generates each character as needed in the appropriate size from the outline description of it. If a font required for a document is not stored on the printer being used, the font must be transmitted to that printer before the document can be printed. This transmission process is called *downloading fonts*.

Fonts are stored and accessed in several ways:

- *Printer-resident fonts* are stored permanently on a printer. These fonts are installed in read-only memory (ROM) on the printer by the manufacturer. If the printer has a disk, you can install fonts on that disk. Most PostScript printers are shipped with 35 standard fonts.
- A *permanently downloaded font* is transmitted to a printer with a PostScript `exitserver` program. A permanently downloaded font remains in printer memory until the printer is turned off. Memory allocated to a downloaded font reduces the memory available on the server for PostScript print requests. Use of an `exitserver` program requires the printer system password and can be reserved for the printer administrator. You should permanently download a font if most print requests serviced by the printer use that font.
- Fonts that are used infrequently or for special purposes can be stored on a user's system. The user can specify these fonts when submitting the print request. The fonts are appended to the print request and transmitted to the printer. When the print request is processed, the space allocated for the font is freed for other print requests.
- *Host-resident fonts* are stored on a system that is shared by many users. The system that stores the fonts can be a print server or a print client. Each user can request fonts in the document to be printed. This method is useful when numerous fonts are available, or when these fonts are not used by all print requests. If the fonts will be used only on printers attached to a print server, they should be stored on the print server. If the fonts will be used by the users on one system and those users can submit requests to multiple printers on a network, the fonts should be stored on the users' system.

  The LP print service provides a special download filter to manage host-resident fonts. The LP print service also supplies `troff` width tables for the 35 standard PostScript fonts that reside on many PostScript printers. For more information about the `troff`program see the `troff`(1) man page

Managing Printer-Resident Fonts

Most PostScript printers come equipped with fonts that are resident in the printer ROM. Some printers have a disk on which additional fonts are stored. When a printer is installed, you should add the list of printer-resident fonts to the font list for that printer. By identifying printer-resident fonts, you prevent fonts from being transmitted unnecessarily across a network. Each printer has its own list of resident fonts, which is contained in this file:

`/etc/lp/printers/`*printer-name*`/residentfonts`

When the printer is attached to a print server, ensure that the list in the `residentfonts` file includes fonts that are on the print server. The file must also include fonts that are available for downloading to the printer.

You must edit the files that contain the list of printer-resident fonts by using a text editor such as `vi`.

Downloading Host-Resident Fonts

When a PostScript document contains a request for fonts not loaded on the printer, the download filter manages this request. The download filter uses PostScript document structuring conventions to determine which fonts to download.

LP print filters are either fast or slow. A *fast filter* quickly prepares a file for printing. Also, a fast filter must have access to the printer while the filter is processing. A *slow filter* takes longer to convert a file. However, a slow filter does not need to access the printer while the filter is processing. An example of a slow filter is a filter that converts an ASCII document to a PostScript document.

The *download filter* is a fast filter. This filter downloads fonts automatically if the fonts are on the print server. The download filter can also be used to send fonts to a print server. To do so, you can create a new filter table entry that calls the download filter as a slow filter by using the `lp -y` command. Alternatively, you can force selection of this filter by changing the input type.

The download filter performs five tasks:

1. The filter searches the PostScript document to determine which fonts are requested. These requests are documented with the following PostScript structuring comments: `%%DocumentFonts:` *font1 font2 …* in the header comments.
2. The download filter searches the list of printer-resident fonts to determine if the requested font must be downloaded.
3. If the font is not resident on the printer, the download filter searches the host-resident font directory table to determine if the requested font is available. This filter performs this search by getting the appropriate file name from the map.
4. If the font is available, the download filter takes the file for that font and appends it to the file to be printed.
5. The download filter sends the font definition file and the source file to the PostScript printer. The *source file* is the file to be printed.

Installing and Maintaining Host-Resident Fonts

Some fonts reside on the host system and are transmitted to the printer as needed for particular print requests. As the administrator, you make PostScript fonts available to all users on a system. To do so, you must know how and where to install these fonts.

Because fonts are requested by name and stored in files, the LP print service keeps a *map file*. A map file shows the correspondence between the names of fonts and the names of the files that contain those fonts. Both the map file and the font list must be updated when you install host-resident fonts.

The fonts available for use with PostScript printers are stored in directories that you create. These directories are called `/usr/share/lib/hostfontdir/`*typeface*`/`*font*. The variable *typeface* is replaced by a name, such as `palatino` or `helvetica`. The variable *font* is replaced by a an attribute name, such as `bold` or `italic`.

▼ How to Install Downloaded PostScript Fonts

Steps

1. **Log in as superuser, `lp` user, or assume an equivalent role on the print server or print client.**

2. **Change directory to the `/etc/lp/printers/`*printer-name* directory.**

   ```
   # cd /etc/lp/printers/printer-name
   ```

   The variable *printer-name* specifies the name of the printer on which you want to install downloaded PostScript fonts.

3. **Create the `residentfonts` file, if this file does not already exist.**

   ```
   # touch residentfonts
   ```

   This file might not exist if you are adding permanently downloaded fonts for the first time.

4. **Edit the `residentfonts` file. Add all the printer-resident fonts and the fonts to be permanently downloaded.**

▼ How to Install Host-Resident PostScript Fonts

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server or print client.**

2. **Create the `hostfontdir` directory.**

   ```
   # cd /usr/share/lib
   ```

   a. **Change to the `/usr/share/libdirectory`.**

      ```
      # cd /usr/share/lib
      ```

   b. **Create the `hostfontdir` directory, if this directory does not already exist.**

      ```
      # mkdir hostfontdir
      ```

c. **Change the permissions on the `hostfontdir` directory to 775.**

```
# chmod 775 hostfontdir
```

3. **Create a directory for a new typeface, if the directory does not already exist.**

```
# mkdir typeface
```

4. **Copy the font file to the appropriate directory.**

```
# cp filename /usr/share/lib/hostfontdir/typeface/font
```

5. **Add to the map file the name of the font and the name of the file in which it resides.**

   a. **Change to the `/usr/share/lib/hostfontdir` directory.**

   b. **Edit the `map` file by using a text editor such as the `vi` editor.**

   Add a one-line entry for each font that you want to add to the file. Put the font name first, followed by a space, followed by the name of the file where the font resides. For example:

```
Palatino-Bold /usr/share/lib/hostfontdir/palatino/bold
```

   c. **Save the file.**

   When an example entry exists in the map file on the appropriate system, users can apply the font (for example, Palatino Bold) to their print jobs. When users submit a print request that contains this font, the LP print service appends a copy of the `/usr/share/lib/hostfontdir/palatino/bold` file to that file before sending it to the printer.

6. **If you are using the `troff` program, you must create new width tables for this font in the standard `troff` font directory.**

CHAPTER 6

Customizing the LP Print Service (Tasks)

This chapter provides overview information and procedures for customizing the LP print service.

For more information about printers, see Chapter 1.

For information on the step-by-step procedures associated with the LP Print Service, see "Customizing the LP Print Service (Task Map)" on page 147.

Customizing the LP Print Service (Task Map)

Task	Description	For Instructions
Add a `terminfo` entry of an unsupported printer.	Add an entry to the `terminfo` database to initialize an unsupported printer and establish required settings.	"How to Add a `terminfo` Entry for an Unsupported Printer" on page 152
Create a new print filter.	Create a new print filter for file types that the LP print service cannot interpret.	"How to Create a New Print Filter" on page 166
Set up a custom printer interface program.	You can furnish a custom printer interface program if you have a printer that is not supported by the standard printer interface program.	"How to Set Up a Custom Printer Interface Program" on page 156

Task	Description	For Instructions
Adjust the printer port characteristics.	Make printer port characteristics the are set by the LP print service compatible with the printer communication settings.	"How to Adjust the Printer Port Characteristics" on page 149
Create a new form definition.	Provide a new form definition by specifying information for nine required printer characteristics.	"How to Create a New Printer Form Definition" on page 170

Adjusting Printer Port Characteristics

The printer port characteristics set by the LP print service must be compatible with the printer communication settings. If the default printer port settings provided by the LP print service do not work with a printer, refer to the printer manual from the manufacturer to determine what settings the printer requires from the LP print service. Use the `stty` command to set and display printer communication settings.

The following table shows the default `stty` settings that are used by the LP print service.

TABLE 6–1 Default `stty` Settings Used by the LP Print Service

Option	Description
`-9600`	Sets baud to 9600
`-cs8`	Sets 8-bit bytes
`-cstopb`	Sends one stop bit per byte
`-parity`	Does not generate parity
`-ixon`	Enables XON/XOFF (also known as START/STOP or DC1/DC3)
`-opost`	Does "output post-processing" using all the settings that follow in this table
`-olcuc`	Do not map lowercase to uppercase
`-onlcr`	Changes linefeed to carriage return/linefeed
`-ocrnl`	Does not change carriage returns into linefeeds
`-onocr`	Outputs carriage returns even at column 0

TABLE 6–1 Default `stty` Settings Used by the LP Print Service *(Continued)*

Option	Description
`-nl0`	Provides no delay after linefeeds
`-cr0`	Provides no delay after carriage returns
`-tab0`	Provides no delay after tabs
`-bs0`	Provides no delay after backspace characters
`-vt0`	Provides no delay after vertical tabs
`-ff0`	Provides no delay after form feeds

▼ How to Adjust the Printer Port Characteristics

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**
2. **Adjust the printer port characteristics.**

   `# lpadmin -p printer-name -o "stty=options"`

   `-p` — Specifies the name of the printer for which you are adjusting the port characteristics.

   `-o` "stty=*options*" — Sets the port characteristic (`stty` option) specified by *options*. You can change more than one `stty` option setting with this command. Enclose each option in single quotation marks, and use a space to separate the options. For a complete list of options, see the`stty`(1) man page. Table 6–1 shows the default `stty` settings used by the LP print service.
3. **Verify that the printer port characteristics have been changed.**

   `# lpstat -p printer-name -l`

Example 6–1 Adjusting the Printer Port Characteristics

This example shows how to set the port characteristics for the printer `luna`. The `parenb` option enables parity checking/generation. The `parodd` option sets odd parity generation. The `cs7` option sets the character size to 7 bits.

```
# lpadmin -p luna -o "stty='parenb parodd cs7'"
```

Example 6–2 Setting the Terminal Baud Rate

This example shows how to set the terminal baud rate to 19200 for the printer `venus`.

```
# lpadmin -p venus -o "stty=19200"
```

Adding a terminfo Entry for an Unsupported Printer

The LP print service uses an interface program and the terminfo database to initialize each printer and establish the following:

- Selected page size
- Character pitch
- Line pitch
- Character set

Each printer is identified in the terminfo database with a short name. The short name required by the terminfo database is identical to the name used to set the TERM shell variable. This name is also the printer type that you specify when setting up a printer. For example, the entries for different types of PostScript printers are in the /usr/share/lib/terminfo/P directory. The default entries provided with the SunOS software release are PS (for PostScript) and PSR (for PostScript Reverse).

If you cannot find a terminfo entry for your printer, you still might be able to use the printer with the LP print service without the automatic selection of page size, pitch, and character sets. However, you might have trouble keeping the printer set in the correct modes for each print request.

If no terminfo entry exists for your type of printer and you want to keep the printer set in the correct modes, you can do one of the following:

- Customize the interface program used with the printer.
- Add an entry to the terminfo database.

A terminal or printer entry in the terminfo database contains and defines hundreds of items. The LP print service, however, uses fewer than 50 of these items. The following table lists the required terminfo items for a printer.

TABLE 6–2 Required terminfo Items for a Printer

Item		Description
Booleans:		
	cpix	Changing character pitch changes resolution
	daisy	Printer requires an operator to change character set
	lpix	Changing line pitch changes resolution
Numbers:		
	bufsx	Number of bytes buffered before printing

TABLE 6–2 Required `terminfo` Items for a Printer *(Continued)*

Item		Description
	`cols`	Number of columns in a line
	`cps`	Average print rate in characters per second
	`it`	Tabs initially every *n* spaces
	`lines`	Number of lines on a page
	`orc`	Horizontal resolution, in units per character
	`orhi`	Horizontal resolution, in units per inch
	`orl`	Vertical resolution, in units per line
	`orvi`	Vertical resolution, in units per inch
Strings:		
	`chr`	Changes horizontal resolution
	`cpi`	Changes number of characters per inch
	`cr`	Carriage return
	`csnm`	List of character set names
	`cudl`	Moves carriage down one line
	`cud`	Moves carriage down *n* lines
	`cuf`	Moves carriage to the right *n* columns
	`cvr`	Changes vertical resolution
	`ff`	Ejects page
	`hpa`	Horizontal position absolute
	`ht`	Tabs to next 8-space tab stop
	`if`	Is the name of initialization file
	`iprog`	Is the path name of initialization program
	`is1`	Is a printer initialization string
	`is2`	Is a printer initialization string
	`is3`	Is a printer initialization string
	`lpi`	Changes number of lines per inch
	`mgc`	Clears all margins (top, bottom, and sides)
	`rep`	Repeats a character *n* times
	`rwidm`	Disables double-wide printing

TABLE 6–2 Required `terminfo` Items for a Printer *(Continued)*

Item	Description
`scs`	Selects character set
`scsd`	Starts definition of a character set
`slines`	Set page length to *n* lines per page
`smgl`	Sets left margin at current column
`smglp`	Set left margin
`smgr`	Sets right margin at current column
`smgrp`	Sets right margin
`smglr`	Sets both left and right margins
`msgt`	Sets top margin at current line
`smgtp`	Sets top margin
`smgb`	Sets bottom margin at current line
`smgbp`	Sets bottom margin
`smgtb`	Sets both top and bottom margins
`swidm`	Enables double-wide printing
`vpa`	Sets vertical position to absolute

▼ How to Add a `terminfo` Entry for an Unsupported Printer

Note – Before you create a `terminfo` entry for a printer, make sure that none of the existing `terminfo` entries will support the printer. To do so, try to set up the printer with an entry for a similar printer, if a similar printer exists.

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**
2. **Determine a `terminfo` entry name for the printer.**

   The directories in the `/usr/share/lib/terminfo` directory contain all the valid `terminfo` entries. Use these entries as a guide for choosing a name for the printer.
3. **Create a `terminfo` entry file for the printer.**

   Table 6–2 shows the items you must define in the `terminfo` entry to add a new printer to the LP print service. For more details about the structure of the `terminfo` database, see the `terminfo`(4) man page.

To help you start writing a new `terminfo` entry, use the `infocmp` command to save an existing `terminfo` entry to a file. This command is helpful if there is a `terminfo` entry that is similar to entry you want to create. For example, the following command saves the `ps` entry to the `ps_cust` file, which will become the new `terminfo` entry.

```
infocmp ps > ps_cust
```

4. **Compile the `terminfo` entry file into the `terminfo` database.**

   ```
   # tic terminfo_entry
   ```

   where *terminfo-entry* variable is the `terminfo` entry file you created.

5. **Check for the new `terminfo` entry file in the `/usr/share/lib/terminfo` directory.**

Customizing the Printer Interface Program

If you have a printer that is not supported by the standard printer interface program, you can furnish your own printer interface program. You can copy the standard program and then tell the LP print service to use it for a specified printer. First, you need to understand what is in the standard program. The following section describes the standard program.

A printer interface program should perform the following tasks:

- Initialize the printer port, if necessary. The standard printer interface program uses the `stty` command to initialize the printer port.
- Initialize the printer hardware. The standard printer interface program receives the control sequences from the `terminfo` database and the `TERM` shell variable.
- Print a banner page, if necessary.
- Print the number of copies that are specified by the print request.

Caution – If you have a printer interface program from a release of UNIX System V prior to release 3.2, it will probably work with the SunOS 5.10 or compatible LP print service. However, several `-o` options have been standardized in the SunOS 5.10 or compatible LP print service. These options will be passed to every printer interface program. These options might interfere with similarly named options used by the old interface.

The LP print service, not a printer interface program, is responsible for opening the printer port. The printer port is given to the printer interface program as standard output. In addition, the printer is identified as the “controlling terminal” for the printer interface program so that a “hang-up” of the port causes a SIGHUP signal to be sent to the printer interface program.

The Standard Printer Interface Program

The standard (model) printer interface program is /usr/lib/lp/model/standard. This program is used by the LP print service to set the printing defaults shown in the following table.

TABLE 6–3 Default Printer Port Characteristics

Characteristic	Default Setting
Default filter	None
Character pitch	None
Line pitch	None
Page width	None
Page length	None
Character set	None
stty options	9600 cs8 -cstopb -parenb -parodd ixon -ixany opost -olcuc onlcr -ocrnl -onocr -onlret -ofill nl0 cr0 tab0 bs0 vt0 ff0
Exit code	0

Customizing stty Modes

If you need to change the terminal characteristics, such as baud rate or output options, find the section of the standard printer interface program that begins with the following comment:

```
## Initialize the printer port
```

For more information, see the “Adjusting Printer Port Characteristics” on page 148.

Exit Codes

When printing is complete, your interface program should exit with a code that shows the status of the print job. The exit code is the last entry in the printer interface program.

The following table shows the exit codes and how they are interpreted by the LP print service.

TABLE 6–4 Printer Interface Program Exit Codes

Exit Code	Meaning to the LP Print Service
`0`	The print request has been successfully completed. If a printer fault occurred, it has been cleared.
`1` to `127`	A problem was encountered when printing a request. For example, there were too many nonprintable characters or the request exceeds the printer's capabilities. The LP print service notifies the person who submitted the request that an error occurred when printing it. This error does not affect future print requests. If a printer fault has occurred, it has been cleared.
`128`	This code is reserved for use by the LP print service. Interface programs must not exit with this code.
`129`	A printer fault was encountered when printing the request. This fault does affect future print requests. If the fault recovery for the printer directs the LP print service to wait for the administrator to correct the problem, the LP print service disables the printer. If the fault recovery is to continue printing, the LP print service does not disable the printer. However, the print service will try printing again in a few minutes.
`>129`	These codes are reserved for use by the LP print service. Interface programs must not exit with codes in this range.

If the program exits with a code of 129, `root` is alerted of a printer fault. The LP print service must also reprint the request from the beginning, after the fault has been cleared. If you do not want the entire request to be reprinted, you can have the interface program send a fault message to the LP print service. (First wait for the fault to be cleared.) When the fault is cleared, the interface program can resume printing the file. When printing is finished, the printer interface program can give a `0` exit code, just as if the fault had never occurred. An added advantage of this approach: The interface program can detect when the fault is cleared automatically. Thus, the administrator does not need to re-enable the printer.

Fault Messages

You can use the `lp.tell` program to send fault messages to the LP print service. This program is referenced by the `LPTELL` shell variable in the standard printer interface code. The program takes standard input and sends it to the LP print service. The LP print service puts standard input into the message that alerts the administrator to the printer fault. If its standard input is empty, the `lp.tell` program does not initiate an alert. For an example of how the `lp.tell` program is used, examine the standard printer interface code immediately after the following comment:

```
# Set up the $LPTELL program to capture fault messages here
```

If you use the special exit code `129` or the `lp.tell` program, the printer interface program does not need to disable the printer itself. The interface program can disable the printer directly. However, doing so overrides the fault-alerting mechanism. Alerts are sent only if the LP print service detects that the printer has a fault, and the special exit code and the `lp.tell` program are its main detection tools.

If the LP print service has to interrupt printing of a file at any time, it kills the interface program with a signal TERM (trap number 15). For more information, see the `kill`(1) and `signal`(3C) man pages. If the printer interface program dies from receipt of any other signal, the LP print service assumes that future print requests will not be affected. As a result, the LP print service continues to use the printer. The LP print service notifies the user who submitted the request that the request has not been finished successfully.

When the interface is first invoked, the signals HUP, INT, QUIT, and PIPE (trap numbers 1, 2, 3, and 13) are ignored. The standard interface changes this default behavior, so that the signals are trapped at appropriate times. The standard interface interprets receipt of these signals as warnings that the printer has a problem. When the standard interface receives a signal, it issues a fault alert.

Using a Customized Printer Interface Program

You can create a customized printer interface program and use it instead of the standard printer interface program on the print server. To do so, you use the `lpadmin` command to register the program with the LP print service for a specific printer.

▼ How to Set Up a Custom Printer Interface Program

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**

2. **Copy the standard printer interface program.**

   ```
   # cp /var/spool/lp/model/standard custom-interface
   ```

3. **(Optional) If you already have a custom printer interface program, go to Step 5**

4. **Change the copy of the standard printer interface program to meet your needs.**

   Refer to the description of the program in "The Standard Printer Interface Program" on page 154 to determine what you need to change.

5. **Set up the custom printer interface program for a specific printer.**

   ```
   # lpadmin -p printer-name -i custom-interface
   ```

   `-p` Specifies the printer that will use the custom printer interface program.

   `-i` Is the name of the custom printer interface program.

The custom printer interface program is registered with the LP print service, and will be used by that printer when users submit print requests.

6. **Verify that the custom printer interface program has been added in the** **`/etc/lp/printers/`***printer-name***`/configuration`** **file.**

Example 6–3 Setting Up a Custom Printer Interface Program

This example shows how to set up a custom printer interface program named `custom` for the printer `luna`.

```
# lpadmin -p luna -i custom
```

Example 6–4 Copying a Custom Printer Interface Program

This example shows how to copy a custom printer interface program from the printer `venus` to the printer `asteroid`.

```
# lpadmin -p asteroid -e venus
```

Creating a New Print Filter

A filter is used by the LP print service each time it has to print a type of file that the printer cannot interpret. Creating a new print filter is not easy, and it usually requires extensive experimentation. The process of defining a new print filter consists of two steps:

- Writing a print filter program
- Creating a print filter definition

A print filter can be as simple or as complex as needed. Filters contain input types, output types, and complex options that provide a language to process command-line arguments within the filter.

If you have nonPostScript printers, you have to create and add print filters as required. First, you need to understand what print filters are and the requirements that must be met by a filter program.

Writing a Print Filter Program

The LP print service provides filter programs in the `/usr/lib/lp/postscript` directory. These filters cover most PostScript printing situations, where the destination printer requires the data to be in PostScript format. A print filter program must be a binary executable file.

Types of Filters

There are two types of print filters: fast filters and slow filters.

Fast filters do not require much processing time to prepare a file for printing. They must have access to the printer when they run. To be capable of detecting printer faults, a print filter must be a fast filter. Any filter that uses the `PRINTER` keyword as a filter option must be installed as a fast filter.

Slow filters require a great deal of processing time to prepare a file for printing. They do not require access to the printer when they run. Slow filters are run in the background so they do not tie up the printer, allowing other files that do not need slow filtering to be printed.

Converting Files

The LP print service uses print filters to convert files from one content type to another content type. You can specify the accepted file content types for each printer. The user specifies the file content type when submitting a print request. Then, the LP print service finds a printer that can print files of that content type. Because many applications can generate files for various printers, this method is often sufficient. However, some applications can generate files that cannot be printed on any available printers.

Each time the LP print service receives a request to print a type of file that is in a format that cannot be accepted directly by a printer, the LP print service tries to match the content type of the print request with the content type of the available (or specified) printer. If there is a match, the file can be sent directly to the printer without filtering. If no match is found, or if the content type specifies that a filter be used, the LP print service tries to match the content type of the file with the input content type of available filters. In addition, the LP print service tries to match the output type of the filter with the content type of the printer. When an appropriate filter is found, the print request is passed through the filter.

Handling Special Printing Modes

A print filter handles special printing modes and requests to print specific pages. A special printing mode is needed to print any characteristics of print requests that require a customized filter. Filters handle the following characteristics:

- Printer type
- Character pitch
- Line pitch
- Page length
- Page width
- Pages to print
- Character set
- Form name
- Number of copies

The LP print service provides default settings for these characteristics. However, a print filter can handle some characteristics more efficiently. For example, some printers can handle multiple copies more efficiently than the LP print service. So, in this case, you can provide a filter for multiple-copy page control.

Detecting Printer Faults

Each printer has its own way of detecting printer faults and transmitting fault signals to the LP print service. The LP print service only checks for hang-ups (loss of carrier) and excessive delays in printing.

Some printers provide good fault coverage and can send a message that describes the reason for a fault. Other printers indicate a fault by using signals other than the signals for loss of carrier signal or shut off of data flow. A filter is required to interpret this additional printer fault information.

A filter can also put a print request on hold, wait for a printer fault to clear, and then resume printing. With this capability, the print request that was interrupted does not need to be reprinted in its entirety. Only a filter that knows the control sequences used by a printer can determine where to break a file into pages. Consequently, only such a filter can find the place in the file where printing should start after a fault is cleared.

When a print filter generates messages, those messages are handled by the LP print service. If alerts are enabled, the LP print service sends alerts to the system administrator. For further information, see "Setting Up Printer Fault Alerts" on page 97.

Requirements for a Print Filter Program

A print filter can be simple or complex, but it has to meet the following requirements:

- The filter should get the contents of a file from its standard input and send the converted file to the standard output.
- A program cannot be used as a filter if it references external files. You might be tempted to use a program such as`troff`, `nroff`, or a similar word processing program as a filter. The LP print service does not recognize references to other files, known as *include files*, from a filter program. Because `troff` and `nroff` programs allow include files, they can fail when used as filters. If the program needs other files to complete its processing. The program t should not be used as a filter.
- The filter should not depend on files that normally would not be accessible to a user. If a filter fails when run directly by a user, it will fail when run by the LP print service.
- A slow filter can send messages about errors in the file to standard error. a fast filter should not. Error messages from a slow filter are collected and sent to the user who submitted the print request.
- If a slow filter dies because it received a signal, the print request is stopped and the user who submitted the request is notified. Likewise, if a slow filter exits with a nonzero exit code, the print request is stopped and the user is notified. The exit codes from fast filters are treated differently.

If you want the filter to detect printer faults, it should also meet the following requirements:

- If possible, the filter should wait for a fault to be cleared before exiting. The filter should also continue to print at the top of the page where printing stopped after the fault is cleared. If you do not want use the continuation feature, the LP print service stops the filter before alerting the administrator.
- The filter should send printer fault messages to its standard error as soon as the fault is recognized. The filter does not have to exit, but can wait for the fault to be cleared.
- The filter should not send messages about errors in the file to standard error. These messages should be included in the standard output, where they can be read by the user.
- The filter should exit with a zero exit code if the file is finished printing (even if errors in the file have prevented it from being printed correctly).
- The filter should exit with a nonzero exit code, only if a printer fault has prevented the filter from finishing a print request.
- When added to the filter table, the filter must be added as a fast filter.

Creating a Print Filter Definition

A print filter definition tells the LP print service about the filter, what print filter program to run, what kind of conversion the print filter does, and so on. A set of filter descriptor files are provided in the `/etc/lp/fd` directory. These files describe the characteristics of the filters (for example, fast filter or slow filter), and point to the filter programs (for example, `/usr/lib/lp/postscript/postdaisy`).

When defining a new print filter, you must create a print filter definition. A print filter definition contains the following information used by the LP print service:

- Name of the filter program to run
- Input types the filter program accepts
- Output types the filter program produces
- Printer types to which the filter program can send jobs
- Names of specific printers to which the filter program can send jobs
- Filter types (either fast or slow)
- Options

You can type the characteristics as direct input to the `lpfilter` command. You also can create a file that specifies the filter's characteristics, and use the file name as input to the `lpfilter` command arguments. Such a file is called a *filter descriptor file* and should be located in the `/etc/lp/fd` directory. These files are not the filters themselves. Rather, these files point to the filters.

Whether you store the information in a file, or type the information directly on the command line, use the following format:

```
Command: command-pathname [options]
Input types: input-type-list
Output types: output-type-list
Printer types: printer-type-list
Printers: printer-list
Filter type: fast or slow
Options: template-list
```

Note – If you provide more than one definition (that is, more than one line) for any filter characteristic other than `Options`, only the second definition is used by the print service.

The information can be arranged in any order, and not all the information is required. When you do not specify values, the values shown in the following table are assigned by default. These values are not very useful, which is why you should specify values.

TABLE 6–5 Default Values for `lpfilter` Command

Item	Default Value
Input types	`any`
Output type	`any`
Printer types	`any`
Printers	`any`
Filter type	`slow`

`lpfilter` Command

Use the full path of the filter program. If there are any fixed options that the program always needs, include them here.

`lpfilter` Input Types Argument

Input types is a list of file content types that the print filter can process. The LP print service does limit the number of input types, but most filters can accept only one input type. Several file types can be similar enough, such that the filter can deal with them. You can use whatever names you prefer, with a maximum of 14 alphanumeric characters and dashes. Do not use underscores as part of the input type name.

The LP print service uses these names to match a filter to a file type, so follow a consistent naming convention. For example, if more than one filter can accept the same input type, use the same name for that input type when you specify it for each filter. Inform your users of these names so that they know how to identify the file type when submitting a file for printing.

`lpfilter` Output Types Argument

Output types is list of file types that the filter can produce as output. For each input type, the filter produces a single output type. The output type can vary, however, from job to job. The name of the output type is restricted to 14 alphanumeric characters and dashes.

The output type names should either match the types of available (local or remote) printers, or match the input types handled by other filters. The LP print service groups filters in a shell pipeline if it finds that several passes by different filters are needed to convert a file. You will unlikely need this level of sophistication, but the LP print service allows it. Try to find a set of filters that takes as input types all the different files the users might want printed, and that converts those files directly into file types the printer can handle.

`lpfilter` Printer Types Argument

Printer types is a list of the types of printers into which the print filter can convert files. For most printers and filters, you can leave this part of the filter definition blank, because this list is identical to the list of output types. However, this list can be different. For example, you could have a printer with a single printer type for purposes of initialization. However, that printer can recognize several different file content types. Essentially, this printer has an internal filter that converts the various file types into a filter type that it can handle. Thus, a filter might produce one of several output types that match the file types that the printer can handle. The print filter should be marked as working with that printer type.

As another example, you might have two different models of printers that are listed as accepting the same file types. Due to slight differences in manufacture, however, one printer deviates in the results it produces. You label the printers as being of different printer types, say A and B, where B is the printer that deviates. You create a filter that adjusts files to account for the deviation produced by printers of type B. Because this filter is needed only for those printer types, you would list this filter as working only on type B printers.

`lpfilter` Printers Argument

A print filter normally can work with all printers that accept its output, so you can usually skip this part of the filter definition.

You might, however, have some printers that are inappropriate for the output that the filter produces. For example, you might want to dedicate one printer for fast turnaround, only sending files that require no filtering to that printer. Other printers of identical type can be used for files that need extensive filtering before they can be printed.

`lpfilter` Filter Type Argument

The LP print service recognizes fast filters and slow filters, as described in "Types of Filters" on page 158.

Slow filters that are invoked by printing modes, by using the `lp -y` command, must be run on the system from which the print request originated. The LP print service cannot pass values for modes to print servers. It can, however, match a file content type (specified after the `-T` option of the `lp` command) to a content type on a print server. Therefore, if you want to activate special modes on a print server, you must specify content types that permit the LP print service to match input types and output types.

`lpfilter` Options Argument

Options specify how different types of information are converted into command-line arguments to the filter command. This information can include specifications from a user (with the print request), the printer definition, and the specifications implemented by any filters that are used to process the request.

Defining Print Filter Options With Templates

There are 13 sources of information for defining print filter options, each of which is represented by a *keyword*. Each option is defined in a *template*. A template is a statement in a filter definition that defines an option to be passed to the filter command, based on the value of one of the filter characteristics

The options specified in a filter definition can include none, all, or any subset of the 13 keywords. In addition, a single keyword can be defined more than once, if multiple definitions are required for a complete filter definition. The following table contains descriptions of the 13 keywords available for defining options in a print filter definition.

TABLE 6–6 Keywords for Print Filter Options

Option Characteristic	Keyword	Patterns	Example
Content type (input)	`INPUT`	*content-type*	`troff`
Content type (output)	`OUTPUT`	*content-type*	`postscript, impress`
Printer type	`TERM`	*printer-type*	`att495`
Printer name	`PRINTER`	*printer-name*	`lp1`
Character pitch	`CPI`	*scaled-decimal*	`10`
Line pitch	`LPI`	*scaled-decimal*	`6`
Page length	`LENGTH`	*scaled-decimal*	`66`
Page width	`WIDTH`	*scaled-decimal*	`80`
Pages to print	`PAGES`	*page-list*	`1-5,13-20`
Character set	`CHARSET`	*character-set*	`finnish`
Form name	`FORM`	*form-name*	`invoice2`
Number of copies	`COPIES`	*integer*	`3`
Special modes	`MODES`	*mode*	`landscape`

A print filter definition can include more than one template. Multiple templates are entered on a single line and separated with commas, or they are entered on separate lines, preceded by the Options: prefix.

The format of a template is as follows: *keywordpattern = replacement*

The *keyword* identifies the type of option being registered for a particular characteristic of the filter.

The *pattern* is a specific option for the keyword.

The *replacement* is what happens when the keyword has the noted value.

For an example of how an option is defined for a particular filter, suppose you want to have the print service scheduler assign print requests to filters following this criteria:

- If the type of `OUTPUT` to be produced by the filter is `impress`, then pass the `-I` option to the filter.
- If the type of `OUTPUT` to be produced by the filter is `postscript`, then pass the `-P` option to the filter.

To specify these criteria, provide the following templates as options to the `lpfilter` command:

```
Options: OUTPUT impress=-I, OUTPUT postscript=-P
```

If the `Options` line becomes too long, put each template on a separate line, as follows:

```
Options: OUTPUT impress=-I
Options: OUTPUT postscript=-P
```

In both templates, the *keyword* is defined as `OUTPUT`. In the first template, the pattern is `impress`, and the value of the *replacement* is `--I`. In the second template, the value of *pattern* is `postscript`, and the value of *replacement* is `-P`.

To determine which values to supply for each type of template (that is, for the *pattern* and *replacement* arguments for each keyword), consider the following:

- The values for the `INPUT` templates come from the file content type that needs to be converted by the filter.
- The values for the `OUTPUT` templates come from the output type that has to be produced by the filter.
- The value for the `TERM` template is the printer type.
- The value for the `PRINTER` template is the name of the printer that will print the final output.
- The values for the `CPI`, `LPI`, `LENGTH`, and `WIDTH` templates come from the user's print request, the form being used, or the default values for the printer.
- The value for the `PAGES` template is a list of pages that should be printed. Typically, the value consists of a list of page ranges separated by commas. Each page range consists of a pair of numbers separated by a dash, or a single number. (For example, 1–5,6,8,10 indicates pages 1 through 5, plus pages 6, 8, and 10.) However, whatever value was given in the `-P` option to a print request is passed unchanged.
- The value for the `CHARSET` template is the name of the character set to be used.

- The value for the FORM template is the name of the form requested by the lp -f command (the command used to submit a print request).
- The value of the COPIES template is the number of copies of the file to print. If the filter uses this template, the LP print service reduces to one the number of copies of the filtered file it prints. This “single copy” includes the multiple copies that are produced by the filter.
- The value of the MODES template comes from the lp -y command. Because a user can specify several -y options, the MODES template might have several values. The values are applied in the left-to-right order given by the user.

The *replacement* option shows how the value of a template should be given to the filter program. This part typically a literal option, sometimes with the placeholder asterisk (*) included to show where the value goes. The *pattern* and *replacement* options also can use the regular expression syntax of the ed command for more complex conversion of user input options into filter options. All regular expression syntax of ed is supported, including the \(... \) and \n constructions. These constructions can be used to extract portions of the *pattern* option for copying into the *replacement* option, and the & option, which can be used to copy the entire *pattern* option into the *replacement* option. For more information, see the ed(1) man page.

Note – If a comma or an equal sign (=) is included in a *pattern* or a *replacement* option, precede it with a backslash (\). A backslash in front of any of these characters is removed when the *pattern* or *replacement* option is used.

▼ How to Create a New Print Filter

Steps

1. **Log in as superuser, lp, or assume an equivalent role on the print server.**
2. **Create a print filter program.**

   For information on print filter programs, see “Writing a Print Filter Program” on page 158. By convention, filter programs for PostScript printers are located in the /usr/lib/lp/postscript directory. You should put programs you create under the /usr/lib/lp directory in a directory of your choosing.
3. **Create a print filter definition.**

   For information on print filter definitions, see “Creating a Print Filter Definition” on page 161. You should save the print filter definition in a text file. By convention, filter definitions are located in the /etc/lp/fd directory and are identified with the .fd suffix.
4. **Add the print filter to a print server.**

   For instructions, see “How to Add a Print Filter” on page 128.

Example 6–5 Creating a New Print Filter

This example shows a print filter definition that is used to convert `N37` or `Nlp` to `simple`.

```
Input types: N37, Nlp, simple
Output types: simple
Command: /usr/bin/col
Options: MODES expand = -x
Options: INPUT simple = -p -f
```

In this example, the print filter program is named `col`. Once you add the new print filter to a print server, a user's print requests are handled as follows:

- When a user types the following command:

  ```
  $ lp -y expand report.doc
  ```

  The print filter program is run with the following arguments to convert the file:

  ```
  /usr/bin/col -x -p -f
  ```

- When a user types the following command:

  ```
  $ lp -T N37 -y expand report.doc
  ```

  The print filter program is run with the following arguments to convert the file:

  ```
  /usr/bin/col -x
  ```

Example 6–6 Creating a New Print Filter by Converting One Input Type to Another Output Type

This example shows a print filter definition that is used to convert from `troff` to a PostScript file.

```
Input types: troff
 Output types: postscript
 Printer types: PS
 Filter type: slow
 Command: /usr/lib/lp/postscript/dpost
 Options: LENGTH * = -l*
 Options: MODES port = -pp, MODES land = -pl
 Options: MODES group \=\([1-9]\) = -n\l
```

In this example, the filter program is named `dpost`. The filter program takes one input type, `troff`, produces a `postscript` output, and works with any printer of type `PS` (PostScript). Users need to give just the abbreviation `port` or `land` when they ask for the paper orientation to be in either portrait mode or landscape mode. Because these options are not intrinsic to the LP print service, users must specify them by using the `lp -y` command.

After you add the new print filter to a print server, print requests will be handled as follows:

- When a user types the following command to submit a `troff` file type for printing on a PostScript printer (type `PS`), with requests for landscape orientation and a page length of 60 lines:

  ```
  $ lp -T troff -o length=60 -y land -d luna ch1.doc
  ```

  The print filter program `dpost` is run with the following arguments to convert the file:

  ```
  /usr/lib/lp/postscript/dpost -l60 -pl
  ```

- When a user enters the following command:

  ```
  $ lp -T troff -y group=4 -d luna ch1.doc
  ```

  The print filter program `dpost` is run with the following arguments to convert the file:

  ```
  /usr/lib/lp/postscript/dpost -n4
  ```

Creating a New Printer Form

When you want to provide a new form, you must define its characteristics by specifying information for nine required printer characteristics (such as page length and page width) as input to the `lpforms` command. The LP print service uses this information to perform the following tasks:

- Initialize the printer so that printing is done properly on the form
- Send reminders to the system administrator about how to handle the form

The form name can be anything you choose, as long as it does not contain more than 14 alphanumeric characters or any underscores. The information must use the following format:

```
Page length: scaled-number
Page width: scaled-number
Number of pages: integer
Line pitch: scaled-number
Character pitch: scaled-number
Character set choice: character-set-name [,mandatory]
Ribbon color: ribbon-color
Comment:
informal notes about the form
Alignment pattern: [content-type] alignment pattern
```

The optional word `[,mandatory]`, means that the user cannot override the character set choice in the form. The *content-type* can be given, although this information is optional, with an alignment pattern. If this value is given, the print service uses it to determine, as necessary, how to filter and print the file.

Aside from following two exceptions, the information can appear in any order.:

- The `Alignment pattern` (which must always be last)
- The `comment` (which must always follow the line with the `Comment:` prompt).

  If the comment contains a line that begins with a key phrase, such as `Page length`, `Page width`, and so on, precede that line with a > character so that the key phrase is not at the beginning of the line. The initial > character is stripped from the comment and is not displayed.

Not all of the information must be given. When you do not specify values, the default values in the following table are assigned. Before running the `lpforms` command, gather the following information about the new form that is described in this table:

TABLE 6–7 Default Values for Printer Forms

Value	Default	Description
Page length	66 lines	The length of the form, or the length of each page in a multipage form. This information can be the number of lines, or the size in inches or centimeters.
Page width	80 columns	The width of the form, in characters, inches, or centimeters.
Number of pages	1	The number of pages in a multipage form. The LP print service uses this number with a print filter (if available) to restrict the alignment pattern to a length of one form. See the description of alignment pattern below. If no filter is available, the LP print service does not truncate the output.
Line pitch	6 lines per inch	A measurement of how close lines appear on the form. Line pitch is also called *leading*. Character pitch is the distance between two lines, from baseline to baseline, measured by either lines per inch or lines per centimeter.
Character pitch	10 characters per inch	A measurement of how close together characters appear on the form. It is the distance between characters, measured by either characters per inch or characters per centimeter.
Character set choice	Any	The character set, print wheel, or font cartridge that should be used when this form is used. Users can choose a different character set for their own print requests when using this form. Or, you can require that only one character set be used.

TABLE 6–7 Default Values for Printer Forms *(Continued)*

Value	Default	Description
Ribbon color	Any	If the form should always be printed using a certain color ribbon, the LP print service can give a mount alert message indicating which color to use.
Comment	(No default)	Any remarks that might help users understand the form. For example, the remarks could indicate the name of the form, its revision, its purpose, or restrictions on its use.
Alignment pattern	(No default)	A sample file that the LP print service uses to fill one blank form. When mounting the form, you can print this pattern on the form to align it properly. You can also define a content type for this pattern so that the print service knows how to print it.

Note – The LP print service does not try to mask sensitive information in the alignment pattern. If you do not want sensitive information printed on sample forms, for example when you align checks, then you should mask the appropriate data. The LP print service keeps the alignment pattern stored in a safe place, where only users logged in as superuser or `lp` user can read it.

When you have gathered the information for the form, specify it as input to the `lpforms` command. You should record this information first in a separate file so that you can edit it before specifying it with the `lpforms` command. You can then use the file as input instead of typing each piece of information separately after a prompt.

▼ How to Create a New Printer Form Definition

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role on the print server.**
2. **Create a form definition file.**

   For a description on creating printer forms, see "Creating a New Printer Form" on page 168. You should save the printer definition in a text file.
3. **Add the form to the LP print service.**

   ```
   # lpadmin -p printer-name -M -f form-name
   ```
4. **Add the form to a print server.**

For instructions, see "How to Add a Form" on page 132.

CHAPTER 7

LP Print Service (Reference)

This chapter provides the following reference information about the LP print service.

Where to Find Printer Tasks

Printer Task	For More Information
Set up printers with Solaris Print Manager and the `lp` commands	Chapter 3
Administer printers with the `lp` commands after the printers are setup	Chapter 4
Manage character sets, filters, forms, and fonts	Chapter 5
Customize the LP print service	Chapter 6

Definition of the LP Print Service

The *LP print service* is a set of software utilities that allows users to print files while users continue to work. Originally, the print service was called the *LP spooler*. *LP* represents line printer, but the meaning now includes many other types of printers, such as laser printers. *Spool* is an acronym for system peripheral operation off-line.

The print service consists of the LP print service software, any print filters you might provide, and the hardware, such as the printer, system, and network connections.

Structure of the LP Print Service

This section describes the directory structure, files and logs of the LP print service.

LP Print Service Directories

The files of the LP print service are distributed among the directories that are shown in the following table.

TABLE 7–1 Directories for the LP Print Service

Directory	Contents
`/usr/bin`	The LP print service user commands.
`/etc/lp`	A hierarchy of LP server configuration files.
`/usr/share/lib`	The `terminfo` database directory.
`/usr/lib/print`	The `lp` conversion scripts, `in.lpd` daemon, and the `printd` daemon. The `printd` daemon transfers all pending jobs in the`/var/spool/print` directory once per minute. When no jobs are remaining to transfer, the `printd` daemon exits.
`/usr/sbin`	The LP print service administrative commands
`/usr/lib/lp`	The `lpsched` program, binary files, PostScript filters, and the `model` directory, which contains the standard printer interface program.

TABLE 7–1 Directories for the LP Print Service *(Continued)*

Directory	Contents
`/var/lp/logs`	LP log files such as `lpsched` *n* (which includes messages from `lpsched`) and `requests` *n* (which includes information about completed print requests).
`/var/spool/lp`	The spooling directory where files are queued for printing.
`/var/spool/print`	The staging area for LP print service client-side requests.

LP Print Service Configuration Files

The `lpsched` daemon stores configuration information in the `/etc/lp` directory, as described in the following table.

Caution – The configuration files listed in this table are private interfaces. These files are subject to change in future releases. You should not build software that relies on these files being in their current locations or that relies on the data being in the format currently used.

TABLE 7–2 Contents of the `/etc/lp` Directory

File	Type	Description
`classes`	Directory	Files identifying classes provided by the `lpadmin -c` command.
`fd`	Directory	Description of existing filters.
`filter.table`	File	Print filter look-up table.
`forms`	Directory	Location to put files for each form. Initially, this directory is empty.
`interfaces`	Directory	Printer interface program files.
`logs`	Link to `/var/lp/logs`	Log files of printing activities.
`model`	Link to `/usr/lib/lp/model`	The standard printer interface program.
`printers`	Directory	Directories for each local printer. Each directory contains configuration information and alert files for an individual printer.
`pwheels`	Directory	Print wheel files or cartridge files.

TABLE 7–2 Contents of the /etc/lp Directory *(Continued)*

File	Type	Description
ppd	Directory	Each local queue that is configured with a PPD file has a copy of the PPD file placed here.

These configuration files serve a function similar to the /etc/printcap file on lpd-based print servers.

Note – You can check the contents of the configuration files, but you should not edit these files directly. Instead, use the lpadmin command to make configuration changes. Your changes are written to the configuration files in the /etc/lp directory. The lpsched daemon administers and updates the configuration files.

The /etc/lp/printers directory has a subdirectory for each local printer that is known to the system. The following example shows the /etc/lp/printers subdirectories of printers sparc1 and luna.

```
$ ls -l /etc/lp/printers
drwxrwxr-x 2 lp lp 512 Jan 23 23:53 luna
drwxrwxr-x 2 lp lp 512 Jan 11 17:50 sparc1
```

The following table describes the files within each printer-specific directory.

File Name	Description
alert.sh	Shell to execute in response to alerts
alert.vars	Alert variables
configuration	Configuration file
users.deny	List of users to whom printer access is denied
comment	Printer description

The configuration file for the printer luna, /etc/lp/printers/luna/configuration, would typically appear as follows:

```
Banner: on: Always
Content types: PS
Device: /dev/term/b
Interface: /usr/lib/lp/model/standard
Printer type: PS
Modules: default
```

The terminfo Database

The /usr/share/lib directory contains the terminfo database directory. This directory contains definitions for many types of terminals and printers. The LP print service uses information in the terminfo database to perform the following tasks:

- Initializes a printer
- Establishes a selected page size, character pitch, line pitch, and character set
- Communicates the sequence of codes to a printer

Each printer is identified in the terminfo database with a short name. For a description of the structure of the terminfo database, see "Printer Type" on page 51. If necessary, you can add entries to the terminfo database, but doing so is tedious and time-consuming For more information, see "Adding a terminfo Entry for an Unsupported Printer" on page 150.

Daemons and LP Internal Files

The /usr/lib/lp directory contains daemons and files used by the LP print service, as described in the following table.

TABLE 7–3 Contents of the /usr/lib/lp Directory

File	Type	Description
bin	Directory	Contains files for generating printing alerts, slow filters, and queue management programs.
lpsched	Daemon	Manages scheduling of LP print requests.
model	Directory	Contains the standard printer interface program.
postscript	Directory	Contains all PostScript filter programs provided by the LP print service. These filters come with descriptor files in the /etc/lp/fd directory. These files tell the LP print service the characteristics of the filters and where to locate them.

LP Print Service Log Files

The LP print service maintains two sets of log files that are described in the following table.

Log File Name	Description
syslogd	Set lpr.debug in /etc/syslog.conf to enable LP print service logging
/var/spool/lp	A list of current requests that are in the print queue
/var/lp/logs/requests	An ongoing history of print requests

Print Queue Logs

The scheduler for each system keeps a log of print requests in the /var/spool/lp/tmp/*system* and /var/spool/lp/requests/*system* directories. Each print request has two files, one file in each directory, that contain information about the request. The information in the /var/spool/lp/requests/*system* directory can be accessed only by superuser or lpr. The information in the /var/spool/lp/tmp/*system* directory can be accessed only by the user who submitted the request, superuser, or lp.

The following example shows the contents of the /var/spool/lp/tmp/starbug directory:

```
$ ls /var/spool/lp/tmp/starbug
5     5-0
# cat 5-0
C 1
D print1
F /etc/profile
P 20
T /etc/profile
t simple
U root
s 0000
v 2
```

These files remain in their directories only as long as the print request is in the queue. Once the print request is finished, the information in the files is combined and appended to the /var/lp/logs/requests file. This file is described in the next section.

Use the information in the /var/spool/lp/logs directory if you need to track the status of a print request that is currently in the queue.

History Log Files

The LP print service records a history of printing services in two log files, lpsched and requests. These log files are located in the /var/lp/logs directory. You can use the information in these log files to diagnose and troubleshoot printing problems. An example of the contents of the /var/lp/logs directory is as follows:

```
# cd /var/lp/logs
# ls
lpsched.1    requests    requests.2
lpsched      lpsched.2   requests.1
#
```

The two most important log files for troubleshooting are the following:

- The lpsched log file — Contains information about local printing requests.
- The requests log file — Contains information about print requests that are completed and no longer in the print queue.

The requests log file has a simple structure so that you can extract data using common UNIX shell commands. Requests are listed in the order they are printed. They are also separated by lines showing their request IDs. Each line below the separator line, the line that starts with =, is marked with a single letter that identifies the kind of information contained in that line. Each letter is separated from the data by a single space.

The following example shows the contents of a requests log file:

```
# pwd
/var/lp/logs
# tail requests.2
= print1-3, uid 0, gid 1, size 206662, Wed Mar 14 08:56:30 MST 2003
z print1
C 1
D print1
F /usr/dict/words
P 20
T /usr/dict/words
t simple
U root
s 0x0014
v 2
#
```

The following table shows the letter codes and the content of their corresponding lines in the requests log file.

TABLE 7–4 Codes in the `requests` Log File

Letter	Content of Line
`=`	The separator line, which contains the following items: request ID, user ID (UID), and group IDs (GIDs) of the user, the total number of bytes in the original (unfiltered) file size, and the time when the request was queued.
`z`	The name of the printer.
`C`	The number of copies printed.
`D`	The printer or class destination, or the word `any`.
`F`	The name of the file printed. The line is repeated for each file printed. Files were printed in the order shown.
`f`	(Optional) The name of the form used.
`H`	(Optional) One of three types of special handling: resume, hold, and immediate.
`N`	(Optional) The type of alert used when the print request was successfully completed. The type is the letter `M` if the user was notified by email or `W` if the user was notified by a message to the terminal.
`O`	(Optional) The printer-dependent `-o` options. For example, `nobanner`.
`P`	The priority of the print request.
`p`	The list of pages printed.
`r`	(Optional) A single-letter line that is included if the user asked for "raw" processing of the files with the `lp -r` command.
`S`	(Optional) The character set, print wheel, or cartridge used.
`T`	The title placed on the banner page.
`t`	The type of content found in the files.
`U`	The name of the user who submitted the print request.
`s`	The outcome of the request, shown as a combination of individual bits expressed in hexadecimal format. Several bits are used internally by the print service. The bits and what they mean are described in the following table.
`v`	■ 0 — Identifies that the print job originated from the `in.lpd` print daemon or another system on the network. ■ 2 — Identifies that the print job originated from the `/usr/bin/lp` or `/usr/ucb/lpr` commands. ■ 1 — Identifies that the print job originated from an older version of the print software.
`x`	(Optional) The slow filter used for the print request.

TABLE 7–4 Codes in the `requests` Log File *(Continued)*

Letter	Content of Line
`Y`	(Optional) The list of special modes for the print filters used to print the request.
`z`	(Optional) The printer used for the request. This printer differs from the destination (the `D` line) if the request was queued for any printer or a class of printers, or if the request was moved to another destination.

The following table shows the outcome codes in the LP `requests` log file and their descriptions.

TABLE 7–5 Outcome Codes in the `requests` Log File

Outcome Code	Description
`0x0001`	The request was held pending resume.
`0x0002`	Slow filtering is running.
`0x0004`	Slow filtering finished successfully.
`0x0008`	The request is on the printer.
`0x0010`	Printing finished successfully.
`0x0020`	The request was held pending user change.
`0x0040`	The request was canceled.
`0x0080`	The request will print next.
`0x0100`	The request failed filtering or printing.
`0x0200`	The request is in transit to a remote printer (obsolete).
`0x0400`	The user will be notified.
`0x0800`	A notification is running.
`0x1000`	A remote system has accepted the request (obsolete).
`0x2000`	The administrator placed a hold on the request.
`0x4000`	The printer had to change filters.
`0x8000`	The request is temporarily stopped.

Spooling Directories

Files queued for printing are stored in the `/var/spool/lp` directory until they are printed, which might be only seconds. The following table shows the contents of the `/var/spool/lp` directory.

TABLE 7–6 Contents of the `/var/spool/lp` Directory

File	Type	Description
`SCHEDLOCK`	File	Lock file for the scheduler. Check for this file if the scheduler terminates and will not restart.
`admins`	Directory	Link to `/etc/lp.`
`bin`	Directory	Link to `/usr/lib/lp/bin.`
`logs`	Link	Link to `../lp/logs` where completed print requests are logged.
`model`	Link	Link to `/usr/lib/lp/model.`
`requests`	Directory	Directory that contains subdirectories for each configured printer where print requests are logged until printed. Users cannot access this log.
`system`	Directory	A print status file for the system.
`temp`	Link	Link to `/var/spool/lp/tmp/`*hostname*, which contains the spooled requests.
`tmp`	Directory	Directory for each configured printer where print requests are logged until printed. Changes to existing print requests are also recorded in this directory.

LP Print Service Commands

The following table lists frequently used LP print service commands. You must be superuser, `lp`, or assume an equivalent role to use the `1M` commands.

TABLE 7–7 Quick Reference to LP Print Service Commands

Command	Task	Man Page
`enable`	Activate a printer	`enable`(1)
`cancel`	Cancel a print request	`cancel`(1)
`lp`	Send one or more file or files to a printer	`lp`(1)
`lpstat`	Report the status of the LP print service	`lpstat`(1)
`disable`	Deactivate on or more printers	`enable`(1)

TABLE 7–7 Quick Reference to LP Print Service Commands *(Continued)*

Command	Task	Man Page
`accept`	Permit print requests to be queued for a specific destination	`accept`(1M)
`reject`	Prevent print requests from being queued for a specific destination	`accept`(1M)
`lpadmin`	Set up or change a printer configuration	`lpadmin`(1M)
`lpfilter`	Set up or change filter definitions	`lpfilter`(1M)
`lpforms`	Set up or change preprinted forms	`lpforms`(1M)
`lpadmin`	Mount a form	`lpadmin`(1M)
`lpmove`	Move output requests from one destination to another destination	`lpmove`(1M)
`lpsched`	Start the LP print service scheduler	`lpsched`(1M)
`lpshut`	Stop the LP print service scheduler	`lpshut`(1M)
`lpusers`	Set or change the default priority and priority limits that can be requested of the LP print service by users	`lpusers`(1M)

Functions of the LP Print Service

The LP print service performs the following functions:

- Administers files and schedules local print requests
- Receives and schedules network requests
- Filters files, if necessary, so they print properly
- Starts programs that interface with the printers
- Tracks the status of print jobs
- Tracks forms that are mounted on the printer
- Tracks print wheels that are currently mounted
- Delivers alerts to mount new forms or different print wheels
- Delivers alerts about printing problems

How LP Administers Files and Schedules Local Print Requests

The LP print service has a scheduler daemon called `lpsched`. The scheduler daemon updates the LP system files with information about printer setup and configuration.

The `lpsched` daemon schedules all local print requests on a print server, as shown in the following figure. Users can issue the requests from an application or from the command line. Also, the scheduler tracks the status of printers and filters on the print server. When a printer finishes a request, the scheduler schedules the next request n the queue on the print server, if a next request exists.

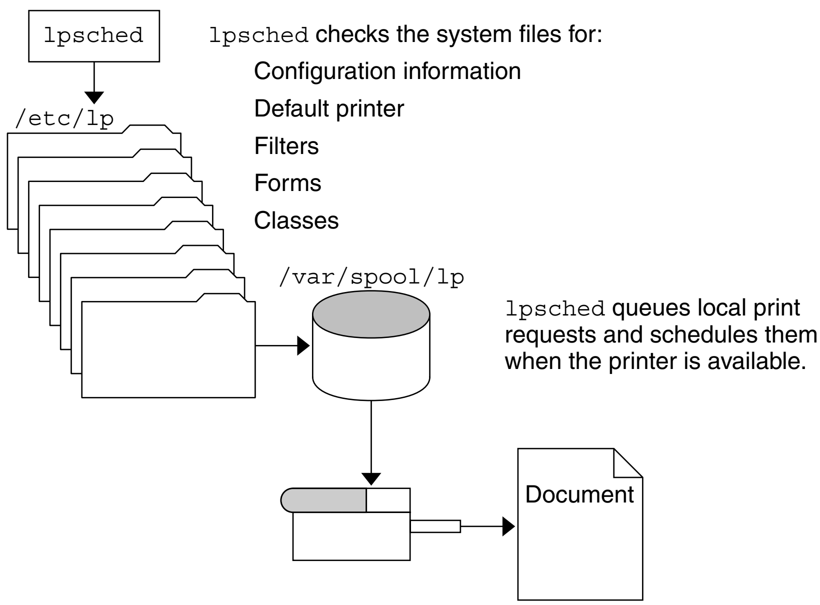

FIGURE 7–1 The `lpsched` Daemon Schedules Local Print Requests

Without rebooting the system, you can stop the scheduler with the `svcadm disable application/print/server` command. Then, restart the scheduler with the `svcadm enable application/print/server` command. The scheduler for each system manages requests that are issued to the system by the `lp` command.

Scheduling Network Print Requests

Each print client communicates directly with a print sever over the network. The communication is done between the requesting command, such as `lp`, `lpstat`, `cancel`, `lpr`, `lpq`, or `lprm`, and the print service on the print server. This communication reduces the print system overhead on client–only systems, improving scalability, performance and accuracy of data.

Print servers listen for print requests with the Internet services daemon (`inetd`). Upon hearing a request for print service from the network, the `inetd` daemon starts a program called the *protocol adaptor* (`in.lpd`). The protocol adaptor translates the print request and communicates it to the print spooler, and returns the results to the requester. This protocol adaptor starts on demand and exits when it has serviced the network request. This process eliminates idle system overhead for printing. This process also eliminates any additional system configuration for network printing support as was the case in previous versions of Solaris printing software.

Filtering Print Files

Print filters are programs on the print server that convert the content of a queued file from one format to another format.

A print filter can be as simple or as complex as needed. The Solaris OS provides print filters in the `/usr/lib/lp/postscript` directory that cover most situations where the destination printer requires the data to be in PostScript format. If you need filters for nonPostScript printers, you have to create the filters and add them to the systems that need filters.

A set of *print filter descriptor files* are provided in the `/etc/lp/fd` directory. These descriptor files describe the characteristics of the filter (for example, fast or slow filter). These description files point to the filter program (for example, to `/usr/lib/lp/postscript/postdaisy`).

What the Printer Interface Program Does

The LP print service interacts with other parts of the Solaris OS. The print service uses a standard printer interface program to do the following:

- Initialize the printer port, if necessary. The standard printer interface program uses the `stty` command to initialize the printer port.
- Initialize the printer. The standard printer interface program uses the `terminfo` database and the `TERM` shell variable to find the appropriate control sequences.
- Print a banner page, if necessary.
- Print the correct number of copies specified by the print request.

The LP print service uses the standard interface program, found in the `/usr/lib/lp/model` directory, unless you specify a different program. You can create custom interface programs> However, you must make sure that the custom program does not terminate the connection to the printer or interfere with proper printer initialization.

How the `lpsched` Daemon Tracks the Status of Print Requests

The `lpsched` daemon keeps a log file of each print request that it processes and notes any errors that occur during the printing process. This log file is kept in the `/var/lp/logs/lpsched` file. Every night, the `lp cron` job renames the `/var/lp/logs/lpsched` file to a new `lpsched.`*n* file and starts a new log file. If errors occur or jobs disappear from the print queue, you can use the log files to determine what the `lpsched` daemon has done with a printing job.

Cleaning Out Log Files

The `lpsched` and `requests` log files in the `/var/lp/logs` directory grow as information is appended. The LP print service uses a default `cron` job to clean out the log files. The `lp cron` job is located in the `/var/spool/cron/crontabs/lp` file. The `cron` job periodically moves the contents of the log files. The contents of *log* are moved to *log*`.1`, and the contents of *log*`.1` are moved to *log*`.2`. The contents of *log*`.2` are lost (that is, replaced by the former contents of *log*`.1`) when *log*`.2` gets overwritten.

▼ How to Change Frequency of Printer Request Log Rotation

The `requests` log file on the printer server is rotated weekly rather than daily. You can change the rotation interval back to daily if the printer server is busy.

Steps

1. **Become superuser or `lp`, or assume an equivalent role on the print server.**

2. **Set the `EDITOR` environment variable.**

```
# EDITOR=vi
# export EDITOR
```

3. **Edit the `lp` `crontab` file.**

```
# crontab -e lp
```

4. **Change the first line of the file that rotates the `requests` log files every Sunday (`0`) to an asterisk (`*`) for daily rotation:**

```
13 3 * * * 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
```

5. **Save the file and exit.**

How Local Printing Works

The print request follows the same path it would if the client and server were separate systems. Requests always flow from client to server following the same path.

The following figure shows what happens when a user submits a request to print a PostScript file on a *local* printer, which is a printer connected to the user's system. The `lpsched` daemon on the local system does all processing. This processing might include the following:

- Matching the printer and content type
- Identifying the default printer, filtering the request
- Starting the specified printer interface program

The printer interface program does the following:

- Prints the banner page.
- Catches printer faults.
- Determines the printer fault policy to either reset, continue, or begin the print job over again.
- Finally, the interface program uses the `lpcat` program to download the print request to the local printer's serial port.

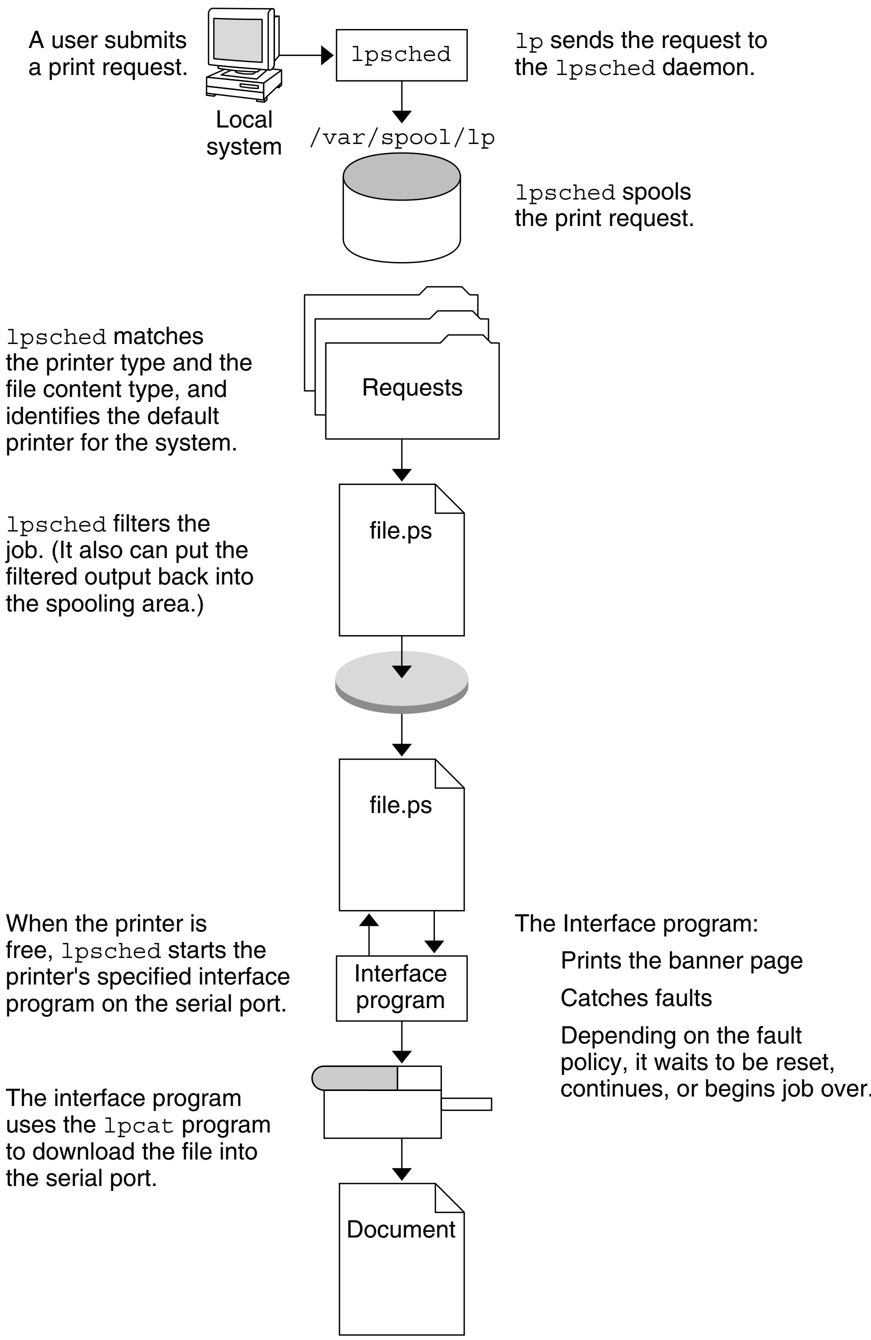

FIGURE 7–2 The Local Printing Process

How Remote Printing Works

The following figure shows what happens when a user on a Solaris print client submits a print request to an lpd-based print server. The command opens a connection and handles its own communications with the print server directly.

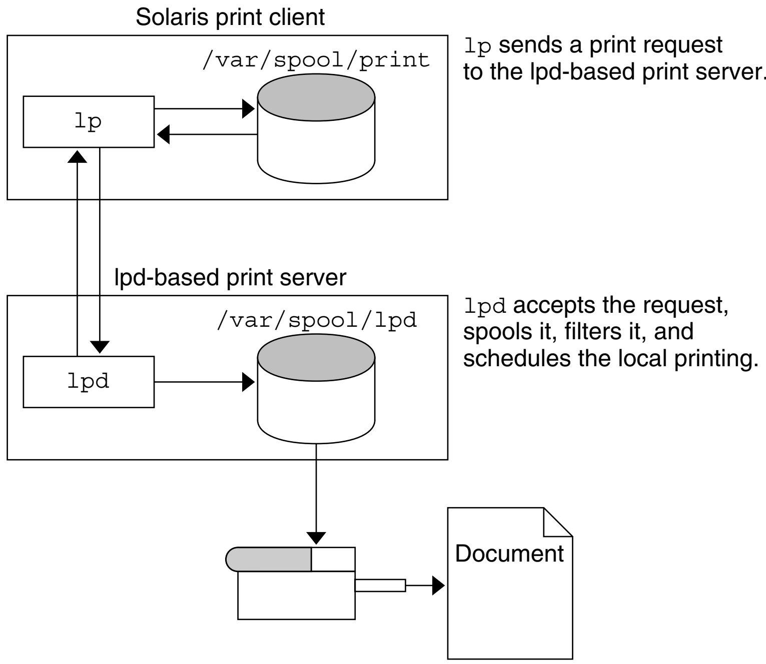

FIGURE 7–3 Printing Between a Solaris Print Client and a `lpd`-based Print Server

The following figure shows an `lpd`-based print client submitting a print request to a Solaris print server. The `lpd` daemon handles the local part of the print request and the connection to the print server. On the print server, the `inetd` process waits for network printing requests and starts a protocol adaptor to service the request. The protocol adaptor communicates with the `lpsched` daemon, which processes the request on the print server.

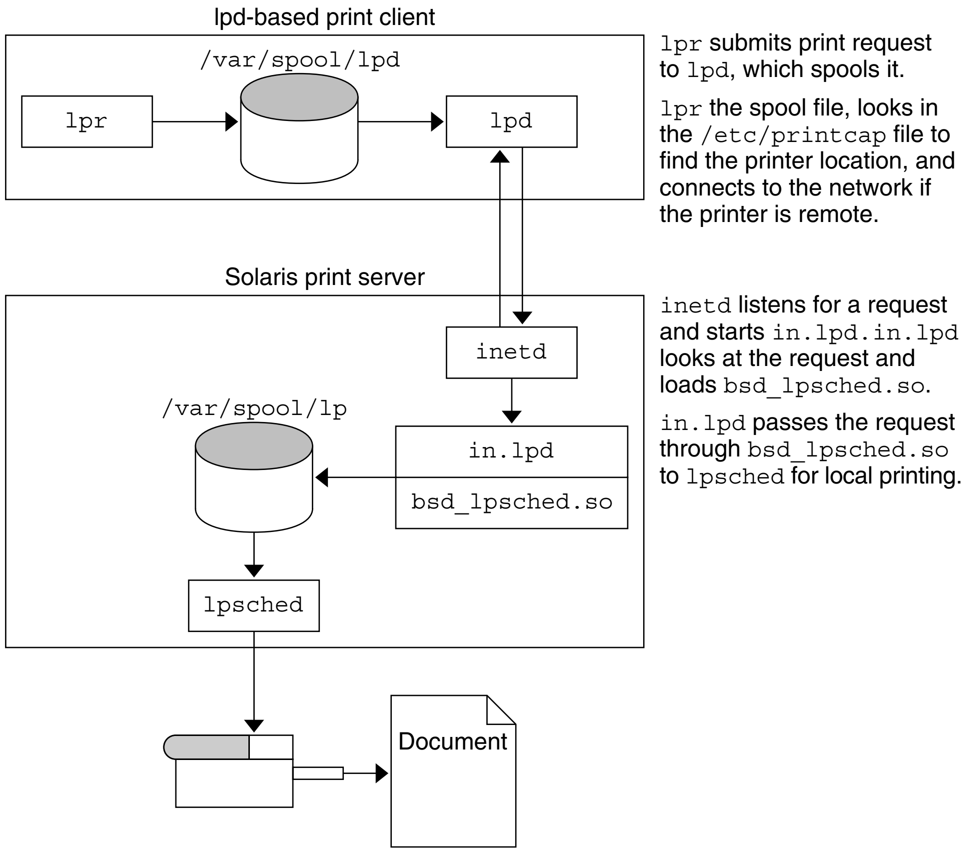

FIGURE 7–4 Printing Between an `lpd`-based Print Client and a Solaris Print Server

The following figure shows what happens when a user on a Solaris print client submits a print request to a Solaris print server. The print command on the print client handles the local part of each print request by communicating directly with the print server.

The `inetd` process on the print server monitors network printing requests and starts a protocol adaptor to communicate with the `lpsched` daemon on the print server, which processes the print request.

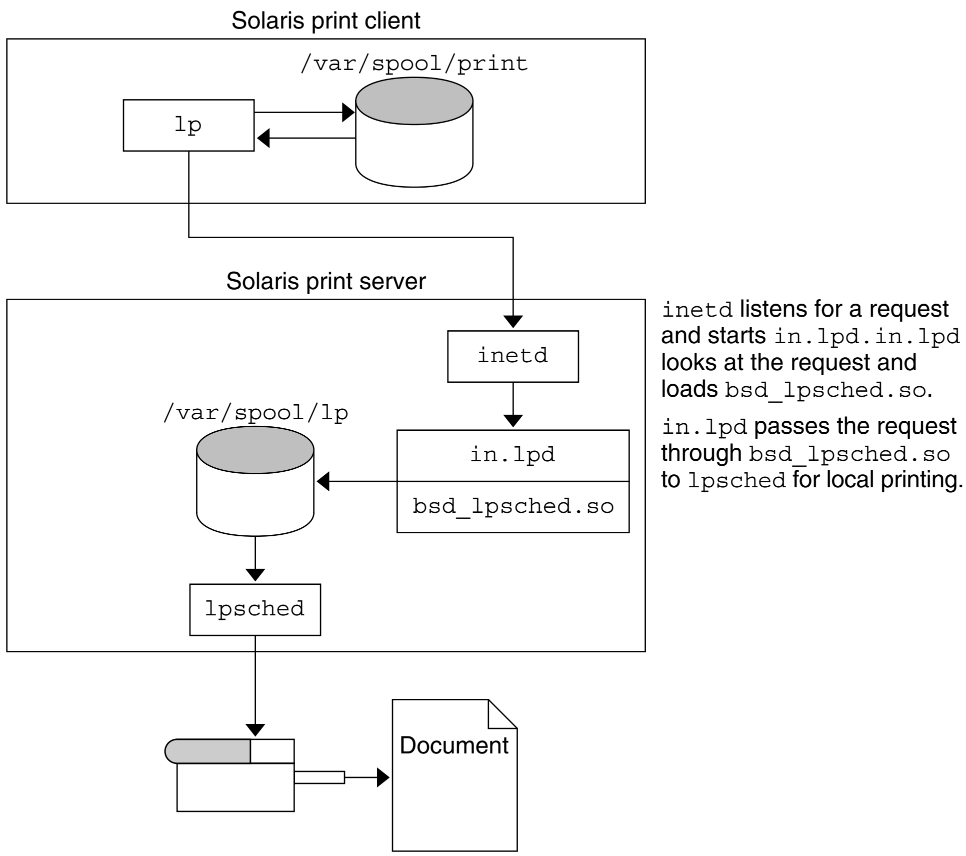

FIGURE 7–5 Printing Between a Solaris Print Client and a Solaris Print Server

CHAPTER 8

Managing Terminals and Modems (Overview)

This chapter provides overview information for managing terminals and modems. This is a list of the overview information in this chapter:

For step-by-step instructions on how to set up terminals and modems with the Serial Ports tool, see Chapter 9.

For step-by-step instructions on how to set up terminals and modems with the Service Access Facility (SAF), see Chapter 10.

Terminals, Modems, Ports, and Services

Terminals and modems provide both local and remote access to system and network resources. Setting up terminals and modem access is an important responsibility of a system administrator. This section explains some of the concepts behind modem and terminal management in the Solaris Operating System.

Terminals

Your system's bit-mapped graphics display is not the same as an alphanumeric terminal. An alphanumeric terminal connects to a serial port and displays only text. You don't have to perform any special steps to administer the graphics display.

Modems

Modems can be set up in three basic configurations:

- Dial-out
- Dial-in
- Bidirectional

A modem connected to your home computer might be set up to provide *dial-out* service. With dial-out service, you can access other computers from your own home. However, nobody outside can gain access to your machine.

Dial-in service is just the opposite. Dial-in service allows people to access a system from remote sites. However, it does not permit calls to the outside world.

Bidirectional access, as the name implies, provides both dial-in and dial-out capabilities.

Ports

A *port* is a channel through which a device communicates with the operating system. From a hardware perspective, a port is a “receptacle” into which a terminal or modem cable might be physically connected.

However, a port is not strictly a physical receptacle, but an entity with hardware (pins and connectors) and software (a device driver) components. A single physical receptacle often provides multiple ports, allowing connection of two or more devices.

Common types of ports include serial, parallel, small computer systems interface (SCSI), and Ethernet.

A *serial port*, using a standard communications protocol, transmits a byte of information bit-by-bit over a single line.

Devices that have been designed according to RS-232-C or RS-423 standards, this include most modems, alphanumeric terminals, plotters, and some printers. These devices can be connected interchangeably, using standard cables, into serial ports of computers that have been similarly designed.

When many serial port devices must be connected to a single computer, you might need to add an *adapter board* to the system. The adapter board, with its driver software, provides additional serial ports for connecting more devices than could otherwise be accommodated.

Services

Modems and terminals gain access to computing resources by using serial port software. Serial port software must be set up to provide a particular "service" for the device attached to the port. For example, you can set up a serial port to provide bidirectional service for a modem.

Port Monitors

The main mechanism for gaining access to a service is through a *port monitor*. A port monitor is a program that continuously monitors for requests to log in or access printers or files.

When a port monitor detects a request, it sets whatever parameters are required to establish communication between the operating system and the device requesting service. Then, the port monitor transfers control to other processes that provide the services needed.

The following table describes the two types of port monitors included in the Solaris Operating System.

TABLE 8–1 Port Monitor Types

Man Page	Port Monitor	Description
`listen`(1M)	`listen`	Controls access to network services, such as handling remote print requests prior to the Solaris 2.6 release. The default Solaris Operating System no longer uses this port monitor type.
`ttymon`(1M)	`ttymon`	Provides access to the login services needed by modems and alphanumeric terminals. The Serial Ports tool automatically sets up a `ttymon` port monitor to process login requests from these devices.

You might be familiar with an older port monitor called `getty`. The new `ttymon` port monitor is more powerful. A single `ttymon` port monitor can replace multiple occurrences of `getty`. Otherwise, these two programs serve the same function. For more information, see the `getty`(1M) man page.

Tools for Managing Terminals and Modems

The following table lists the tools for managing terminals and modems.

TABLE 8–2 Tools For Managing Terminals and Modems

Managing Terminals and Modems Method	Tool	For More Information
The most comprehensive	Service Access Facility (SAF) commands	"Service Access Facility" on page 196
The quickest setup	Solaris Management Console's Serial Ports tool	Chapter 9 and Solaris Management Console online help

Serial Ports Tool

The Serial Ports tool sets up the serial port software to work with terminals and modems by calling the pmadm command with the appropriate information. The tool also provides the following:

- Templates for common terminal and modem configurations
- Multiple port setup, modification, or deletion
- Quick visual status of each port

Service Access Facility

The SAF is the tool used for administering terminals, modems, and other network devices. In particular, the SAF enables you to set up the following:

- ttymon and listen port monitors by using the sacadm command
- ttymon port monitor services by using the pmadm and ttyadm commands
- listen port monitor services by using the pmadm and nlsadmin commands
- Troubleshoot tty devices
- Troubleshoot incoming network requests for printing service
- Troubleshoot the Service Access Controller by using the sacadm command

The SAF is an open-systems solution that controls access to system and network resources through tty devices and local-area networks (LANs). The SAF is not a program, but a hierarchy of background processes and administrative commands.

CHAPTER 9

Setting Up Terminals and Modems (Tasks)

This chapter provides step-by-step instructions for setting up terminals and modems using Solaris Management Console's Serial Ports tool.

For overview information about terminals and modems, see Chapter 8. For overview information about managing system resources, see Chapter 11.

For information about the procedures associated with setting up terminals and modems using Solaris Management Console's Serial Ports tool, see "Setting Terminals and Modems (Task Map)" on page 197

Setting Terminals and Modems (Task Map)

Task	Description	For Instructions
Set up a terminal.	Set up a terminal by using the Solaris Management Console Serial Ports tool. Configure the terminal by choosing the appropriate option from the Action menu.	"How to Set Up a Terminal" on page 201

Task	Description	For Instructions
Set up a modem.	Set up a modem by using the Solaris Management Console Serial Ports tool. Configure the modem by choosing the appropriate option from the Action menu.	"How to Set Up a Modem" on page 202
Initialize a port.	To initialize a port, use the Solaris Management Console Serial Ports tool. Choose the appropriate option from the Action menu.	"How to Initialize a Port" on page 202

Setting Up Terminals and Modems With Serial Ports Tool (Overview)

You can set up serial ports with the Solaris Management Console's Serial Ports tool. Select a serial port from the Serial Ports window and then choose a Configure option from the Action menu to configure the following:

- Terminal
- Modem – Dial–In
- Modem – Dial–Out
- Modem – Dial–In/Dial–Out
- Initialize Only – No Connection

The Configure options provide access to the templates for configuring these services. You can view two levels of detail for each serial port: Basic and Advanced. You can access the Advanced level of detail for each serial port after it is configured by selecting the serial port and selecting the Properties option from the Action menu. After a serial port is configured, you can disable or enable the port with the SAF commands. For information on using the SAF commands, see Chapter 10.

For information on using the Serial Ports command–line interface, see the`smserialport`(1M) man page.

Setting Up Terminals

The following table describes the menu items (and their default values) when you set up a terminal by using the Serial Ports tool.

TABLE 9–1 Terminal Default Values

Detail	Item	Default Value
Basic	Port	—
	Description	Terminal
	Service Status	Enabled
	Baud Rate	`9600`
	Terminal Type	`tvi925`
	Login Prompt	`ttyn login:`
Advanced	Carrier Detection	Software
	Option: Connect on Carrier	No
	Option: Bidirectional	Yes
	Option: Initialize Only	No
	Timeout (seconds)	Never
	Port Monitor	`zsmon`
	Service Program	`/usr/bin/login`

Setting Up Modems

The following table describes the three modem templates that are available when you set up a modem using the Serial Ports tool.

TABLE 9–2 Modem Templates

Modem Configuration	Description
Dial-In Only	Users can dial in to the modem but cannot dial out.
Dial-Out Only	Users can dial out from the modem but cannot dial in.
Dial-In and Out (Bidirectional)	Users can either dial in or dial out from the modem.

The following table describes the default values of each template.

TABLE 9–3 Modem Template Default Values

Detail	Item	Modem - Dial-In Only	Modem - Dial-Out Only	Modem - Dial In and Out
Basic	Port Name	—	—	—

TABLE 9–3 Modem Template Default Values *(Continued)*

Detail	Item	Modem - Dial-In Only	Modem - Dial-Out Only	Modem - Dial In and Out
	Description	Modem – Dial In Only	Modem – Dial Out Only	Modem – Dial In and Out
	Service Status	Enabled	Enabled	Enabled
	Baud Rate	9600	`9600`	`9600`
	Login Prompt	`ttyn login:`	`ttyn login:`	`ttyn login:`
Advanced	Carrier Detection	Software	Software	Software
	Option: Connect on Carrier	No	No	No
	Option: Bidirectional	No	No	Yes
	Option: Initialize Only	No	Yes	No
	Timeout (seconds)	Never	Never	Never
	Port Monitor	`zsmon`	`zsmon`	`zsmon`
	Service Program	`/usr/bin/login`	`/usr/bin/login`	`/usr/bin/login`

The following table describes the default values for the Initialize Only template.

TABLE 9–4 Initialize Only - No Connection Default Values

Detail	Item	Default Value
Basic	Port Name	—
	Description	Initialize Only - No Connection
	Service Status	Enabled
	Baud Rate	9600
	Login Prompt	`ttyn login:`
Advanced	Carrier Detection	Software
	Option: Connect on Carrier	No
	Option: Bidirectional	Yes
	Option: Initialize Only	Yes
	Timeout (seconds)	Never
	Port Monitor	`zsmon`

TABLE 9–4 Initialize Only - No Connection Default Values *(Continued)*

Detail	Item	Default Value
	Service Program	`/usr/bin/login`

How to Set Up a Terminal, a Modem, and Initialize a Port (Tasks)

▼ How to Set Up a Terminal

Steps

1. **Start the Solaris Management Console, if it's not already running.**

   ```
   % /usr/sadm/bin/smc &
   ```

   For information on starting the Solaris Management Console, see "Starting the Solaris Management Console" in *System Administration Guide: Basic Administration*.

2. **Click This Computer icon in the Navigation pane.**

3. **Click Devices and Hardware —>Serial Ports.**

   The Serial Ports menu is displayed.

4. **Select the port that will be used with a terminal.**

5. **Choose Configure—>Terminal from the Action menu.**

   The Configure Serial Port window is displayed in Basic Detail mode.

   For a description of the Terminal menu items, see Table 9–1.

6. **Click OK.**

7. **To configure the advanced items, select the port configured as a terminal. Then, select Properties from the Action menu.**

8. **Change the values of template entries, if desired.**

9. **Click OK to configure the port.**

10. **Verify that the terminal service has been added.**

    ```
    $ pmadm -l -s ttyn
    ```

▼ How to Set Up a Modem

Steps

1. **Start the Solaris Management Console, if it's not already running.**

   ```
   % /usr/sadm/bin/smc &
   ```

   For information on starting the Solaris Management Console, see "Starting the Solaris Management Console" in *System Administration Guide: Basic Administration*.

2. **Click This Computer icon in the Navigation pane.**

3. **Click Devices and Hardware—>Serial Ports.**

   The Serial Ports menu is displayed.

4. **Select the port that will be used with a modem.**

5. **Choose one of the following Configure options from the Action menu.**

   a. **Configure—>Modem (Dial In)**

   b. **Configure—>Modem (Dial Out)**

   c. **Configure—>Modem (Dial In/Out)**

   The Configure Serial Port window is displayed in Basic Detail mode.

   For a description of the Modem menu items, see Table 9–3.

6. **Click OK.**

7. **To configure the advanced items, select the port configured as a modem. Then, select Properties from the Action menu.**

8. **Change the values of template entries, if desired.**

9. **Click OK to configure the port.**

10. **Verify that the modem service has been configured.**

    ```
    $ pmadm -l -s ttyn
    ```

▼ How to Initialize a Port

Steps

1. **Start the Solaris Management Console, if it's not already running.**

   ```
   % /usr/sadm/bin/smc &
   ```

   For information on starting the Solaris Management Console, see "Starting the Solaris Management Console" in *System Administration Guide: Basic Administration*.

2. **Click This Computer icon in the Navigation pane.**

3. **Click Devices and Hardware—>Serial Ports.**

   The Serial Ports menu is displayed.

4. **Select the port to be initialized.**

5. **Choose Configure—>Initialize Only – No Connection**

   The Serial Port window is displayed in Basic Detail mode.

   For a description of the Initialize Only menu items, see Table 9–4.

6. **Click OK.**

7. **To configure the advanced items, select the port configured as initialize only. Then, select Properties from the Action menu.**

8. **Change the values of template entries, if desired.**

9. **Click OK to configure the port.**

10. **Verify that the modem service has been initialized.**

```
$ pmadm -l -s ttyn
```

Troubleshooting Terminal and Modem Problems

If users are unable to log in over serial port lines after you have added a terminal or modem and set up the proper services, consider the following possible causes of failure:

- Check with the user.

  Malfunctions in terminals and modem use are typically reported by a user who has failed to log in or dial in. For this reason, begin troubleshooting by checking for a problem on the desktop.

  Some common reasons for login failure include:

  - Login ID or password is incorrect
  - Terminal is waiting for X-ON flow control key (Control-Q)
  - Serial cable is loose or unplugged
  - Terminal configuration is incorrect
  - Terminal is shut off or otherwise has no power

- Check the terminal.

  Continue to troubleshoot by checking the configuration of the terminal or modem. Determine the proper *ttylabel* for communicating with the terminal or modem. Verify that the terminal or modem settings match the *ttylabel* settings.

- Check the terminal server.

  If the terminal checks out, continue to search for the source of the problem on the terminal or modem server. Use the `pmadm` command to verify that a port monitor has been configured to service the terminal or modem and that it has the correct *ttylabel* associated with it. For example:

  ```
  $ pmadm -l -t ttymon
  ```

  Examine the `/etc/ttydefs` file and double–check the label definition against the terminal configuration. Use the `sacadm`command to check the port monitor's status. Use `pmadm` to check the service associated with the port the terminal uses.

- Check the serial connection.

  If the Service Access Controller is *starting* the TTY port monitor *and* the following is true:

  - The `pmadm` command reports that the service for the terminal's port is *enabled*.
  - The terminal's configuration matches the port monitor's configuration.

  Then, continue to search for the problem by checking the serial connection. A serial connection comprises serial ports, cables, and terminals. Test each of these parts by using one part with two other parts that are known to be reliable.

  Test all of the following:

  - Serial ports
  - Modems
  - Cables
  - Connectors

- Do not use the Serial Ports tool to modify serial port settings if the serial port is being used as a console. The correct procedure for changing console settings is by modifying the following line in the `/etc/inittab` file:

  ```
  co:234:respawn:/usr/lib/saf/ttymon -g -h -p "‘uname -n‘ console
  login: " -T terminal_type -d /dev/console -l console -m
  ldterm,ttcompat
  ```

CHAPTER 10

Managing Serial Ports With the Service Access Facility (Tasks)

This chapter describes how to manage serial port services using the Service Access Facility (SAF).

This is a list of the overview information in this chapter.

- "Using the Service Access Facility" on page 206
- "Overall SAF Administration (sacadm)" on page 207
- "Port Monitor Service Administration (pmadm)" on page 208
- "TTY Monitor and Network Listener Port Monitors" on page 211

For information about the step-by-step procedures associated with managing serial ports with the SAF, see the following:

- "Managing Serial Ports With the Service Access Facility (Task Map)" on page 205
- "Administering ttymon services (Task Map)" on page 217

For reference information about the SAF, see "Service Access Facility Administration (Reference)" on page 221.

Managing Serial Ports With the Service Access Facility (Task Map)

Task	Description	For Instructions
Add a ttymon port monitor.	Use the sacadm command to add a ttymon port monitor.	"How to Add a ttymon Port Monitor" on page 213

Task	Description	For Instructions
View a ttymon port monitor status.	Use the sacadm command to view ttymon port monitor status.	"How to View ttymon Port Monitor Status" on page 214
Stop a ttymon port monitor.	Use the sacadm command to stop a ttymon port monitor.	"How to Stop a ttymon Port Monitor" on page 215
Start a ttymon port monitor.	Use the sacadm command to start a ttymon port monitor.	"How to Start a ttymon Port Monitor" on page 215
Disable a ttymon port monitor.	Use the sacadm command to disable a ttymon port monitor.	"How to Disable a ttymon Port Monitor" on page 215
Enable a ttymonport monitor.	Use the sacadm command to enable a ttymon port monitor.	"How to Enable a ttymon Port Monitor" on page 216
Removing a ttymon port monitor	Use the sacadm command to remove a ttymon port monitor.	"How to Remove a ttymon Port Monitor" on page 216

Using the Service Access Facility

You can set up terminals and modems with the Solaris Management Console's Serial Ports tool or the SAF commands.

The SAF is a tool that is used to administer terminals, modems, and other network devices. The top-level SAF program is the Service Access Controller (SAC). The SAC controls port monitors that you administer through the sacadm command. Each port monitor can manage one or more ports.

You administer the services associated with ports through the pmadm command. While services provided through the SAC can differ from network to network, the SAC and its administrative commands, sacadm and pmadm, are network independent.

The following table describes the SAF control hierarchy. The sacadm command is used to administer the SAC, which controls the ttymon and listen port monitors.

The services of ttymon and listen are in turn controlled by the pmadm command. One instance of ttymon can service multiple ports. One instance of listen can provide multiple services on a network interface.

TABLE 10–1 SAF Control Hierarchy

Function	Program	Description
Overall administration	`sacadm`	Command for adding and removing port monitors
Service Access Controller	`sac`	SAF's master program
Port monitors	`ttymon`	Monitors serial port login requests
	`listen`	Monitors requests for network services
Port monitor service administrator	`pmadm`	Command for controlling port monitors services
Services	logins, remote procedure calls	Services to which the SAF provides access
Console administration	console login	The console is automatically set up with an entry in the `/etc/inittab` file that uses `ttymon`-express mode. Do not use the `pmadm` or `sacadm` command to manage the console directly. For more information, see "`ttymon` and the Console Port" on page 211.

Overall SAF Administration (`sacadm`)

The `sacadm` command is the top level of the SAF. The `sacadm` command primarily is used to add and remove port monitors such as `ttymon` and `listen`. Other `sacadm` functions include listing the current status of port monitors and administering port monitor configuration scripts.

Service Access Controller (SAC Program)

The Service Access Controller program (SAC) oversees all port monitors. A system automatically starts the SAC upon entering multiuser mode.

When the SAC program is invoked, it first looks for, and interprets, each system's configuration script. You can use the configuration script to customize the SAC program environment. This script is empty by default. The modifications made to the SAC environment are inherited by all the "children" of the SAC. This inherited environment might be modified by the children.

After the SAC program has interpreted the per-system configuration script, the SAC program reads its administrative file and starts the specified port monitors. For each port monitor, the SAC program runs a copy of itself, forking a child process. Each child process then interprets its per-port monitor configuration script, if such a script exists.

Any modifications to the environment specified in the per-port monitor configuration script affect the port monitor and will be inherited by all its children. Finally, the child process runs the port monitor program by using the command found in the SAC program administrative file.

SAC Initialization Process

The following steps summarize what happens when SAC is first started:

1. The SAC program is spawned by `init` at run–level two.
2. The SAC program reads `/etc/saf/_sysconfig`, the per-system configuration script.
3. The SAC program reads `/etc/saf/_sactab`, the SAC administrative file.
4. The SAC program forks a child process for each port monitor it starts.
5. Each port monitor reads `/etc/saf/pmtag/_config`, the per-port monitor configuration script.

Port Monitor Service Administration (`pmadm`)

The `pmadm` command enables you to administer port monitors' services. In particular, you use the `pmadm` command to add or remove a service and to enable or disable a service. You can also install or replace per-service configuration scripts, or print information about a service.

Each instance of a service must be uniquely identified by a port monitor and a port. When you use the `pmadm` command to administer a service, you specify a particular port monitor with the *pmtag* argument, and a particular port with the *svctag* argument.

For each port monitor type, the SAF requires a specialized command to format port monitor-specific configuration data. This data is used by the `pmadm` command. For `ttymon` and `listen` type port monitors, these specialized commands are `ttyadm` and `nlsadmin`, respectively.

`ttymon` Port Monitor

Whenever you attempt to log in by using a directly connected modem or alphanumeric terminal, `ttymon` goes to work.

As shown in the figure 12–1, the `init` process is the first process to be started at boot time. Consulting its administrative file (`/etc/inittab`), the `init` process starts other processes as they are needed. Listed among those processes is the SAC.

The SAC, in turn, automatically starts the port monitors designated in its administrative file, `/etc/saf/_sactab`. The following figure shows only a single `ttymon` port monitor.

After the `ttymon` port monitor has been started, it monitors the serial port lines for service requests.

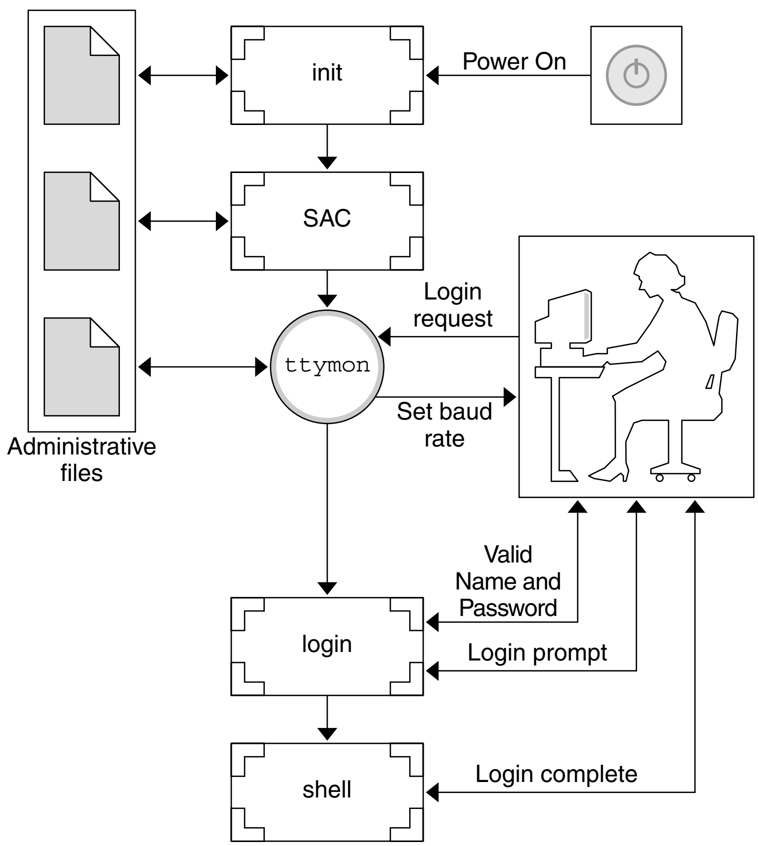

FIGURE 10–1 How `ttymon` Helps Process a Login Request

When someone attempts to log in by using an alphanumeric terminal or a modem, the serial port driver passes the activity to the operating system. The `ttymon` port monitor notes the serial port activity, and attempts to establish a communications link. The `ttymon` port monitor determines which data transfer rate, line discipline, and handshaking protocol are required to communicate with the device.

After the proper parameters for communication with the modem or terminal are established, the `ttymon` port monitor passes these parameters to the login program and transfers control to it.

Port Initialization Process

When an instance of the `ttymon` port monitor is invoked by the SAC, `ttymon` starts to monitor its ports. For each port, the `ttymon` port monitor first initializes the line disciplines, if they are specified, and the speed and terminal settings. The values used for initialization are taken from the appropriate entry in the `/etc/ttydefs` file.

The `ttymon` port monitor then writes the prompt and waits for user input. If the user indicates that the speed is inappropriate by pressing the Break key, the `ttymon` port monitor tries the next speed and writes the prompt again.

If *autobaud* is enabled for a port, the `ttymon` port monitor tries to determine the baud rate on the port automatically. Users must press Return before the `ttymon` port monitor can recognize the baud rate and print the prompt.

When valid input is received, the `ttymon` port monitor does the following tasks:

- Interprets the per-service configuration file for the port
- Creates an `/etc/utmpx` entry, if required
- Establishes the service environment
- Invokes the service associated with the port

After the service terminates, the `ttymon` port monitor cleans up the `/etc/utmpx` entry, if this entry exists, and returns the port to its initial state.

Bidirectional Service

If a port is configured for bidirectional service, the `ttymon` port monitor does the following:

- Allows users to connect to a service
- Allows the `uucico`, `cu`, or `ct` commands to use the port for dialing out, if the port is free
- Waits to read a character before printing a prompt
- Invokes the port's associated service, without sending the prompt message, when a connection is requested, if the connect-on-carrier flag is set

TTY Monitor and Network Listener Port Monitors

Though the SAF provides a generic means for administering any future or third-party port monitors, only two port monitors are implemented in the Solaris Operating System: `ttymon` and `listen`.

TTY Port Monitor (`ttymon`)

The `ttymon` port monitor is STREAMS-based and does the following:

- Monitors ports
- Sets terminal modes, baud rates, and line disciplines
- Invokes the login process

The `ttymon` port monitor provides Solaris users the same services that the `getty` port monitor did under previous versions of SunOS 4.1 software.

The `ttymon` port monitor runs under the SAC program and is configured with the `sacadm` command. Each instance of `ttymon` can monitor multiple ports. These ports are specified in the port monitor's administrative file. The administrative file is configured by using the `pmadm` and `ttyadm` commands.

`ttymon` and the Console Port

Console services are not managed by the Service Access Controller nor by any explicit `ttymon` administration file. An entry in the `/etc/inittab` file is used to manage the console port by using the `ttymon` port monitor in *express* mode. Express mode is a special `ttymon` mode that is invoked directly by a command that requires login service.

The default console entry in the `/etc/inittab` file is as follows:

```
co:234:respawn:/usr/lib/saf/ttymon -g -h -p "`uname -n` console login: "
 -T terminal-type -d /dev/console -l console -m ldterm,ttcompat
```

`co:234:respawn:`
: `co` identifies the entry as the console. `234` identifies the run levels for the action. `respawn` means that the console entry should be restarted if it fails or doesn't exist at run levels 2, 3, and 4.

`/usr/lib/saf/ttymon -g -h`
: The `-g` option is used so that he correct baud rate and terminal setting can be set on a port and connect to a login service without being preconfigured by the SAC. The `-h` option forces a line hang-up by setting the line speed to zero before setting the default or specified speed.

`-p "'uname -n' console login:`
: The`- p` option identifies the prompt string for the console port.

`-t` *terminal-type*
: the `-t` option dentifies the terminal type of the console.

`-d /dev/console -l console -m ldterm,ttcompat`
: The `-d` option identifies the console device. The `-l` option identifies the `ttylabel` in the `/etc/ttydefs` file. The `-m` option identifies the STREAMS modules to be pushed.

`ttymon`-Specific Administrative Command (`ttyadm`)

The `ttymon` administrative file is updated by the `sacadm` and `pmadm` commands, as well as by the `ttyadm` command. The `ttyadm` command formats `ttymon`-specific information and writes it to standard output, providing a means for presenting formatted `ttymon`-specific data to the `sacadm` and `pmadm` commands.

Thus, the `ttyadm` command does not administer `ttymon` directly. The `ttyadm` command complements the generic administrative commands, `sacadm` and `pmadm`. For more information, see the `ttyadm`(1M) man page.

Network Listener Service (`listen`)

The `listen` port monitor runs under the SAC and does the following:

- Monitors the network for service requests
- Accepts requests when they arrive
- Invokes servers in response to those service requests

The `listen` port monitor is configured by using the `sacadm` command. Each instance of `listen` can provide multiple services. These services are specified in the port monitor's administrative file. This administrative file is configured by using the `pmadm` and `nlsadmin` commands.

The network listener process can be used with any connection-oriented transport provider that conforms to the Transport Layer Interface (TLI) specification. In the Solaris Operating System, `listen` port monitors can provide additional network services not provided by the `inetd` service.

Special `listen`-Specific Administrative Command (`nlsadmin`)

The `listen` port monitor's administrative file is updated by the `sacadm` and `pmadm` commands, as well as by the `nlsadmin` command. The `nlsadmin` command formats `listen`-specific information and writes it to standard output, providing a means of presenting formatted `listen`-specific data to the `sacadm` and `pmadm` commands.

Thus, the `nlsadmin` command does not administer `listen` directly. The command complements the generic administrative commands, `sacadm` and `pmadm`.

Each network, configured separately, can have at least one instance of the network listener process associated with it. The `nlsadmin` command controls the operational states of `listen` port monitors.

The `nlsadmin` command can establish a `listen` port monitor for a given network, configure the specific attributes of that port monitor, and *start* and *kill* the monitor. The `nlsadmin` command can also report on the `listen` port monitors on a machine.

For more information, see the `nlsadmin`(1M) man page.

Administering `ttymon` Port Monitors

Use the `sacadm` command to add, list, remove, kill, start, enable, disable, enable, and remove a `ttymon` port monitor.

▼ How to Add a `ttymon` Port Monitor

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Add a `ttymon` port monitor.**

   ```
   # sacadm -a -p mbmon -t ttymon -c /usr/lib/saf/ttymon -v `ttyadm
   -V` -y "TTY Ports a & b"
   ```

   `-a` Specifies the *add* port monitor option.

   `-p` Specifies the *pmtag* `mbmon` as the port monitor tag.

   `-t` Specifies the port monitor *type* as `ttymon`.

-c Defines the *command* string used to start the port monitor.

-v Specifies the *version* number of the port monitor.

-y Defines a comment to describe this instance of the port monitor.

▼ How to View `ttymon` Port Monitor Status

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **View the status of a `ttymon` port monitor.**

   ```
   # sacadm -l -p mbmon
   ```

   -l Specifies the *list* port monitor status flag.

   -p Specifies the *pmtag* `mbmon` as the port monitor tag.

Example 10–1 Viewing `ttymon` Port Monitor Status

This example shows how to view a port monitor named, `mbmon`.

```
# sacadm -l -p mbmon
PMTAG   PMTYPE   FLGS  RCNT  STATUS     COMMAND
mbmon   ttymon   -     0     STARTING   /usr/lib/saf/ttymon #TTY Ports a & b
```

PMTAG	Identifies the port monitor name, `mbmon`.
PMTYPE	Identifies the port monitor type, `ttymon`.
FLGS	Indicates whether the following flags are set: ■ d — Do not enable the new port monitor. ■ x — Do not start the new port monitor. ■ dash (-) — No flags are set.
RCNT	Indicates the return count value. A return count of `0` indicates that the port monitor is not to be restarted if it fails.
STATUS	Indicates the current status of the port monitor.
COMMAND	Identifies the command used to start the port monitor.
#TTY Ports a & b	Identifies any comment used to describe the port monitor.

▼ How to Stop a `ttymon` Port Monitor

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Stop a `ttymon` port monitor.**

   ```
   # sacadm -k -p mbmon
   ```

   `-k` Specifies the *kill* port monitor status flag.

   `-p` Specifies the *pmtag* `mbmon` as the port monitor tag.

▼ How to Start a `ttymon` Port Monitor

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Start a killed `ttymon` port monitor.**

   ```
   # sacadm -s -p mbmon
   ```

   `-s` Specifies the *start* port monitor status flag.

   `-p` Specifies the *pmtag* `mbmon` as the port monitor tag.

▼ How to Disable a `ttymon` Port Monitor

Disabling a port monitor prevents new services from starting, without affecting existing services.

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Disable a `ttymon` port monitor.**

   ```
   # sacadm -d -p mbmon
   ```

`-d` Specifies the *disable* port monitor status flag.

`-p` Specifies the *pmtag* `mbmon` as the port monitor tag.

▼ How to Enable a `ttymon` Port Monitor

Enabling a `ttymon` port monitor allows it to service new requests.

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Enable a `ttymon` port monitor.**

   ```
   # sacadm -e -p mbmon
   ```

   `-e` Specifies the *enable* port monitor status flag.

   `-p` Specifies the *pmtag* `mbmon` as the port monitor tag.

▼ How to Remove a `ttymon` Port Monitor

Removing a port monitor deletes all the configuration files associated with it.

Note – Port monitor configuration files cannot be updated or changed by using the `sacadm` command. To re-configure a port monitor, *remove* it and then *add* a new one.

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Remove a `ttymon` port monitor.**

   ```
   # sacadm -r -p mbmon
   ```

   `-r` Specifies the *remove* port monitor status flag.

   `-p` Specifies the *pmtag* `mbmon` as the port monitor tag.

Administering `ttymon` services (Task Map)

Task	Description	For Instructions
Add a `ttymon` service.	Use the `pmadm` command to add a service.	"How to Add a Service" on page 217
View the Status of a TTY Port Service.	Use the `pmadm`command to view the status of a TTY port.	"How to View the Status of a TTY Port Service" on page 218
Enable a port monitor service.	Use the `pmadm` command with the `-e` option to enable a port monitor.	"How to Enable a Port Monitor Service" on page 220
Disable a port monitor service.	Use the `pmadm` command with the `-d` option to disable a port monitor.	"How to Disable a Port Monitor Service" on page 221

Administering `ttymon` Services

Use the `pmadm` command to add services, list the services of one or more ports associated with a port monitor, and enable or disable a service.

▼ How to Add a Service

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Add a standard terminal service to the `mbmon` port monitor.**

```
# pmadm -a -p mbmon -s a -i root -v `ttyadm -V` -m "`ttyadm -i 'Terminal
 disabled' -l contty -m ldterm,ttcompat -S y -d /dev/term/a
-s /usr/bin/login`"
```

Note – In this example, the input wraps automatically to the next line. Do not use a Return key or line feed.

`-a` Specifies the *add* port monitor status flag.

`-p` Specifies the *pmtag* `mbmon` as the port monitor tag.

`-s` Specifies the *svctag* `a` as the port monitor *service* tag.

`-i` Specifies the *identity* to be assigned to *svctag* when the service runs.

`-v` Specifies the *version* number of the port monitor.

`-m` Specifies the `ttymon`-specific configuration data formatted by `ttyadm`.

The preceding `pmadm` command contains an embedded `ttyadm` command. The options in this embedded command are as follows:

`-b` Specifies the *bidirectional* port flag.

`-i` Specifies the *inactive* (disabled) response message.

`-l` Specifies which TTY *label* in the `/etc/ttydefs` file to use.

`-m` Specifies the STREAMS *modules* to push before invoking this service.

`-d` Specifies the full path name to the *device* to use for the TTY port.

`-s` Specifies the full path name of the *service* to invoke when a connection request is received. If arguments are required, enclose the command and its arguments in quotation marks (").

▼ How to View the Status of a TTY Port Service

Use the `pmadm` command as shown in this procedure to list the status of a TTY port or all the ports that are associated with a port monitor.

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **List one service of a port monitor.**

   ```
   # pmadm -l -p mbmon -s a
   ```

   `-l` Lists service information on the system.

   `-p` Specifies the *pmtag* `mbmon` as the port monitor tag.

-s Specifies the *svctag* a as the port monitor *service* tag.

Example 10–2 Viewing the Status of a TTY Port Monitor Service

This example lists all services of a port monitor.

```
# pmadm -l -p mbmon
PMTAG  PMTYPE  SVCTAG FLAGS ID    <PMSPECIFIC>
mbmon  ttymon  a      -     root  /dev/term/a - - /usr/bin/login - contty
ldterm,ttcompat login:  Terminal disabled tvi925 y  #
```

PMTAG	Identifies the port monitor name, mbmon, that is set by using the pmadm -p command.
PMTYPE	Identifies the port monitor type, ttymon.
SVCTAG	Indicates the service tag value that is set by using the pmadm -s command.
FLAGS	Identifies whether the following flags are set by using the pmadm -f command: ■ x — Do not enable the service. ■ u — Create a utmpx entry for the service. ■ dash (-) — No flags are set.
ID	Indicates the identity assigned to the service when it is started. This value is set by using the pmadm -i command.
<PMSPECIFIC>	*Information*
/dev/term/a	Indicates the TTY port path name that is set by using the ttyadm -d command.
-	Indicates whether the following flags are set by using the ttyadm -c -b -h -I -r command: ■ c — Sets the connect on carrier flag for the port. ■ b — Sets the port as bidirectional, allowing both incoming and outgoing traffic. ■ h — Suppresses an automatic hangup immediately after an incoming call is received. ■ I — Initializes the port. ■ r— Forces ttymon to wait until it receives a character from the port before it prints the login: message. ■ dash (-) — No flags are set
-	Indicates a value that is set by using the ttyadm -r *count* option. This option determines when ttymon displays a prompt after receiving data

	from a port. If *count* is 0, `ttymon` waits until it receives any character. If *count* is greater than 0, `ttymon` waits until *count* new lines have been received. No value is set in this example.
`/usr/bin/login`	Identifies the full path name of the service to be invoked when a connection is received. This value is set by using the `ttyadm -s` command.
`-`	Identifies the `ttyadm -t` command's time-out value. This option specifies that `ttymon` should close a port if the open on the port succeeds, and no input data is received in *timeout* seconds. There is no time-out value in this example.
`contty`	Identifies the TTY label in the `/etc/ttydefs` file. This value is set by using the `ttyadm -l` command.
`ldterm,ttcompat`	Identifies the STREAMS modules to be pushed. These modules are set by using the `ttyadmin -m` command.
`login: Terminal disabled`	Identifies an inactive message to be displayed when the port is disabled. This message is set by using the `ttyadm -i` command.
`tvi925`	Identifies the terminal type, if set, by using the `ttyadm -T` command. The terminal type is `tvi925` in this example.
`y`	Identifies the software carrier value that is set by using the `ttyadm -S` command. `n` turns the software carrier off. `y` turns the software carrier on. The software carrier is turned on in this example.
`#`	Identifies any comment specified with the `pmadm -y` command. There is no comment in this example.

▼ How to Enable a Port Monitor Service

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Enable a disabled port monitor service.**

```
# pmadm -e -p mbmon -s a
```

`-e` Specifies the *enable* flag.

`-p` Specifies the *pmtag* `mbmon` as the port monitor tag.

`-s` Specifies the *svctag* `a` as the port monitor *service* tag.

▼ How to Disable a Port Monitor Service

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Disable a port monitor service.**

```
# pmadm -d -p mbmon -s a
```

`-d` Specifies the *disable* flag.

`-p` Specifies the *pmtag* `mbmon` as the port monitor tag.

`-s` Specifies the *svctag* `a` as the port monitor *service* tag.

Service Access Facility Administration (Reference)

Files Associated With the SAF

The SAF uses configuration files that can be modified by using the `sacadm` and `pmadm` commands. You should not need to manually edit the configuration files.

File Name	Description
`/etc/saf/_sysconfig`	Per-system configuration script.

File Name	Description
/etc/saf/_sactab	The SAC's administrative file that contains configuration data for the port monitors that the SAC controls
/etc/saf/*pmtag*	Home directory for port monitor *pmtag*
/etc/saf/*pmtag*/_config	Per-port monitor configuration script for port monitor pmtag if it exists
/etc/saf/*pmtag*/*_pmtab*	Port monitor *pmtag*'s administrative file that contains port monitor-specific configuration data for the services *pmtag* provides
/etc/saf/*pmtag*/*svctag*	Per-service configuration script for service *svctag*
/var/saf/log	The SAC's log file
/var/saf/*pmtag*	Directory for files created by *pmtag*, for example, log files

/etc/saf/_sactab File

The information in the /etc/saf/_sactab file is as follows:

```
# VERSION=1
zsmon:ttymon::0:/usr/lib/saf/ttymon     #
```

# VERSION=1	Indicates the Service Access Facility version number.
zsmon	Is the name of the port monitor.
ttymon	Is the type of port monitor.
::	Indicates whether the following two flags are set: ■ d — Do not enable the port monitor. ■ x — Do not start the port monitor. No flags are set in this example.
0	Indicates the return code value. A return count of 0 indicates that the port monitor is not be restarted if the port monitor fails.
/usr/lib/saf/ttymon	Indicates the port monitor path name.

/etc/saf/*pmtab/_pmtab* File

The /etc/saf/*pmtab/_pmtab* file, such as /etc/saf/zsmon/_pmtab, is similar to the following:

```
# VERSION=1
ttya:u:root:reserved:reserved:reserved:/dev/term/a:I::/usr/bin/login::9600:
ldterm,ttcompat:ttya login\: ::tvi925:y:#
```

`# VERSION=1`	Indicates the Service Access Facility version number.
`ttya`	Indicates the service tag.
`x,u`	Identifies whether the following flags are set: ■ `x` — Do not enable the service. ■ `u` — Create a `utmpx` entry for the service.
`root`	Indicates the identity assigned to the service tag.
`reserved`	This field is reserved for future use.
`reserved`	This field is reserved for future use.
`reserved`	This field is reserved for future use.
`/dev/term/a`	Indicates the TTY port path name.
`/usr/bin/login`	Identifies the full path name of the service to be invoked when a connection is received.
`:c,b,h,I,r:`	Indicates whether the following flags are set: `c` — Sets the connect on carrier flag for the port. `b` — Sets the port as bidirectional, allowing both incoming and outgoing traffic. `h` — Suppresses an automatic hangup immediately after an incoming call is received. `I` — Initializes the port. `r` — Forces `ttymon` to wait until it receives a character from the port before `ttymon` prints the `login:` message.
`9600`	Identifies the TTY label defined in the `/etc/ttydefs` file.
`ldterm,ttcompat`	Identifies the STREAMS modules to be pushed.
`ttya login\:`	Identifies the prompt to be displayed.
`:y/n:`	Indicates yes or no response.
message	Identifies any inactive (disabled) response message.
`tvi925`	Identifies the terminal type.
`y`	Indicates whether the software carrier is set (`y/n`).

Service States

The `sacadm` command controls the states of services. The following table describes the possible states of services.

State	Description
Enabled	*Default state* – When the port monitor is added, the service operates.
Disabled	*Default state* – When the port monitor is removed, the service stops.

To determine the state of any particular service, use the following:

```
# pmadm -l -p portmon-name -ssvctag
```

Port Monitor States

The `sacadm` command controls the states of the `ttymon` and `listen` port monitors. The following table describes the possible port monitor states.

State	Description
`Started`	*Default state* – When the port monitor is added, it is automatically started.
`Enabled`	*Default state* – When the port monitor is added, it is automatically ready to accept requests for service.
`Stopped`	*Default state* – When the port monitor is removed, it is automatically stopped.
`Disabled`	*Default state* – When the port monitor is removed, it automatically continues existing services and refuses to add new services.
`Starting`	*Intermediate state* – The port monitor is in the process of starting.
`Stopping`	*Intermediate state* – The port monitor has been manually terminated, but it has not completed its shutdown procedure. The port monitor is on the way to becoming stopped.
`Notrunning`	*Inactive state* – The port monitor has been killed. All ports previously monitored are inaccessible. An external user cannot tell whether a port is `disabled` or `notrunning`.
`Failed`	*Inactive state* – The port monitor is unable to start and remain running.

To determine the state of any particular port monitor, use the following command:

```
# sacadm -l -p portmon-name
```

Port States

Ports can be enabled or disabled depending on the state of the port monitor that controls the ports.

States	Description
Serial (`ttymon`) port states	
`Enabled`	The `ttymon` port monitor sends a prompt message to the port and provides login service to it.
`Disabled`	Default state of all ports if `ttymon` is killed or disabled. If you specify this state, `ttymon` sends out the `disabled` message when it receives a connection request.

CHAPTER 11

Managing System Resources (Overview)

This chapter provides a brief description of features that are new in the Solaris Operating System and a road map to help you manage system resources.

Using these features, you can display general system information, monitor disk space, set disk quotas and use accounting programs. You can also schedule the `cron` and `at` commands to automatically run routine commands.

This section does not cover Solaris resource management that enables you to allocate, monitor, and control system resources in a flexible way.

For information on the procedures that are associated with managing system resources without Solaris resource management, see "Managing System Resources (Road Map)" on page 228.

For information on managing system resources with Solaris resource management, see Chapter 1, "Introduction to Solaris 10 Resource Manager," in *System Administration Guide: N1 Grid Containers, Resource Management, and Solaris Zones*.

What's New in Managing System Resources?

This section describes new features for managing system resources.

New psrinfo Command Option to Identify Chip Multithreading Features

The psrinfo command has been modified to provide information about physical processors, in addition to information about virtual processors. This enhanced functionality has been added to identify chip multithreading (CMT) features. The new -p option reports the total number of physical processors that are in a system. Using the psrinfo -pv command will list all the physical processors that are in the system, as well as the virtual processors that are associated with each physical processor. The default output of the psrinfo command continues to display the virtual processor information for a system.

For more information, see the psrinfo(1M) man page.

For information about the procedures associated with this feature, see "How to Display a System's Physical Processor Type" on page 235.

New localeadm Command

The new localeadm command allows you to change the locales on your system without reinstalling the OS or manually adding and removing packages. This command also allows you to query your system to determine which locales are installed. To run the localeadm command, you must have superuser privileges or assume an equivalent role through role-based access control (RBAC).

For more information, see the localeadm(1M) man page.

For more information in this guide, see Chapter 12.

Managing System Resources (Road Map)

Task	Description	Instructions
Displaying and changing system information	Use various commands to display and change system information, such as general system information, the language environment, the date and time, and the system's host name.	Chapter 12

Task	Description	Instructions
Managing disk use	Identify how disk space is used and take steps to remove old and unused files.	Chapter 13
Managing quotas	Use UFS file system quotas to manage how much disk space is used by users.	Chapter 14
Scheduling system events	Use cron and at jobs to help schedule system routines that can include clean up of old and unused files.	Chapter 15
Managing system accounting	Use system accounting to identify how users and applications are using system resources.	Chapter 16
Managing system resources with Solaris Resource Management	Use resource manager to control how applications use available system resources and to track and charge resource usage.	Chapter 1, "Introduction to Solaris 10 Resource Manager," in *System Administration Guide: N1 Grid Containers, Resource Management, and Solaris Zones*

CHAPTER 12

Displaying and Changing System Information (Tasks)

This chapter describes the tasks that are required to display and change the most common system information.

For information about the procedures associated with displaying and changing system information, see the following:

- "Displaying System Information (Task Map)" on page 231
- "Changing System Information (Task Map)" on page 240

For overview information about managing system resources, see Chapter 11.

Displaying System Information (Task Map)

Task	Description	For Instructions
Determine whether a system can run the 64–bit Solaris OS	Use the `uname -m` command to verify if a system is an UltraSPARC system.	"How to Determine if a System Can Run the 64–Bit Solaris Operating System" on page 233
Determine whether a system has 64–bit capabilities enabled.	Use the `isainfo` command to determine if a system has 64'-bit capabilities enabled.	"How to Determine if a System Has 64–Bit Solaris Capabilities Enabled" on page 234

Task	Description	For Instructions
Display a system's physical processor type.	Use the `psrinfo -p` command to list the total number of physical processors on a system. Use the `psrinfo -pv` command to list all physical processors on a system and the virtual processors that is associated with each physical processor.	"How to Display a System's Physical Processor Type" on page 235
Display a system's logical processor type.	Use the `psrinfo -v` command to display a system's virtual processor type.	"How to Display a System's Virtual Processor Type" on page 236
Display Solaris Release Information	Display the contents of the `/etc/release` file to identify your Solaris release version.	"How to Display Solaris Release Information" on page 236
Display General System Information.	Use the `showrev` command to display general system information.	"How to Display General System Information" on page 236
Display a system's Host ID number.	Use the `hostid` command to display your system's host id.	"How to Display a System's Host ID Number" on page 237
Display a System's Installed Memory	Use the `prtconf` command to display information about your system's installed memory.	"How to Display a System's Installed Memory" on page 238
Display a system's date and time.	Use the `date` command to display your system's date and time.	"How to Display the Date and Time" on page 238
Display locales that are installed on a system.	Use the `localeadm` command to display locales that are installed on your system.	"How to Display Locales Installed on a System" on page 238
Determine if a locale is installed on a system.	Use the `-q` option of the `localeadm` command and a locale to determine if a locale is installed on your system.	"How to Determine if a Locale is Installed on a System" on page 239

Displaying System Information

The following table describes commands that enable you to display general system information.

TABLE 12–1 Commands for Displaying System Information

Command	System Information Displayed	Man Page
`date`	Date and time	`date`(1)
`hostid`	Host ID number	`hostid`(1)
`isainfo`	The number of bits supported by *native* applications on the running system, which can be passed as a token to scripts	`isainfo`(1)
`isalist`	Processor type for x86 based systems	`psrinfo`(1M)
`localeadm`	Locales installed on the system	`localeadm`(1M)
`prtconf`	Installed memory	`prtconf`(1M)
`psrinfo`	Processor type	`psrinfo`(1M)
`showrev`	Host name, host ID, release, kernel architecture, application architecture, hardware provider, domain, and kernel version	`showrev`(1M)
`uname`	Operating system name, release, version, node name, hardware name, and processor type	`uname`(1)

▼ How to Determine if a System Can Run the 64–Bit Solaris Operating System

Steps

1. **Currently, the only platform capable of supporting the 64–bit Solaris Operating System is an UltraSPARC® system. To verify if a system is an UltraSPARC system, use the following command:**

   ```
   $ uname -m
   sun4u
   ```

   If the output of the `uname -m` command is `sun4u`, then the machine is an UltraSPARC system.

2. **Alternately, you can verify that the machine is an UltraSPARC system by using the `psrinfo` command:**

```
# psrinfo -v
Status of virtual processor 0 as of: 02/25/2004 16:06:58
  on-line since 02/03/2004 15:07:52.
  The sparcv9 processor operates at 650 MHz,
        and has a sparcv9 floating point processor.
```

   If the processor type is `sparcv9`, then the platform is only capable of running the 64-bit Solaris Operating System. The current release of the Solaris OS does not include a 32–bit Solaris kernel.

▼ How to Determine if a System Has 64–Bit Solaris Capabilities Enabled

Step ● **Use the `isainfo` command to determine if a system has 64-bit capabilities enabled, which means that the system is booted with the 64–bit kernel.**

```
# isinfo options
```

The `isainfo` command, run without specifying any options, displays the name or names of the native instruction sets for applications supported by the current OS version.

`-v` Prints detailed information about the other options

`-b` Prints the number of bits in the address space of the native instruction set.

`-n` Prints the name of the native instruction set used by portable applications supported by the current version of the OS.

`-k` Prints the name of the instruction set or sets that are used by the OS kernel components such as device drivers and STREAMS modules.

Example 12–1 Determining If a System Has 64–bit Solaris Capabilities Enabled

The `isainfo` command output for an UltraSPARC system that is running previous releases of the Solaris OS using a 32-bit kernel appears as follows:

```
$ isainfo -v
32-bit sparc applications
```

This output means that this system can support only 32–bit applications.

The current release of the Solaris OS only ships a 64–bit kernel on SPARC based systems. The `isainfo` command output for an UltraSPARC system that is running a 64–bit kernel appears as follows:

```
$ isainfo -v
64-bit sparcv9 applications
```

```
32-bit sparc applications
```

This output means that this system is capable of supporting both 32–bit and 64–bit applications.

Use the `isainfo -b` command to display the number of bits supported by native applications on the running system.

The output from a SPARC based, x86 based, or UltraSPARC system that is running the 32–bit Solaris Operating System appears as follows:

```
$ isainfo -b
32
```

The `isainfo`command output from a 64–bit UltraSPARC system that is running the 64–bit Solaris Operating System appears as follows:

```
$ isainfo -b
64
```

The command returns `64` only. Even though a 64–bit UltraSPARC system can run both types of applications, 64–bit applications are the best kind of applications to run on a 64–bit system.

▼ How to Display a System's Physical Processor Type

Step ● **Use the `psrinfo -p` command to display the total number of physical processors on a system.**

```
$ psrinfo -p
1
```

Use the `psrinfo -pv` command to display information about each physical processor on a system, and the virtual processor associated with each physical processor.

```
$ psrinfo -pv
The UltraSPARC-IV physical processor has 2 virtual processors (8, 520)
The UltraSPARC-IV physical processor has 2 virtual processors (9, 521)
The UltraSPARC-IV physical processor has 2 virtual processors (10, 522)
The UltraSPARC-IV physical processor has 2 virtual processors (11, 523)
The UltraSPARC-III+ physical processor has 1 virtual processor (16)
The UltraSPARC-III+ physical processor has 1 virtual processor (17)
The UltraSPARC-III+ physical processor has 1 virtual processor (18)
The UltraSPARC-III+ physical processor has 1 virtual processor (19)
```

When you use the `psrinfo -pv` command on an x86 based system, the following output is displayed:

```
$ psrinfo -pv
The i386 physical processor has 2 virtual processors (0, 2)
The i386 physical processor has 2 virtual processors (1, 3)
```

▼ How to Display a System's Virtual Processor Type

Step ● **Use the `psrinfo -v` command to display information about a system's processor type. This command also displays the number of processors on a system, as shown in the following example.**

```
$ psrinfo -v
Status of virtual processor 0 as of: 04/16/2004 10:32:13
  on-line since 03/22/2004 19:18:27.
  The sparcv9 processor operates at 650 MHz,
  and has a sparcv9 floating point processor.
```

On an x86 based system, use the `isalist` command to display the virtual processor type.

```
$ isalist
pentium_pro+mmx pentium_pro pentium+mmx pentium i486 i386 i86
```

▼ How to Display Solaris Release Information

Step ● **Display the contents of the `/etc/release` file to identify your Solaris release version.**

```
$ more /etc/release
                        Solaris 10 s10_51 SPARC
           Copyright 2004 Sun Microsystems, Inc.  All Rights Reserved.
                        Use is subject to license terms.
                            Assembled 21 January 2004
```

▼ How to Display General System Information

Step ● **To display general system information, use the `showrev` command.**

$ **showrev** *options*

`-a`	Prints all system revision information available.
`-c (command)`	Prints the revision information about command
`-p`	Prints only the revision information about patches.

`-R (root_path)`	Defines the full path name of a directory to use as the root_path.
`-s (host name)`	Performs this operation on the specified host name
`-w`	Prints only the OpenWindows revision information.

You can also use the `uname` command to display system information. The following example shows the `uname` command output. The `-a` option displays the operating system name as well as the system node name, operating system release, operating system version, hardware name, and processor type.

```
$ uname
SunOS
$ uname -a
SunOS starbug 5.10 Generic sun4u sparc SUNW,Ultra-5_10
$
```

Example 12–2 Displaying General System Information

The following example shows the `showrev` command output. The `-a` option displays all available system information.

```
$ showrev -a
Hostname: touchstone
Hostid: 8099dfb9
Release: 5.10
Kernel architecture: sun4u
Application architecture: sparc
Hardware provider: Sun_Microsystems
Domain: boulder.Central.Sun.COM
Kernel version: SunOS 5.10 s10_46

OpenWindows version:
Solaris X11 Version 6.6.2 20 October 2003

No patches are installed
```

▼ How to Display a System's Host ID Number

Step ● **To display the host ID number in hexadecimal format, use the `hostid` command.**

Example 12–3 Displaying a System's Host ID Number

The following example shows sample output from the `hostid` command.

```
$ hostid
80a5d34c
```

▼ How to Display a System's Installed Memory

Step ● **To display the amount of memory that is installed on your system, use the `prtconf` command.**

Example 12–4 Displaying a System's Installed Memory

The following example shows sample output from the `prtconf` command. The `grep Memory` command selects output from the `prtconf` command to display memory information only.

```
# prtconf | grep Memory
Memory size: 128 Megabytes
```

▼ How to Display the Date and Time

Step ● **To display the current date and time according to your system clock, use the `date` command.**

Example 12–5 Displaying the Date and Time

The following example shows sample output from the `date` command.

```
$ date
Wed Jan 21 17:32:59 MST 2004
$
```

▼ How to Display Locales Installed on a System

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Display the locales currently installed on your system using the `localeadm` command. The `-l` option displays the locales that are installed on the system. For example:**

```
# localeadm -l
Checking for installed pkgs. This could take a while.

Checking for Australasia region (aua)
```

```
(1of2 pkgs)
|......|
.
.
.
The following regions are installed on concordance on Wed Dec 17 15:13:00 MST 2003


POSIX (C)

Central Europe (ceu)
[ Austria, Czech Republic, Germany, Hungary, Poland, Slovakia,
Switzerland (German), Switzerland (French) ]

Done.
```

▼ How to Determine if a Locale is Installed on a System

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Determine if a locale is installed on your system using the `localeadm` command. The `-q` option and a locale queries the system to see if that locale is installed on the system. To see if the Central European region (ceu) is installed on your system, for example:**

   ```
   # localeadm -q ceu
   locale/region name is ceu
   Checking for Central Europe region (ceu)
   .
   .
   .
   The Central Europe region (ceu) is installed on this system
   ```

Changing System Information (Task Map)

Task	Directions	For Instructions
Manually set a system's date and time.	Manually set your system's date and time by using the `date` *mmddHHMM[[cc]yy]* command-line syntax.	"How to Set a System's Date and Time Manually" on page 240
Set up a message-of-the-day.	Set up a message-of-the-day on your system by editing the `/etc/motd` file.	"How to Set Up a Message-of-the-Day" on page 241
Change a system's host name.	Change your system's host name by editing the following files: ■ `/etc/nodename` ■ `/etc/hostname.`*host-name* ■ `/etc/inet/hosts`	"How to Change a System's Host Name" on page 242
Add a locale to a system.	Use the `localeadm` command to add a locale to your system.	How to Add a Locale to a System
Remove a locale from a system.	Use the `-r` option of the `localeadm` command and the locale to remove of locale from your system.	How to Remove a Locale From a System

Changing System Information

This section describes commands that enable you to change general system information.

▼ How to Set a System's Date and Time Manually

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide:*

Security Services.

2. **Enter the new date and time.**

   ```
   # date mmddHHMM[[cc]yy]
   ```

mm	Month, using two digits.
dd	Day of the month, using two digits.
HH	Hour, using two digits and a 24-hour clock.
MM	Minutes, using two digits.
cc	Century, using two digits.
yy	Year, using two digits.

   See the `date`(1) man page for more information.

3. **Verify that you have reset your system's date correctly by using the `date` command with no options.**

Example 12–6 Setting a System's Date and Time Manually

The following example shows how to use the `date` command to manually set a system's date and time.

```
# date
Wed Mar  3 14:04:19 MST 2004
# date 0121173404
Thu Jan 21 17:34:34 MST 2004
```

▼ How to Set Up a Message-of-the-Day

Edit the message-of-the-day file, `/etc/motd`, to include announcements or inquiries to all users of a system when they log in. Use this feature sparingly, and edit this file regularly to remove obsolete messages.

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Edit the `/etc/motd` file and add a message of your choice.**

   Edit the text to include the message that will be displayed during user login. Include spaces, tabs, and carriage returns.

3. **Verify the changes by displaying the contents of the `/etc/motd` file.**

```
$ cat /etc/motd
Welcome to the UNIX Universe. Have a nice day.
```

Example 12–7 Setting Up a Message-of-the-Day

The default message-of-the-day, which is provided when you install Solaris software, contains SunOS version information.

```
$ cat /etc/motd
Sun Microsystems Inc.   SunOS 5.10       Generic  May 2004
```

The following example shows an edited `/etc/motd` file that provides information about system availability to each user who logs in.

```
$ cat /etc/motd
The system will be down from 7:00 a.m to 2:00 p.m. on
Saturday, July 7, for upgrades and maintenance.
Do not try to access the system during those hours.
Thank you.
```

▼ How to Change a System's Host Name

A system's host name is specified in several different locations.

Remember to update your name service database to reflect the new host name.

Use the following procedure to change or rename a system's host name.

You can also use the `sys-unconfig` command to reconfigure a system, including the host name. For more information, see the `sys-unconfig`(1M) man page.

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Change the system's host name in the following files:**

   - `/etc/nodename`
   - `/etc/hostname.`*host-name*
   - `/etc/inet/hosts`

3. **(Optional) If you are using a name service, change the system's host name in the `hosts` file.**

4. **Reboot the system to activate the new host name.**

   ```
   # init 6
   ```

▼ How to Add a Locale to a System

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Add the packages for the locale you want to install on your system using the `localeadm` command. The `-a` option and a locale identifies the locale that you want to add. The `-d` option and a device identifies the device containing the locale packages you want to add. To add the Central European region (ceu) to your system, for example:**

   ```
   # localeadm -a ceu -d /net/install/latest/Solaris/Product

   locale/region name is ceu

   Devices are /net/install/latest/Solaris/Product
   .
   .
   .
   One or more locales have been added.
   To update the list of locales available at
   .
   .
   .
   ```

▼ How to Remove a Locale From a System

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Remove the packages for the locale installed on your system using the `localeadm` command. The `-r` option and a locale identifies the locale that you want to remove from the system. To remove the Central European region (ceu) from your system, for example:**

   ```
   # localeadm -r ceu
   locale/region name is ceu
   Removing packages for Central Europe (ceu)
   .
   .
   .
   One or more locales have been removed.
   ```

```
To update the list of locales available
at the login screen's "Options->Language" menu,
.
.
.
```

CHAPTER 13

Managing Disk Use (Tasks)

This chapter describes how to optimize disk space by locating unused files and large directories.

For information on the procedures associated with managing disk use, see "Managing Disk Use (Task Map)" on page 245.

Managing Disk Use (Task Map)

Task	Description	For Instructions
Display information about files and disk space.	Display information about how disk space is used by using the `df` command.	"How to Display Information About Files and Disk Space" on page 247
Display the size of files.	Display information about the size of files by using the `ls` command with the `-lh` options.	"How to Display the Size of Files" on page 249
Find large files.	The `ls -s` command allows you to sort files by size, in descending order.	"How to Find Large Files" on page 250
Find files that exceed a specified size limit.	Locate and display the names of files that exceed a specified size by using the `find` command with the `-size` option and the value of the specified size limit.	"How to Find Files That Exceed a Specified Size Limit" on page 251

Task	Description	For Instructions
Display the size of directories, subdirectories, and files.	Display the size of one or more directories, subdirectories, and files by using the du command.	"How to Display the Size of Directories, Subdirectories, and Files" on page 252
Display ownership of local UFS file systems.	Display ownership of files by using the quot -a command.	"How to Display the User Ownership of Local UFS File Systems" on page 253
List the newest files.	Display the most recently created or changed files first, by using the ls -t command	"How to List the Newest Files" on page 255
Find and remove old or inactive files.	Use the find command with the -atime and -mtime options to locate files that have not been accessed for a specified number of days. You can remove these files by using therm 'cat *filename*' command.	"How to Find and Remove Old or Inactive Files" on page 255
Clear out temporary directories.	Locate temp directories, then use the rm -r * command to remove the entire directory.	"How to Clear Out Temporary Directories" on page 256
Find and delete core files.	Find and delete core files by using the find > -name core -exec rm {} \; command.	"How to Find and Delete core Files" on page 257
Delete crash dump files.	Delete crash dump files that are located in the /var/crash/ directory by using the rm * command.	"How to Delete Crash Dump Files" on page 258

Displaying Information About Files and Disk Space

This table summarizes the commands available for displaying information about file size and disk space.

Command	Description	Man Page
`df`	Reports the number of free disk blocks and files	`df`(1M)
`du`	Summarizes disk space allocated to each subdirectory	`du`(1)
`find -size`	Searches recursively through a directory based on the size specified with the `-size` option	`find`(1)
`ls -lh`	Lists the size of a file in the power of 1024 scaling	`ls`(1)

▼ How to Display Information About Files and Disk Space

Step ● **Display information about how disk space is used by using the `df` command.**

```
$ df [directory]  [-h] [-t]
```

`df` With no options, lists all mounted file systems and their device names, the number of 512-byte blocks used, and the number of files.

directory Specifies the directory whose file system you want to check.

`-h` Displays disk space in the power of 1024 scaling.

`-t` Displays the total blocks as well as the blocks used for all mounted file systems.

Example 13–1 Displaying Information About File Size and Disk Space

In the following example, all the file systems listed are locally mounted except for `/usr/dist`, which is mounted remotely from the system `venus`.

```
$ df
/                  (/dev/dsk/c0t0d0s0 ):  287530 blocks    92028 files
/usr               (/dev/dsk/c0t0d0s6 ): 1020214 blocks   268550 files
/devices           (/devices          ):       0 blocks        0 files
/proc              (/proc             ):       0 blocks      878 files
/dev/fd            (fd                ):       0 blocks        0 files
/etc/mnttab        (mnttab            ):       0 blocks        0 files
/var/run           (swap              ):  396016 blocks     9375 files
/tmp               (swap              ):  396016 blocks     9375 files
/opt               (/dev/dsk/c0t0d0s5 ):  381552 blocks    96649 files
/export/home       (/dev/dsk/c0t0d0s7 ):  434364 blocks   108220 files
/usr/dist          (venus:/usr/dist   ):14750510 blocks  2130134 files
```

Example 13–2 Displaying File Size Information in 1024 Bytes

In the following example, file system information is displayed in 1024 bytes.

```
$ df -h
Filesystem             size   used  avail capacity  Mounted on
/dev/dsk/c0t0d0s0      1.9G    58M   1.8G     4%    /
/dev/dsk/c0t0d0s6      2.5G   765M   1.7G     1%    /usr
/devices                 0K     0K     0K     0%    /devices
/proc                    0K     0K     0K     0%    /proc
fd                       0K     0K     0K     0%    /dev/fd
mnttab                   0K     0K     0K     0%    /etc/mnttab
swap                   2.1G    24K   2.1G     1%    /var/run
swap                   2.1G     0K   2.1G     0%    /tmp
/dev/dsk/c0t0d0s5      1.9G    12M   1.8G     1%    /opt
/dev/dsk/c0t0d0s7      1.9G    10K   1.8G     1%    /export/home
```

Although /proc and /tmp are local file systems, they are not UFS file systems. /proc is a PROCFS file system, /var/run and /tmp are TMPFS file systems, and /etc/mnttab is an MNTFS file system.

Example 13–3 Displaying Total Number of Blocks and Files Allocated for a File System

The following example shows a list of all mounted file systems, device names, total 512-byte blocks used, and the number of files. The second line of each two-line entry displays the total number of blocks and files that are allocated for the file system.

```
$ df -t
/                  (/dev/dsk/c0t0d0s0 ):   287530 blocks    92028 files
                                  total:   385614 blocks    96832 files
/usr               (/dev/dsk/c0t0d0s6 ):  1020214 blocks   268550 files
                                  total:  2381102 blocks   300288 files
/devices           (/devices          ):        0 blocks        0 files
/proc              (/proc             ):        0 blocks      879 files
                                  total:        0 blocks      924 files
/dev/fd            (fd                ):        0 blocks        0 files
                                  total:        0 blocks       72 files
/etc/mnttab        (mnttab            ):        0 blocks        0 files
```

```
                                    total:        0 blocks        1 files
/var/run           (swap               ):   396112 blocks     9375 files
                                    total:   396112 blocks     9395 files
/tmp               (swap               ):   396112 blocks     9375 files
                                    total:   396128 blocks     9395 files
/opt               (/dev/dsk/c0t0d0s5 ):   381552 blocks    96649 files
                                    total:   385614 blocks    96832 files
/export/home       (/dev/dsk/c0t0d0s7 ):   434364 blocks   108220 files
                                    total:   434382 blocks   108224 files
/usr/dist          (venus:/usr/dist   ): 14750510 blocks  2130134 files
                                    total: 41225162 blocks  2482176 files
```

Checking the Size of Files

im

You can check the size of files and sort them by using the `ls` command. You can find files that exceed a size limit by using the `find` command. For more information, see the`ls`(1) and `find`(1) man pages.

Note – If you run out of space in the `/var` directory, do not symbolically link the `/var` directory to a directory on a file system with more disk space. Doing so, even as a temporary measure, might cause problems for certain Solaris daemon processes and utilities.

▼ How to Display the Size of Files

Steps

1. **Change to the directory where the files you want to check are located.**

2. **Display the size of the files.**

   `$ ls [-lh] [-s]`

   `-l` Displays a list of files and directories in long format, showing the sizes in bytes. (See the example that follows.)

   `-h` Scales file sizes and directory sizes into Kbytes, Mbytes, Gbytes, or Tbytes when the file or directory size is larger than 1024 bytes. This option also modifies the output displayed by the `-o`, `-n`, `-@`, and `-g` options to display file or directory sizes in the new format. For more information, see the`ls`(1) man page.

`-s` Displays a list of the files and directories, showing the sizes in blocks.

Example 13–4 Displaying the Size of Files

The following example shows that the `lastlog` and `messages` files are larger than the other files in the `/var/adm` directory.

```
$ cd /var/adm
$ ls -lh
total 148
drwxrwxr-x   5 adm      adm          512 Nov 26 09:39 acct/
-rw-------   1 uucp     bin            0 Nov 26 09:25 aculog
drwxr-xr-x   2 adm      adm          512 Nov 26 09:25 exacct/
-r--r--r--   1 root     other       342K Nov 26 13:56 lastlog
drwxr-xr-x   2 adm      adm          512 Nov 26 09:25 log/
-rw-r--r--   1 root     root         20K Nov 26 13:55 messages
drwxr-xr-x   2 adm      adm          512 Nov 26 09:25 passwd/
drwxrwxr-x   2 adm      sys          512 Nov 26 09:39 sa/
drwxr-xr-x   2 root     sys          512 Nov 26 09:49 sm.bin/
-rw-rw-rw-   1 root     bin            0 Nov 26 09:25 spellhist
drwxr-xr-x   2 root     sys          512 Nov 26 09:25 streams/
-rw-r--r--   1 root     bin         3.3K Nov 26 13:56 utmpx
-rw-r--r--   1 root     root           0 Nov 26 10:17 vold.log
-rw-r--r--   1 adm      adm          19K Nov 26 13:56 wtmpx
```

The following example shows that the `lpsched.1` file uses two blocks.

```
$ cd /var/lp/logs
$ ls -s
total 2            0 lpsched       2 lpsched.1
```

▼ How to Find Large Files

Steps

1. **Change to the directory that you want to search.**
2. **Display the size of files in blocks from largest to smallest.**

   ```
   $ ls -s | sort -nr | more
   ```

   The `sort -nr` command sorts the list of files by block size from largest to smallest.

Example 13–5 Finding Large Files

In the following example, the `lastlog` and `messages` files are the largest files in the `/var/adm` directory.

```
$ cd /var/adm
$ ls -s | sort -nr | more
  48 lastlog
  30 messages
  24 wtmpx
  18 pacct
   8 utmpx
   2 vold.log
   2 sulog
   2 sm.bin/
   2 sa/
   2 passwd/
   2 pacct1
   2 log/
   2 acct/
   0 spellhist
   0 aculog
total 144
```

▼ How to Find Files That Exceed a Specified Size Limit

Step ● **To locate and display the names of files that exceed a specified size, use the `find` command.**

$ **find** *directory* **-size +***nnn*

directory	Identifies the directory that you want to search.
`-size +`*nnn*	Is a number of 512-byte blocks. Files that exceed this size are listed.

Example 13–6 Finding Files That Exceed a Specified Size Limit

The following example shows how to find files larger than 400 blocks in the current working directory. The `-print` option displays the output of the `find` command.

```
$ find . -size +400 -print
./Howto/howto.doc
./Howto/howto.doc.backup
./Howto/howtotest.doc
./Routine/routineBackupconcepts.doc
./Routine/routineIntro.doc
./Routine/routineTroublefsck.doc
./.record
./Mail/pagination
./Config/configPrintadmin.doc
./Config/configPrintsetup.doc
./Config/configMailappx.doc
```

```
./Config/configMailconcepts.doc
./snapshot.rs
```

Checking the Size of Directories

You can display the size of directories by using the `du` command and options. Additionally, you can find the amount of disk space used by user accounts on local UFS file systems by using the `quot` command. For more information about these commands, see the `du`(1) and `quot`(1M) man pages.

▼ How to Display the Size of Directories, Subdirectories, and Files

Step ● **Display the size of one or more directories, subdirectories, and files by using the `du` command. Sizes are displayed in 512-byte blocks.**

```
$ du [-as] [directory ...]
```

`du`	Displays the size of each directory that you specify, including each subdirectory beneath it.
`-a`	Displays the size of each file and subdirectory, and the total number of blocks that are contained in the specified directory.
`-s`	Displays the total number of blocks that are contained in the specified directory.
`-h`	Displays the size of each directory in 1024–byte blocks.
`-H`	Displays the size of each directory in 1000–byte blocks.
[*directory ...*]	Identifies one or more directories that you want to check. Separate multiple directories in the command-line syntax with spaces.

Example 13–7 Displaying the Size of Directories, Subdirectories, and Files

The following example shows the sizes of two directories.

```
$ du -s /var/adm /var/spool/lp
130     /var/adm
40      /var/spool/lp
```

The following example shows the sizes of two directories and includes the sizes of all the subdirectories and files that are contained within each directory. The total number of blocks that are contained in each directory is also displayed.

```
$ du /var/adm /var/spool/lp
2       /var/adm/exacct
2       /var/adm/log
2       /var/adm/streams
2       /var/adm/acct/fiscal
2       /var/adm/acct/nite
2       /var/adm/acct/sum
8       /var/adm/acct
2       /var/adm/sa
2       /var/adm/sm.bin
258     /var/adm
4       /var/spool/lp/admins
2       /var/spool/lp/requests/printing.Eng.Sun.COM
4       /var/spool/lp/requests
4       /var/spool/lp/system
2       /var/spool/lp/fifos
24      /var/spool/lp
```

The following example shows directory sizes in 1024–byte blocks.

```
$ du -h /usr/share/audio
 796K   /usr/share/audio/samples/au
 797K   /usr/share/audio/samples
 798K   /usr/share/audio
```

▼ How to Display the User Ownership of Local UFS File Systems

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Display users, directories, or file systems, and the number of 1024-byte blocks used.**

   ```
   # quot [-a] [filesystem ...]
   ```

   `-a` Lists all users of each mounted UFS file system and the number of 1024-byte blocks used.

   filesystem Identifies a UFS file system. Users and the number of blocks used are displayed for that file system.

Note – The `quot` command works only on local UFS file systems.

Example 13–8 Displaying the User Ownership of Local UFS File Systems

In the following example, users of the root (/) file system are displayed. In the subsequent example, users of all mounted UFS file systems are displayed.

```
# quot /
/dev/rdsk/c0t0d0s0:
43340   root
 3142   rimmer
   47   uucp
   35   lp
   30   adm
    4   bin
    4   daemon

# quot -a
/dev/rdsk/c0t0d0s0 (/):
43340   root
 3150   rimmer
   47   uucp
   35   lp
   30   adm
    4   bin
    4   daemon
/dev/rdsk/c0t0d0s6 (/usr):
460651  root
206632  bin
  791   uucp
   46   lp
    4   daemon
    1   adm
/dev/rdsk/c0t0d0s7 (/export/home):
    9   root
```

Finding and Removing Old or Inactive Files

Part of the job of cleaning up heavily loaded file systems involves locating and removing files that have not been used recently. You can locate unused files by using the `ls` or `find` commands. For more information, see the `ls`(1) and `find`(1) man pages.

Other ways to conserve disk space include emptying temporary directories such as the the directories located in `/var/tmp` or `/var/spool`, and deleting `core` and crash dump files. For more information about crash dump files, refer to Chapter 24.

▼ How to List the Newest Files

Step ● **List files, displaying the most recently created or changed files first, by using the `ls -t` command.**

```
$ ls -t [directory]
```

`-t` Sorts files by latest time stamp first.

directory Identifies the directory that you want to search.

Example 13–9 Listing the Newest Files

The following example shows how to use the `ls -tl` command to locate the most recently created or changed files within the `/var/adm` directory. The `sulog` file was created or edited most recently.

```
$ ls -tl /var/adm
total 134
-rw-------   1 root     root        315 Sep 24 14:00 sulog
-r--r--r--   1 root     other    350700 Sep 22 11:04 lastlog
-rw-r--r--   1 root     bin        4464 Sep 22 11:04 utmpx
-rw-r--r--   1 adm      adm       20088 Sep 22 11:04 wtmpx
-rw-r--r--   1 root     other         0 Sep 19 03:10 messages
-rw-r--r--   1 root     other         0 Sep 12 03:10 messages.0
-rw-r--r--   1 root     root      11510 Sep 10 16:13 messages.1
-rw-r--r--   1 root     root          0 Sep 10 16:12 vold.log
drwxr-xr-x   2 root     sys         512 Sep 10 15:33 sm.bin
drwxrwxr-x   5 adm      adm         512 Sep 10 15:19 acct
drwxrwxr-x   2 adm      sys         512 Sep 10 15:19 sa
-rw-------   1 uucp     bin           0 Sep 10 15:17 aculog
-rw-rw-rw-   1 root     bin           0 Sep 10 15:17 spellhist
drwxr-xr-x   2 adm      adm         512 Sep 10 15:17 log
drwxr-xr-x   2 adm      adm         512 Sep 10 15:17 passwd
```

▼ How to Find and Remove Old or Inactive Files

Steps 1. **Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Find files that have not been accessed for a specified number of days and list them in a file.**

```
# find directory -type f[-atime +nnn] [-mtime +nnn] -print > filename &
```

directory	Identifies the directory you want to search. Directories below this directory are also searched.
`-atime +`*nnn*	Finds files that have not been accessed within the number of days (nnn) that you specify.
`-mtime +`*nnn*	Finds files that have not been modified within the number of days (nnn) that you specify.
filename	Identifies the file that contains the list of inactive files.

3. **Remove the inactive files found listed in the previous step.**

```
# rm `cat filename`
```

where *filename* identifies the file that was created in the previous step. This file contains the list of inactive files.

Example 13–10 Finding and Removing Old or Inactive Files

The following example shows files in the `/var/adm` directory and the subdirectories that have not been accessed in the last 60 days. The `/var/tmp/deadfiles` file contains the list of inactive files. The `rm` command removes these inactive files.

```
# find /var/adm -type f -atime +60 -print > /var/tmp/deadfiles &
# more /var/tmp/deadfiles
/var/adm/aculog
/var/adm/spellhist
/var/adm/wtmpx
/var/adm/sa/sa13
/var/adm/sa/sa27
/var/adm/sa/sa11
/var/adm/sa/sa23
/var/adm/sulog
/var/adm/vold.log
/var/adm/messages.1
/var/adm/messages.2
/var/adm/messages.3
# rm `cat /var/tmp/deadfiles`
#
```

▼ How to Clear Out Temporary Directories

Steps 1. **Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2. **Change to the directory that you want to clean out.**

```
# cd directory
```

Caution – Ensure that you are in the correct directory before completing Step 3. Step 3 deletes all files in the current directory.

---

3. **Delete the files and subdirectories in the current directory.**

```
# rm -r *
```

4. **Change to other directories that contain unnecessary, temporary or obsolete subdirectories and files. Delete these subdirectories and files by repeating Step 3.**

Example 13–11 Clearing Out Temporary Directories

The following example shows how to clear out the `mywork` directory, and how to verify that all files and subdirectories were removed.

```
# cd mywork
# ls
filea.000
fileb.000
filec.001
# rm -r *
# ls
#
```

▼ How to Find and Delete `core` Files

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2. **Change to the directory where you want to search for `core` files.**

3. **Find and remove any `core` files in this directory and its subdirectories.**

```
# find . -name core -exec rm {} \;
```

Example 13–12 Finding and Deleting `core` Files

The following example shows how to find and remove `core` files from the `jones` user account by using the `find` command.

```
# cd /home/jones
# find . -name core -exec rm {} \;
```

▼ How to Delete Crash Dump Files

Crash dump files can be very large. If you have enabled your system to store these files, do not retain them for longer than necessary.

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Change to the directory where crash dump files are stored.**

   ```
   # cd /var/crash/system
   ```

   where *system* identifies a system that created the crash dump files.

Caution – Ensure you are in the correct directory before completing Step 3. Step 3 deletes all files in the current directory.

---

3. **Remove the crash dump files.**

   ```
   # rm *
   ```

4. **Verify that the crash dump files were removed.**

   ```
   # ls
   ```

Example 13–13 Deleting Crash Dump Files

The following example shows how to remove crash dump files from the system `venus`, and how to verify that the crash dump files were removed.

```
# cd /var/crash/venus
# rm *
# ls
```

CHAPTER 14

Managing Quotas (Tasks)

This chapter describes how to set up and administer quotas for disk space and inodes.

For information associated with managing quotas, see the following:

- "Setting Up Quotas (Task Map)" on page 262
- "Maintaining Quotas (Task Map)" on page 266

What Are Quotas?

Quotas enable system administrators to control the size of UFS file systems. Quotas limit the amount of disk space and the number of inodes, which roughly corresponds to the number of files, that individual users can acquire. For this reason, quotas are especially useful on the file systems where user home directories reside. As a rule, the `public` and `/tmp` file systems usually do not benefit significantly by establishing quotas.

Using Quotas

Once quotas are in place, they can be changed to adjust the amount of disk space or the number of inodes that users can consume. Additionally, quotas can be added or removed as system needs change. For instructions on changing quotas or the amount of time that quotas can be exceeded, disabling individual quotas, or removing quotas from file systems, see "Changing and Removing Quotas" on page 269.

In addition, quota status can be monitored. Quota commands enable administrators to display information about quotas on a file system, or search for users who have exceeded their quotas. For procedures that describe how to use these commands, see "Checking Quotas" on page 267.

Setting Soft Limits and Hard Limits for Quotas

You can set both soft limits and hard limits. The system does not allow a user to exceed his or her hard limit. However, a system administrator might set a soft limit, which the user can temporarily exceed. The soft limit must be less than the hard limit.

Once the user exceeds the soft limit, a quota timer begins. While the quota timer is ticking, the user is allowed to operate above the soft limit but cannot exceed the hard limit. Once the user goes below the soft limit, the timer is reset. However, if the user's usage remains above the soft limit when the timer expires, the soft limit is enforced as a hard limit. By default, the soft limit timer is set to seven days.

The `timeleft` field in the `repquota` and `quota` commands shows the value of the timer.

For example, let's say a user has a soft limit of 10,000 blocks and a hard limit of 12,000 blocks. If the user's block usage exceeds 10,000 blocks and the seven-day timer is also exceeded, the user cannot allocate more disk blocks on that file system until his or her usage drops below the soft limit.

The Difference Between Disk Block and File Limits

A file system provides two resources to the user, blocks for data and inodes for files. Each file consumes one inode. File data is stored in data blocks. Data blocks are usually made up of 1Kbyte blocks.

Assuming no directories exist, a user can exceed his or her inode quota by creating all empty files without using any blocks. A user can also use one inode, yet exceed his or her block quota, by creating one file that is large enough to consume all the data blocks in the user's quota.

Setting Up Quotas

Setting up quotas involves these general steps:

1. Ensuring that quotas are enforced each time the system is rebooted by adding a quota option to the `/etc/vfstab` file entries. Also, creating a `quotas` file in the top-level directory of the file system.

2. After you create a quota for one use, copying the quota as a prototype to set up other user quotas.
3. Before you turn quotas on, checking the consistency of the proposed quotas with the current disk usage to make sure that there are no conflicts.
4. Turning on the quotas on for one or more file systems.

For specific information about these procedures, see "Setting Up Quotas (Task Map)" on page 262.

The following table describes the commands that you use to set up disk quotas.

TABLE 14–1 Commands for Setting Up Quotas

Command	Task	Man Page
`edquota`	Sets the hard limits and soft limits on the number of inodes and the amount of disk space for each user.	`edquota`(1M)
`quotacheck`	Examines each mounted UFS file system, comparing the file system's current disk usage against information stored in the file system's disk quota file. Then, resolves inconsistencies	`quotacheck`(1M)
`quotaon`	Activates the quotas for the specified file systems.	`quotaon`(1M)
`quota`	Displays users' disk quotas on mounted file systems to verify that the quotas have been correctly set up.	`quota`(1M)

Guidelines for Setting Up Quotas

Before you set up quotas, you need to determine how much disk space and how many inodes to allocate to each user. If you want to ensure that the total file system space is never exceeded, you can divide the total size of the file system between the number of users. For example, if three users share a 100-Mbyte slice and have equal disk space needs, you could allocate 33 Mbytes to each user.

In environments where not all users are likely to push their limits, you might want to set individual quotas so that they add up to more than the total size of the file system. For example, if three users share a 100-Mbyte slice, you could allocate 40 Mbytes to each user.

When you have established a quota for one user by using the `edquota` command, you can use this quota as a prototype to set the same quota for other users on the same file system.

Before you turn on the quotas, do the following:

- First, configure the UFS file systems for the quotas.
- Establish quotas for each user, and run the `quotacheck` command to check for consistency between current disk usage and quota files.
- Run the `quotacheck` command periodically if systems are rebooted infrequently.

The quotas you set up with the `edquota` command are not enforced until you turn them on by using the `quotaon` command. If you have properly configured the quota files, the quotas are turned on automatically each time a system is rebooted and the file system is mounted.

Setting Up Quotas (Task Map)

Task	Description	For Instructions
1. Configure a file system for quotas.	Edit the `/etc/vfstab` file so that quotas are activated each time the file system is mounted. Also, create a `quotas` file.	"How to Configure File Systems for Quotas" on page 262
2. Set up quotas for a user.	Use the `edquota` command to create disk quotas and inode quotas for a single user account.	"How to Set Up Quotas for a User" on page 263
3. (Optional) Set up quotas for multiple users.	Use the `edquota` command to apply prototype quotas to other user accounts.	"How to Set Up Quotas for Multiple Users" on page 264
4. Check for consistency.	Use the `quotacheck` command to compare quotas to current disk usage for consistency across one or more file systems.	"How to Check Quota Consistency" on page 265
5. Turn on quotas.	Use the `quotaon` command to initiate quotas on one or more file systems.	"How to Turn On Quotas" on page 266

▼ How to Configure File Systems for Quotas

Steps 1. **Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Edit the `/etc/vfstab` file and add `rq` to the `mount options` field for each UFS file system that will have quotas.**

3. **Change directory to the root of the file system that will have quotas.**

4. **Create a file named `quotas`.**

```
# touch quotas
```

5. **Change permissions to read/write for superuser access only.**

```
# chmod 600 quotas
```

Example 14–1 Configuring File Systems for Quotas

The following `/etc/vfstab` example shows that the `/export/home` directory from the system `pluto` is mounted as an NFS file system on the local system. You can tell that quotas are enabled by the `rq` entry under the `mount options` column.

```
# device    device    mount     FS    fsck   mount   mount
# to mount       to fsck  point     type   pass   at boot options
# pluto:/export/home -   /export/home nfs    -     yes    rq
```

The following example line from the `/etc/vfstab` file shows that the local `/work` directory is mounted with quotas enabled, signified by the `rq` entry under the `mount options` column.

```
#device     device        mount  FS   fsck  mount    mount
#to mount         to fsck      point  type  pass  at boot options
#/dev/dsk/c0t4d0s0 /dev/rdsk/c0t4d0s0 /work ufs  3    yes     rq
```

See Also

- "How to Set Up Quotas for a User" on page 263
- "How to Set Up Quotas for Multiple Users" on page 264
- "How to Check Quota Consistency" on page 265
- "How to Turn On Quotas" on page 266

▼ How to Set Up Quotas for a User

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Use the quota editor to create a temporary file that contains one line of quota information for each mounted UFS file system that has a `quotas` file in the file system's root directory.**

```
# edquota username
```

where *username* is the user for whom you want to set up quotas.

3. **Change the number of 1-Kbyte disk blocks, both soft and hard, and the number of inodes, both soft and hard, from the default of 0, to the quotas that you specify for each file system.**

4. **Verify the user's quota.**

```
# quota -v username
```

`-v`	Displays the user's quota information on all mounted file systems where quotas exist.
username	Specifies the user name to view quota limits.

Example 14–2 Setting Up Quotas for a User

The following example shows the contents of the temporary file opened by `edquota` on a system where `/files` is the only mounted file system that contains a `quotas` file in the root directory.

```
fs /files blocks (soft = 0, hard = 0) inodes (soft = 0, hard = 0)
```

The following example shows the same line in the temporary file after quotas have been set up.

```
fs /files blocks (soft = 50, hard = 60) inodes (soft = 90, hard = 100)
```

▼ How to Set Up Quotas for Multiple Users

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Use the quota editor to apply the quotas you already established for a prototype user to the additional users that you specify.**

```
# edquota -p prototype-user username ...
```

prototype-user	Is the user name of the account for which you have set up quotas. More than one user name is specified by separating each user name with a space.
username ...	Specifies one or more user names of additional accounts.

Example 14–3 Setting Up Prototype Quotas for Multiple Users

The following example shows how to apply the quotas established for user `bob` to users `mary` and `john`.

```
# edquota -p bob mary john
```

▼ How to Check Quota Consistency

The `quotacheck` command is run automatically when a system is rebooted. You generally do not have to run the `quotacheck` command on an empty file system with quotas. However, if you are setting up quotas on a file system with existing files, you need to run the `quotacheck` command to synchronize the quota database with the files or inodes that already exist in the file system.

Also keep in mind that running the `quotacheck` command on large file systems can be time-consuming.

Note – To ensure accurate disk data, the file systems being checked should be quiescent when you run the `quotacheck` command manually.

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Run a consistency check on UFS file systems.**

   ```
   # quotacheck [-va] filesystem
   ```

   `-v` (Optional) Identifies the disk quotas for each user on a particular file system.

   `-a` Checks all file systems with an `rq` entry in the `/etc/vfstab` file.

   filesystem Specifies the file system to check.

   See the `quotacheck`(1M) man page for more information.

Example 14–4 Checking Quota Consistency

The following example shows how to check quotas for the `/export/home` file system on the `/dev/rdsk/c0t0d0s7` slice. The `/export/home` file system is the only file system with an `rq` entry in the `/etc/vfstab` file.

```
# quotacheck -va
*** Checking quotas for /dev/rdsk/c0t0d0s7 (/export/home)
```

▼ How to Turn On Quotas

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Turn on file system quotas.**

   ```
   # quotaon [-v] -a filesystem ...
   ```

`-v`	Displays a message for each file system after quotas are turned on.
`-a`	Turns on quotas for all file systems with an `rq` entry in the `/etc/vfstab` file.
filesystem ...	Turns on quotas for one or more file systems that you specify. More than one file system is specified by separating each file system name with a space.

Example 14–5 Turning On Quotas

The following example shows how to turn quotas on for the file systems on the `/dev/dsk/c0t4d0s7` and `/dev/dsk/c0t3d0s7` slices.

```
# quotaon -v /dev/dsk/c0t4d0s7 /dev/dsk/c0t3d0s7
/dev/dsk/c0t4d0s7: quotas turned on
/dev/dsk/c0t3d0s7: quotas turned on
```

Maintaining Quotas (Task Map)

Task	Description	For Instructions
Check for exceeded quotas.	Display the quotas and disk use for individual users on file systems on which quotas have been activated by using the `quota` command.	"How to Check for Exceeded Quotas" on page 267
Check for quotas on a file system.	Display the quotas and disk use for all users on one or more file systems by using the `repquota` command.	"How to Check Quotas on a File System" on page 268

Task	Description	For Instructions
Change the soft limit default.	Change the length of time that users can exceed their disk space quotas or inode quotas by using the edquota command.	"How to Change the Soft Limit Default" on page 270
Change quotas for a user.	Use the quota editor, edquota, to change quotas for an individual user.	"How to Change Quotas for a User" on page 271
Disable quotas for a user.	Use the quota editor, edquota, to disable quotas for an individual user.	"How to Disable Quotas for a User" on page 272
Turn off quotas.	Turn off quotas by using the quotaoff command.	"How to Turn Off Quotas" on page 273

Checking Quotas

After you have set up and turned on disk quotas and inode quotas, you can check for users who exceed their quotas. In addition, you can check quota information for entire file systems.

The following table describes the commands that you use to check quotas.

TABLE 14–2 Commands for Checking Quotas

Command	Task
quota(1M)	Displays user quotas and current disk use, and information about users who are exceeding their quotas
repquota(1M)	Displays quotas, files, and the amount of space that is owned for specified file systems

▼ How to Check for Exceeded Quotas

You can display the quotas and disk use for individual users on file systems on which quotas have been activated by using the quota command.

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Display user quotas for mounted file systems where quotas are enabled.**

   ```
   # quota [-v] username
   ```

`-v`	Displays one or more users' quotas on all mounted file systems that have quotas.
username	Is the login name or UID of a user's account.

Example 14–6 Checking for Exceeded Quotas

The following example shows that the user account identified by UID 301 has one 1–Kbyte quota but has not used any disk space.

```
# quota -v 301
Disk quotas for bob (uid 301):
Filesystem  usage  quota limit timeleft files quota  limit timeleft
/export/home   0      1     2              0      2      3
```

`Filesystem`	Is the mount point for the file system.
`usage`	Is the current block usage.
`quota`	Is the soft-block limit.
`limit`	Is the hard-block limit.
`timeleft`	Is the amount of time, in days, left on the quota timer.
`files`	Is the current inode usage.
`quota`	Is the soft-inode limit.
`limit`	Is the hard-inode limit.
`timeleft`	Is the amount of time, in days, left on the quota timer.

▼ How to Check Quotas on a File System

Display the quotas and disk use for all users on one or more file systems by using the `repquota` command.

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Display all quotas for one or more file systems, even if there is no usage.**

   ```
   # repquota [-v]  -a filesystem
   ```

`-v`	Reports on quotas for all users, even those users who do not consume resources.
`-a`	Reports on all file systems.
filesystem	Reports on the specified file system.

Example 14–7 Checking Quotas on a File System

The following example shows output from the `repquota` command on a system that has quotas enabled on only one file system (`/export/home`).

```
# repquota -va
/dev/dsk/c0t3d0s7 (/export/home):
                  Block limits              File limits
User      used   soft   hard  timeleft  used  soft   hard  timeleft
#301  --           0      1   2.0 days         0      2      3
#341  --    57     50     60   7.0 days  2      90    100
```

Block limits	Definition
`used`	Is the current block usage.
`soft`	Is the soft-block limit.
`hard`	Is the hard-block limit.
`timeleft`	Is the amount of time, in days, left on the quota timer.

File limits	Definition
`used`	Is the current inode usage.
`soft`	Is the soft-inode limit.
`hard`	Is the hard-inode limit.
`timeleft`	Is the amount of time, in days, left on the quota timer.

Changing and Removing Quotas

You can change quotas to adjust the amount of disk space or the number of inodes that users can consume. You can also remove quotas, for individual users or from entire file systems, as needed.

The following table describes the commands that you use to change quotas or to remove quotas.

TABLE 14–3 Commands for Changing Quotas and Removing Quotas

Command	Man Page	Description
`edquota`	`edquota`(1M)	Changes the hard limits and soft limits on the number of inodes or amount of disk space for each user. Also, changes the soft limit for each file system with a quota.
`quotaoff`	`quotaon`(1M)	Turns off quotas for specified file systems.

▼ How to Change the Soft Limit Default

By default, users can exceed the soft time limits for their quotas for one week. So, after a week of repeated violations of the soft time limits of either disk space quotas or inode quotas, the system prevents users from using any more inodes or disk blocks.

You can change the length of time that users can exceed their disk space quotas or inode quotas by using the `edquota` command.

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Use the quota editor to create a temporary file that contains soft time limits.**

   ```
   # edquota -t
   ```

   where the `-t` option specifies the editing of the soft time limits for each file system.

3. **Change the time limits from 0 (the default) to the time limits that you specify. So, use numbers and the keywords `month`, `week`, `day`, `hour`, `min`, or `sec`.**

   Note – This procedure does not affect current quota violators.

Example 14–8 Changing the Soft Limit Default

The following example shows the contents of the temporary file opened by the `edquota` command on a system where `/export/home` is the only mounted file system with quotas. The default value, `0`, means that the default time limit of one week is used.

```
fs /export/home blocks time limit = 0 (default), files time limit = 0 (default)
```

The following example shows the same temporary file after the time limit for exceeding the blocks quota has been changed to 2 weeks. Also, the time limit for exceeding the number of files has been changed to 16 days.

```
fs /export/home blocks time limit = 2 weeks, files time limit = 16 days
```

▼ How to Change Quotas for a User

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Use the quota editor to open a temporary file that contains one line for each mounted file system that has a `quotas` file in the file system's root directory.**

   ```
   # edquota username
   ```

   where *username* specifies the user name whose quota you want to change.

   Caution – You can specify multiple users as arguments to the `edquota` command. However, the user that this information belongs to, is not displayed. To avoid confusion, specify only one user name.

   ---

3. **Specify the number of 1-Kbyte disk blocks, both soft and hard, and the number of inodes, both soft and hard.**

4. **Verify that a user's quota has been correctly changed.**

   ```
   # quota -v username
   ```

   `-v` Displays user quota information on all mounted file systems with quotas enabled.

   username Specifies the user name whose quota you want to check.

Example 14–9 Changing Quotas for a User

The following example shows the contents of the temporary file opened by the `edquota` command. This temporary file is opened on a system where `/files` is the only mounted file system containing a `quotas` file in the file system's root directory.

```
fs /files blocks (soft = 0, hard = 0) inodes (soft = 0, hard = 0)
```

The following output shows the same temporary file after quotas have been changed.

```
fs /files blocks (soft = 0, hard = 500) inodes (soft = 0, hard = 100)
```

Example 14–10 Verifying That Hard Quotas Have Been Changed

The following example shows how to verify that the hard quotas for user `smith` have been changed to 500 1-Kbyte blocks, and 100 inodes.

```
# quota -v smith
Disk quotas for smith (uid 12):
Filesystem  usage  quota  limit  timeleft  files  quota  limit  timeleft

  /files     1       0     500              1       0     100
```

▼ How to Disable Quotas for a User

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Use the quota editor to create a temporary file containing one line for each mounted file system that has a `quotas` file in its top-level directory.**

   ```
   # edquota username
   ```

   Where *username* specifies the user name whose quota you want to disable.

   Caution – You can specify multiple users as arguments to the `edquota` command. However, the user that this information belongs to, is not displayed. To avoid confusion, specify only one user name.

   ---

3. **Change the number of 1-Kbyte disk blocks, both soft and hard, and the number of inodes, both soft and hard, to 0.**

   Note – Ensure that you change the values to zero. Do *not* delete the line from the text file.

   ---

4. **Verify that you have disabled a user's quota.**

   ```
   # quota -v username
   ```

`-v`	Displays user quota information on all mounted file systems with quotas enabled.
username	Specifies the user name (UID) whose quota you want to check.

Example 14–11 Disabling Quotas for a User

The following example shows the contents of the temporary file opened by the `edquota` command on a system where `/files` is the only mounted file system that contains a quotas file in the file system's root directory.

```
fs /files blocks (soft = 50, hard = 60) inodes (soft = 90, hard = 100)
```

The following example shows the same temporary file after quotas have been disabled.

```
fs /files blocks (soft = 0, hard = 0) inodes (soft = 0, hard = 0)
```

▼ How to Turn Off Quotas

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Turn off file system quotas.**

   ```
   # quotaoff [-v] -a filesystem ...
   ```

   `-v` Displays a message from each file system when quotas are turned off.

   `-a` Turns off quotas for all file systems.

   filesystem Turns off quotas for one or more file systems that you specify. More than one file system is specified by separating each file system name with a space.

Example 14–12 Turning Off Quotas

The following example shows how to turn off the quotas for the `/export/home` file system.

```
# quotaoff -v /export/home
/export/home: quotas turned off
```

CHAPTER 15

Scheduling System Tasks (Tasks)

This chapter describes how to schedule routine or single (one-time) system tasks by using the crontab and at commands. This chapter also explains how to control access to these commands by using the cron.deny, cron.allow, and at.deny files.

For information on the procedures associated with scheduling system tasks, see the following:

- "Creating and Editing crontab Files (Task Map)" on page 275
- "Using the at Command (Task Map)" on page 288

Creating and Editing crontab Files (Task Map)

Task	Description	For Instructions
Create or edit a crontab file.	Use the crontab -e command to create or edit a crontab file.	"How to Create or Edit a crontab File" on page 281
Verify that a crontab file exists.	Use the ls -l command to verify the contents of the /var/spool/cron/crontabs file.	"How to Verify That a crontab File Exists" on page 282
Display a crontab file.	Use the ls -l command to display the crontab file.	"How to Display a crontab File" on page 283

Task	Description	For Instructions
Remove a `crontab` file	The `crontab` file is set up with restrictive permissions Use the `crontab -r` command, rather than the `rm` command to remove a `crontab` file.	"How to Remove a `crontab` File" on page 284
Deny `crontab` access	To deny users access to `crontab` commands, add user names to the `/etc/cron.d/cron.deny` file by editing this file.	"How to Deny `crontab` Command Access" on page 286
Limit `crontab` access to specified users.	To allow users access to the `crontab` command, add user names to the `/etc/cron.d/cron.allow` file.	"How to Limit `crontab` Command Access to Specified Users" on page 287

Ways to Automatically Execute System Tasks

You can set up many system tasks to execute automatically. Some of these tasks should occur at regular intervals. Other tasks need to run only once, perhaps during off hours such as evenings or weekends.

This section contains overview information about two commands, `crontab` and `at`, which enable you to schedule routine tasks to execute automatically. The `crontab` command schedules repetitive commands. The `at` command schedules tasks that execute once.

The following table summarizes `crontab` and `at` commands, as well as the files that enable you to control access to these commands.

TABLE 15–1 Command Summary: Scheduling System Tasks

Command	What It Schedules	Location of Files	Files That Control Access
`crontab`	Multiple system tasks at regular intervals	`/var/spool/cron/crontabs`	`/etc/cron.d/cron.allow` and `/etc/cron.d/cron.deny`
`at`	A single system task	`/var/spool/cron/atjobs`	`/etc/cron.d/at.deny`

You can also use the Solaris Management Console's Scheduled Jobs tool to schedule routine tasks. For information on using and starting the Solaris Management Console, see Chapter 2, "Working With the Solaris Management Console (Tasks)," in *System Administration Guide: Basic Administration*.

For Scheduling Repetitive Jobs: `crontab`

You can schedule routine system administration tasks to execute daily, weekly, or monthly by using the `crontab` command.

Daily `crontab` system administration tasks might include the following:

- Removing files more than a few days old from temporary directories
- Executing accounting summary commands
- Taking snapshots of the system by using the `df` and `ps` commands
- Performing daily security monitoring
- Running system backups

Weekly `crontab` system administration tasks might include the following:

- Rebuilding the `catman` database for use by the `man -k` command
- Running the `fsck -n` command to list any disk problems

Monthly `crontab` system administration tasks might include the following:

- Listing files not used during a specific month
- Producing monthly accounting reports

Additionally, users can schedule `crontab` commands to execute other routine system tasks, such as sending reminders and removing backup files.

For step-by-step instructions on scheduling `crontab` jobs, see "How to Create or Edit a `crontab` File" on page 281.

For Scheduling a Single Job: `at`

The `at` command allows you to schedule a job for execution at a later time. The job can consist of a single command or a script.

Similar to `crontab`, the `at` command allows you to schedule the automatic execution of routine tasks. However, unlike `crontab` files, `at` files execute their tasks once. Then, they are removed from their directory. Therefore, the `at` command is most useful for running simple commands or scripts that direct output into separate files for later examination.

Submitting an `at` job involves typing a command and following the `at` command syntax to specify options to schedule the time your job will be executed. For more information about submitting `at` jobs, see "Description of the `at` Command" on page 289.

The at command stores the command or script you ran, along with a copy of your current environment variable, in the /var/spool/cron/atjobs directory. Your at job file name is given a long number that specifies its location in the at queue, followed by the .a extension, such as 793962000.a.

The cron daemon periodically executes the atrun program, usually at 15-minute intervals. Then the atrun program executes at jobs at their scheduled times. After the cron daemon executes an at job, the at job's file is removed from the atjobs directory.

For step-by-step instructions on scheduling at jobs, see "How to Create an at Job" on page 290.

Scheduling a Repetitive System Task (cron)

The following sections describe how to create, edit, display, and remove crontab files, as well as how to control access to them.

Inside a crontab File

The cron daemon schedules system tasks according to commands found within each crontab file. A crontab file consists of commands, one command per line, that will be executed at regular intervals. The beginning of each line contains date and time information that tells the cron daemon when to execute the command.

For example, a crontab file named root is supplied during SunOS software installation. The file's contents include these command lines:

```
10 3 * * * /usr/sbin/logadm     (1)
15 3 * * 0 /usr/lib/fs/nfs/nfsfind     (2)
1 2 * * * [ -x /usr/sbin/rtc ] && /usr/sbin/rtc -c > /dev/null 2>&1     (3)
30 3 * * * [ -x /usr/lib/gss/gsscred_clean ] && /usr/lib/gss/gsscred_clean     (4)
```

The following describes the output for each of these command lines:

- The first line runs the logadm command at 3:10 a.m. every day.
- The second line executes the nfsfind script every Sunday at 3:15 a.m.
- The third line runs a script that checks for daylight savings time (and make corrections, if necessary) at 2:10 a.m. daily.

  If there is no RTC time zone, nor an /etc/rtc_config file, this entry does nothing.

x86 only – The /usr/sbin/rtc script can only be run on an x86 based system.

- The fourth line checks for (and removes) duplicate entries in the Generic Security Service table, /etc/gss/gsscred_db, at 3:30 a.m. daily.

For more information about the syntax of lines within a crontab file, see "Syntax of crontab File Entries" on page 280.

The crontab files are stored in the /var/spool/cron/crontabs directory. Several crontab files besides root are provided during SunOS software installation. See the following table:

TABLE 15–2 Default crontab Files

crontab File	Function
adm	Accounting
lp	Printing
root	General system functions and file system cleanup
sys	Performance data collection
uucp	General uucp cleanup

Besides the default crontab files, users can create crontab files to schedule their own system tasks. Other crontab files are named after the user accounts in which they are created, such as bob, mary, smith, or jones.

To access crontab files that belong to root or other users, superuser privileges are required.

Procedures explaining how to create, edit, display, and remove crontab files are described in subsequent sections.

How the cron Daemon Handles Scheduling

The cron daemon manages the automatic scheduling of crontab commands. The role of the cron daemon is to check the /var/spool/cron/crontab directory for the presence of crontab files, normally every 15 minutes. The cron daemon performs the following tasks:

- Checks for new crontab files or changes to existing files
- Reads the execution times that are listed within the files
- Submits the commands for execution at the proper times

In much the same way, the cron daemon controls the scheduling of at files. These files are stored in the /var/spool/cron/atjobs directory.

Syntax of crontab File Entries

A crontab file consists of commands, one command per line, that execute automatically at the time specified by the first five fields of each command line. These five fields, described in the following table, are separated by spaces.

TABLE 15–3 Acceptable Values for crontab Time Fields

Time Field	Values
Minute	0-59
Hour	0-23
Day of month	1-31
Month	1-12
Day of week	0-6 (0 = Sunday)

Follow these guidelines for using special characters in crontab time fields:

- Use a space to separate each field.
- Use a comma to separate multiple values.
- Use a hyphen to designate a range of values.
- Use an asterisk as a wildcard to include all possible values.
- Use a comment mark (#) at the beginning of a line to indicate a comment or a blank line.

For example, the following crontab command entry displays a reminder in the user's console window at 4 p.m. on the first and fifteenth days of every month.

```
0 16 1,15 * * echo Timesheets Due > /dev/console
```

Each command within a crontab file must consist of one line, even if that line is very long. The crontab file does not recognize extra carriage returns. For more detailed information about crontab entries and command options, refer to the crontab(1) man page.

Creating and Editing crontab Files

The simplest way to create a `crontab` file is to use the `crontab -e` command. This command invokes the text editor that has been set for your system environment. The default editor for your system environment is defined in the `EDITOR` environment variable. If this variable has not been set, the `crontab` command uses the default editor, `ed`. Preferably, you should choose an editor that you know well.

The following example shows how to determine if an editor has been defined, and how to set up `vi` as the default.

```
$ which $EDITOR
$
$ EDITOR=vi
$ export EDITOR
```

When you create a `crontab` file, it is automatically placed in the `/var/spool/cron/crontabs` directory and is given your user name. You can create or edit a `crontab` file for another user, or `root`, if you have superuser privileges.

▼ How to Create or Edit a crontab File

Before You Begin If you are creating or editing a `crontab` file that belongs to `root` or another user you must become superuser or assume an equivalent role. Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*

You do not need to become superuser to edit your own `crontab`file.

Steps

1. **Create a new `crontab` file, or edit an existing file.**

   ```
   $ crontab -e [username]
   ```

   where *username* specifies the name of the user's account for which you want to create or edit a `crontab` file. You can create your own `crontab` file without superuser privileges, but you must have superuser privileges to creating or edit a `crontab` file for `root` or another user.

   Caution – If you accidentally type the `crontab` command with no option, press the interrupt character for your editor. This character allows you to quit without saving changes. If you instead saved changes and exited the file, the existing `crontab` file would be overwritten with an empty file.

   ---

2. **Add command lines to the `crontab` file.**

Follow the syntax described in "Syntax of crontab File Entries" on page 280. The crontab file will be placed in the /var/spool/cron/crontabs directory.

3. **Verify your crontab file changes.**

```
# crontab -l [username]
```

Example 15–1 Creating a crontab File

The following example shows how to create a crontab file for another user.

```
# crontab -e jones
```

The following command entry added to a new crontab file automatically removes any log files from the user's home directory at 1:00 a.m. every Sunday morning. Because the command entry does not redirect output, redirect characters are added to the command line after *.log. Doing so ensures that the command executes properly.

```
# This command helps clean up user accounts.
1 0 * * 0 rm /home/jones/*.log > /dev/null 2>&1
```

▼ How to Verify That a crontab File Exists

Step ● **To verify that a crontab file exists for a user, use the ls -l command in the /var/spool/cron/crontabs directory. For example, the following output shows that crontab files exist for users jones and smith.**

```
$ ls -l /var/spool/cron/crontabs
-rw-r--r--  1 root     sys          190 Feb 26 16:23 adm
-rw-------  1 root     staff        225 Mar  1  9:19 jones
-rw-r--r--  1 root     root        1063 Feb 26 16:23 lp
-rw-r--r--  1 root     sys          441 Feb 26 16:25 root
-rw-------  1 root     staff         60 Mar  1  9:15 smith
-rw-r--r--  1 root     sys          308 Feb 26 16:23 sys
```

Verify the contents of user's crontab file by using the crontab -l command as described in "How to Display a crontab File" on page 283.

Displaying crontab Files

The crontab -l command displays the contents of a crontab file much the same way that the cat command displays the contents of other types of files. You do not have to change the directory to /var/spool/cron/crontabs directory (where crontab files are located) to use this command.

By default, the `crontab -l` command displays your own `crontab` file. To display `crontab` files that belong to other users, you must be superuser.

▼ How to Display a `crontab` File

Before You Begin Become superuser or assume an equivalent role to display a `crontab` file that belongs to `root` or another user. Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

You do not need to become superuser or assume an equivalent role to display your own `crontab` file.

Step ● **Display the `crontab` file.**

```
$ crontab -l [username]
```

where *username* specifies the name of the user's account for which you want to display a `crontab` file. Displaying another user's `crontab` file requires superuser privileges.

Caution – If you accidentally type the `crontab` command with no option, press the interrupt character for your editor. This character allows you to quit without saving changes. If you instead saved changes and exited the file, the existing `crontab` file would be overwritten with an empty file.

Example 15–2 Displaying a `crontab` File

This example shows how to use the `crontab -l` command to display the contents of the user's default `crontab` file.

```
$ crontab -l
13 13 * * * chmod g+w /home1/documents/*.book > /dev/null 2>&1
```

Example 15–3 Displaying the Default `root` `crontab` file.

This example shows how to display the default `root` `crontab` file.

```
$ suPassword:
Sun Microsystems Inc.   SunOS 5.10      s10_51  May 2004
# crontab -l
#ident  "@(#)root       1.19    98/07/06 SMI"   /* SVr4.0 1.1.3.1       */
#
# The root crontab should be used to perform accounting data collection.
#
#
10 3 * * * /usr/sbin/logadm
```

```
15 3 * * 0 /usr/lib/fs/nfs/nfsfind
30 3 * * * [ -x /usr/lib/gss/gsscred_clean ] && /usr/lib/gss/gsscred_clean
#10 3 * * * /usr/lib/krb5/kprop_script ___slave_kdcs___
```

Example 15–4 Displaying the `crontab` File of Another User

This example shows how to display the `crontab` file that belongs to another user.

```
$ su
Password:
Sun Microsystems Inc.   SunOS 5.10      s10_51  May 2004
# crontab -l jones
13 13 * * * cp /home/jones/work_files /usr/backup/. > /dev/null 2>&1
```

Removing `crontab` Files

By default, `crontab` file protections are set up so that you cannot inadvertently delete a `crontab` file by using the `rm` command. Instead, use the `crontab -r` command to remove `crontab` files.

By default, the `crontab -r` command removes your own `crontab` file.

You do not have to change the directory to `/var/spool/cron/crontabs` (where `crontab` files are located) to use this command.

▼ How to Remove a `crontab` File

Before You Begin Become superuser or assume an equivalent role to display a `crontab` file that belongs to `root` or another user. Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

You do not need to become superuser or assume an equivalent role to display your own `crontab` file.

Steps

1. **Remove the `crontab` file.**

   ```
   $ crontab -r [username]
   ```

   where *username* specifies the name of the user's account for which you want to remove a `crontab` file. Removing `crontab` files for another user requires superuser privileges.

Caution – If you accidentally type the `crontab` command with no option, press the interrupt character for your editor. This character allows you to quit without saving changes. If you instead saved changes and exited the file, the existing `crontab` file would be overwritten with an empty file.

2. **Verify that the `crontab` file has been removed.**

```
# ls /var/spool/cron/crontabs
```

Example 15–5 Removing a `crontab` File

The following example shows how user `smith` uses the `crontab -r` command to remove his `crontab` file.

```
$ ls /var/spool/cron/crontabs
adm     jones     lp     root     smith    sys     uucp
$ crontab -r
$ ls /var/spool/cron/crontabs
adm     jones     lp     root     sys    uucp
```

Controlling Access to the `crontab` Command

You can control access to the `crontab` command by using two files in the `/etc/cron.d` directory: `cron.deny` and `cron.allow`. These files permit only specified users to perform `crontab` command tasks such as creating, editing, displaying, or removing their own `crontab` files.

The `cron.deny` and `cron.allow` files consist of a list of user names, one user name per line. These access control files work together as follows:

- If `cron.allow` exists, only the users who are listed in this file can create, edit, display, or remove `crontab` files.
- If `cron.allow` does not exist, all users can submit `crontab` files, except for users who are listed in `cron.deny`.
- If neither `cron.allow` nor `cron.deny` exists, superuser privileges are required to run the `crontab` command.

Superuser privileges are required to edit or create the `cron.deny` and `cron.allow` files.

The `cron.deny` file, which is created during SunOS software installation, contains the following user names:

```
$ cat /etc/cron.d/cron.deny
daemon
bin
smtp
nuucp
listen
nobody
noaccess
```

None of the user names in the default `cron.deny` file can access the `crontab` command. You can edit this file to add other user names that will be denied access to the `crontab` command.

No default `cron.allow` file is supplied. So, after Solaris software installation, all users (except users who are listed in the default `cron.deny` file) can access the `crontab` command. If you create a `cron.allow` file, only these users can access the `crontab` command.

▼ How to Deny `crontab` Command Access

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Edit the `/etc/cron.d/cron.deny` file and add user names, one user per line. Include users who will be denied access to the `crontab` commands.**

   ```
   daemon
   bin
   smtp
   nuucp
   listen
   nobody
   noaccess
   username1
   username2
   username3
   .
   .
   .
   ```

3. **Verify that the `/etc/cron.d/cron.deny` file contains the new entries..**

   ```
   # cat /etc/cron.d/cron.deny
   daemon
   bin
   nuucp
   listen
   nobody
   ```

```
noaccess
```

▼ How to Limit `crontab` Command Access to Specified Users

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Create the `/etc/cron.d/cron.allow` file.**

3. **Add the `root` user name into the `cron.allow` file.**

   If you do not add `root` to the file, superuser access to `crontab` commands will be denied.

4. **Add the user names, one user name per line. Include users that will be allowed to use the `crontab` command.**

   `root`
   username1
   username2
   username3
   .
   .
   .

Example 15–6 Limiting `crontab` Command Access to Specified Users

The following example shows a `cron.deny` file that prevents user names `jones`, `temp`, and `visitor` from accessing the `crontab` command.

```
$ cat /etc/cron.d/cron.deny
daemon
bin
smtp
nuucp
listen
nobody
noaccess
jones
temp
visitor
```

The following example shows a `cron.allow` file. The users `root`, `jones`, `lp`, and `smith` are the only users who can access the `crontab` command.

```
$ cat /etc/cron.d/cron.allow
root
jones
```

```
lp
smith
```

How to Verify Limited crontab Command Access

To verify if a specific user can access the crontab command, use the crontab -l command while you are logged into the user account.

```
$ crontab -l
```

If the user can access the crontab command, and already has created a crontab file, the file is displayed. Otherwise, if the user can access the crontab command but no crontab file exists, a message similar to the following message is displayed:

```
crontab: can't open your crontab file
```

Either this user either is listed in the cron.allow file (if the file exists), or the user is not listed in the cron.deny file.

If the user cannot access the crontab command, the following message is displayed whether or not a previous crontab file exists:

```
crontab: you are not authorized to use cron. Sorry.
```

This message means that either the user is not listed in the cron.allow file (if the file exists), or the user is listed in the cron.deny file.

Using the at Command (Task Map)

Task	Description	For Instructions
Create an at job.	Use the atcommand to do the following: ■ Start the at utility from the command line. ■ Type the commands or scripts that you want to execute, one per line. ■ Exit the at utiltily and save the job.	"How to Create an at Job" on page 290

Task	Description	For Instructions
Display the at queue.	User the atq command to display the at queue.	"How to Display the at Queue" on page 291
Verify an at job.	Use the atq command to confirm that at jobs that belong to a specific user have been submitted to the queue.	"How to Verify an at Job" on page 292
Display at jobs.	Use the at -l[*job-id* to display at jobs. that have been submitted to the queue.	"How to Display at Jobs" on page 292
Remove at jobs.	Use the at -r [*job-id*] command to remove at jobs from the queue.	"How to Remove at Jobs" on page 292
Deny access to the at command.	To deny users access to the atcommand, edit the /etc/cron.d/at.deny file.	"How to Deny Access to the at Command" on page 293

Scheduling a Single System Task (at)

The following sections describe how to use the at command to perform the following tasks:

- Schedule jobs (command and scripts) for execution at a later time
- How to display and remove these jobs
- How to control access to the at command

By default, users can create, display, and remove their own at job files. To access at files that belong to root or other users, you must have superuser privileges.

When you submit an at job, it is assigned a job identification number along with the .a extension. This designation becomes the job's file name, as well as its queue number.

Description of the at Command

Submitting an at job file involves these steps:

1. Invoking the at utility and specifying a command execution time
2. Typing a command or script to execute later

Note – If output from this command or script is important, be sure to direct the output to a file for later examination.

For example, the following `at` job removes `core` files from the user account `smith` near midnight on the last day of July.

```
$ at 11:45pm July 31
at> rm /home/smith/*core*
at> Press Control-d
commands will be executed using /bin/csh
job 933486300.a at Tue Jul 31 23:45:00 2004
```

Controlling Access to the `at` Command

You can set up a file to control access to the `at` command, permitting only specified users to create, remove, or display queue information about their `at` jobs. The file that controls access to the `at` command, `/etc/cron.d/at.deny`, consists of a list of user names, one user name per line. The users who are listed in this file cannot access `at` commands.

The `at.deny` file, which is created during SunOS software installation, contains the following user names:

```
daemon
bin
smtp
nuucp
listen
nobody
noaccess
```

With superuser privileges, you can edit the `at.deny` file to add other user names whose `at` command access you want to restrict.

▼ How to Create an `at` Job

Steps **1. Start the `at` utility, specifying the time you want your job executed.**

`$ at [-m] time [date]`

`-m` Sends you email after the job is completed.

time Specifies the hour that you want to schedule the job. Add `am` or `pm` if you do not specify the hours according to the 24-hour clock. Acceptable keywords are `midnight`, `noon`, and `now`. Minutes are optional.

date Specifies the first three or more letters of a month, a day of the week, or the keywords `today` or `tomorrow`.

2. **At the `at` prompt, type the commands or scripts that you want to execute, one per line.**

   You may type more than one command by pressing Return at the end of each line.

3. **Exit the `at` utility and save the `at` job by pressing Control-D.**

   Your `at` job is assigned a queue number, which is also the job's file name. This number is displayed when you exit the `at` utility.

Example 15–7 Creating an `at` Job

The following example shows the `at` job that user `jones` created to remove her backup files at 7:30 p.m. She used the `-m` option so that she would receive an email message after her job completed.

```
$ at -m 1930
at> rm /home/jones/*.backup
at> Press Control-D
job 897355800.a at Thu Jul  12 19:30:00 2004
```

She received a email message which confirmed the execution of her `at` job.

```
Your "at" job "rm /home/jones/*.backup"
completed.
```

The following example shows how `jones` scheduled a large `at` job for 4:00 a.m. Saturday morning. The job output was directed to a file named `big.file`.

```
$ at 4 am Saturday
at> sort -r /usr/dict/words > /export/home/jones/big.file
```

▼ How to Display the `at` Queue

Step ● **To check your jobs that are waiting in the `at` queue, use the `atq` command. This command displays status information about the `at` jobs that you have created.**

```
$ atq
```

▼ How to Verify an `at` Job

Step ● **To verify that you have created an `at` job, use the `atq` command. In the following example, the `atq` command confirms that `at` jobs that belong to `jones` have been submitted to the queue.**

```
$ atq
Rank      Execution Date     Owner     Job          Queue   Job Name
  1st   Jul 12, 2004 19:30   jones  897355800.a     a     stdin
  2nd   Jul 14, 2004 23:45   jones  897543900.a     a     stdin
  3rd   Jul 17, 2004 04:00   jones  897732000.a     a     stdin
```

▼ How to Display `at` Jobs

Step ● **To display information about the execution times of your `at` jobs, use the `at -l` command.**

```
$ at -l [job-id]
```

where the `-l` *job-id* option identifies the identification number of the job whose status you want to display.

Example 15–8 Displaying `at` Jobs

The following example shows output from the `at -l` command, which provides information on the status of all jobs submitted by a user.

```
$ at -l
897543900.a    Sat Jul 14 23:45:00 2004
897355800.a    Thu Jul 12 19:30:00 2004
897732000.a    Tue Jul 17 04:00:00 2004
```

The following example shows the output that is displayed when a single job is specified with the `at -l` command.

```
$ at -l 897732000.a
897732000.a       Tue Jul 17 04:00:00 2004
```

▼ How to Remove `at` Jobs

Before You Begin Become superuser or assume an equivalent role to remove an `at` job that belongs to `root` or another user. Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

You do not need to become superuser or assume an equivalent role to remove your own `at` job.

Steps

1. **Remove the `at` job from the queue before the job is executed.**

   ```
   $ at -r [job-id]
   ```

   where the `-r` *job-id* option specifies the identification number of the job you want to remove.

2. **Verify that the `at` job is removed by using the `at -l` (or the `atq`) command.**

   The `at -l` command displays the jobs remaining in the `at` queue. The job whose identification number you specified should not appear.

   ```
   $ at -l [job-id]
   ```

Example 15–9 Removing `at` Jobs

In the following example, a user wants to remove an `at` job that was scheduled to execute at 4 a.m. on July 17th. First, the user displays the `at` queue to locate the job identification number. Next, the user removes this job from the `at` queue. Finally, the user verifies that this job has been removed from the queue.

```
$ at -l
897543900.a    Sat Jul 14 23:45:00 2003
897355800.a    Thu Jul 12 19:30:00 2003
897732000.a    Tue Jul 17 04:00:00 2003
$ at -r 897732000.a
$ at -l 897732000.a
at: 858142000.a: No such file or directory
```

▼ How to Deny Access to the `at` Command

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Edit the `/etc/cron.d/at.deny` file and add the names of users, one user name per line, that will be prevented from using the `at` commands.**

   ```
   daemon
   bin
   smtp
   nuucp
   listen
   nobody
   ```

```
noaccess
username1
username2
username3
.
.
.
```

Example 15–10 Denying `at` Access

The following example shows an `at.deny` file that has been edited so that the users `smith` and `jones` cannot access the `at` command.

```
$ cat at.deny
daemon
bin
smtp
nuucp
listen
nobody
noaccess
jones
smith
```

▼ How to Verify That `at` Command Access Is Denied

Step ● **To verify that a username was added correctly to the `/etc/cron.d/at.deny` file, use the `at -l` command while logged in as the user. If the user `smith` cannot access the `at` command, the following message is displayed.**

```
# su smith
Password:
$ at -l
at: you are not authorized to use at.  Sorry.
```

Likewise, if the user tries to submit an `at` job, the following message is displayed:

```
$ at 2:30pm
at: you are not authorized to use at.  Sorry.
```

This message confirms that the user is listed in the `at.deny` file.

If `at` command access is allowed, then the `at -l` command returns nothing.

CHAPTER 16

Managing System Accounting (Tasks)

This chapter describes how to set up and maintain system accounting.

This is a list of the overview information in this chapter.

- "What is System Accounting?" on page 296
- "Setting Up System Accounting" on page 301

For information on using extended accounting, see Chapter 4, "Extended Accounting (Overview)," in *System Administration Guide: N1 Grid Containers, Resource Management, and Solaris Zones*.

For information on the step-by-step procedures associated with system accounting, see "System Accounting (Task Map)" on page 300

For reference information on the various system accounting reports, see Chapter 17.

What's New in System Accounting

This section describes new or modified system accounting features.

Solaris Process Accounting and Statistics Improvements

In this release, changes have been made to the internals of the load averaging, `cpu usr/sys/idle`, and accounting functions. Microstate accounting has replaced the old accounting mechanism and is enabled by default all of the time. As a result, you might notice slightly different process usage and timing statistics.

Switching to microstate accounting provides substantially more accurate data about user processes and the amount of time they spend in various states. In addition, this information is used to generate more accurate load averages and statistics from the `/proc` file system.

What is System Accounting?

The SunOS 5.10 system accounting software is a set of programs that enables you to collect and record data about user connect time, CPU time charged to processes, and disk usage. Once you collect this data, you can generate reports and charge fees for system usage.

You can use system accounting on a daily or monthly basis. Or, you can tack disk usage per user.

You can use the accounting programs to perform these tasks:

- Monitor system usage
- Locate and correct performance problems
- Maintain system security

After you set up the system accounting programs, they run mostly on their own.

How System Accounting Works

Automatic accounting is set up by first putting the accounting startup script into root's `crontab` file. The accounting startup script can then be started automatically by the `cron` command.

The following overview describes the system accounting process.

1. Between system startup and shutdown, raw data about system use (such as user logins, running processes, and data storage) are collected in accounting files.
2. Periodically (usually once a day), the `/usr/lib/acct/runacct` script processes the various accounting files and produces both cumulative summary files and daily accounting reports. Then, the `/usr/lib/acct/prdaily` script prints the daily reports.

   For more information about the `runacct` script, see "`runacct` Script" on page 311.
3. Monthly, you can process and print the cumulative `runacct` summary files by executing the `monacct` script. The summary reports produced by the `monacct` script provide an efficient means for billing users on a monthly or other fiscal basis.

System Accounting Components

The system accounting software provides C language programs and shell scripts that organize data into summary files and reports. These programs reside in the `/usr/lib/acct` directory. The accounting reports reside in the `/var/adm/acct` directory.

Daily accounting can help you perform four types of auditing:

- Connect accounting
- Process accounting
- Disk accounting
- Fee calculations

Connect Accounting

Connect accounting enables you to determine the following information:

- The length of time a user was logged in
- How the `tty` lines are being used
- The number of reboots on your system
- How many times the accounting software was turned off and on

To provide this information on connect sessions, the system stores the following data

- Record of time adjustments
- Boot times
- Number of times the accounting software was turned off and on
- Changes in run levels
- The creation of user processes (`login` processes and `init` processes)
- The terminations of processes

These records are produced from the output of system programs such as `date`, `init`, `login`, `ttymon`, and `acctwtmp`. They are stored in the `/var/adm/wtmpx` file. Entries in the `wtmpx` file can contain the following information:

- Login name
- Device name
- Process ID
- Entry type
- Time stamp that denotes when the entry was made

Process Accounting

Process accounting enables you to keep track of the following data about each process that runs on your system:

- User IDs and group IDs of users using the process

- Beginning times and elapsed times of the process
- CPU time for the process (user time and system time)
- Amount of memory used by the process
- Commands run by the process
- The tty that controls the process

Every time a process terminates, the exit program collects this information and writes it to the /var/adm/pacct file.

Disk Accounting

Disk accounting enables you to gather and format the following data about the files each user has on disks:

- User name and user ID of the user
- Number of blocks that are used by the user's files

This data is collected by the /usr/lib/acct/dodisk shell script at intervals that are determined by the entry you add to the /var/spool/cron/crontabs/root file. In turn, the dodisk script invokes the acctdisk and acctdusg commands. These commands gather disk usage by login name.

Caution – Information gathered by running the dodisk script is stored in the /var/adm/acct/nite/disktacct file. This information is overwritten the next time the dodisk script is run. Therefore, avoid running the dodisk script twice in the same day.

The acctdusg command might overcharge for files that are written randomly, which can create holes in the files. This problem occurs because the acctdusg command does not read the indirect blocks of a file when determining the file size. Rather, the acctdusg command determines the file size by checking the current file size value in the file's inode.

Fee Calculations

The chargefee utility stores charges for special services that are provided to a user in the /var/adm/fee file. A special service, for example, is file restoration. Each entry in the file consists of a user login name, user ID, and the fee. This file is checked by the runacct script every day, and new entries are merged into the accounting records. For instructions on running the chargefee script to bill users, see "How to Bill Users" on page 304.

How Daily Accounting Works

Here is a step-by-step summary of how daily accounting works:

1. When the system is switched into multiuser mode, the /usr/lib/acct/startup program is executed. The startup program executes several other programs that invoke daily accounting.
2. The acctwtmp program adds a "boot" record to the /var/adm/wtmpx file. In this record, the system name is shown as the user name in the wtmpx record. The following table summarizes how the raw accounting data is gathered and where it is stored.

File in /var/adm	Information Stored	Written By	Format
wtmpx	Connect sessions	login, init	Binary
	Changes	date	Binary
	Reboots	acctwtmp	Binary
	Shutdowns	shutacct	Binary
pacctn	Processes	Kernel (when the process ends)	Binary
		turnacct switch (which creates a new file when the old file reaches 500 blocks)	Binary
fee	Special charges	chargefee	ASCII
acct/nite/disktacct	Disk space used	dodisk	Binary

3. The turnacct script, invoked with the -on option, begins process accounting. Specifically, the turnacct script executes the accton program with the /var/adm/pacct argument.
4. The remove shell script "cleans up" the saved pacct and wtmpx files that are left in the sum directory by the runacct script.
5. The login and init programs record connect sessions by writing records into the /var/adm/wtmpx file. Date changes (using date with an argument) are also written to the /var/adm/wtmpx file. Reboots and shutdowns using the acctwtmp command are also recorded in the /var/adm/wtmpx file.
6. When a process ends, the kernel writes one record per process, using the acct.h format, in the /var/adm/pacct file.

   Every hour, the cron command executes the ckpacct script to check the size of the /var/adm/pacct file. If the file grows beyond 500 blocks (default), the turnacct switch command is executed. (The program moves the pacct file to the pacct*n* file and creates a new file.) The advantage of having several smaller pacct files becomes apparent when you try to restart the runacct script if a failure occurs when processing these records.

7. The `runacct` script is executed by the `cron` command each night. The `runacct` script processes the accounting files to produce command summaries and usage summaries by user name. These accounting files are processed: `/var/adm/pacct`*n*, `/var/adm/wtmpx`, `/var/adm/fee`, and `/var/adm/acct/nite/disktacct`.
8. The `/usr/lib/acct/prdaily` script is executed on a daily basis by the `runacct` script to write the daily accounting information in the `/var/adm/acct/sum/rprt`*MMDD* files.
9. The `monacct` script should be executed on a monthly basis (or at intervals you determine, such as at the end of every fiscal period). The `monacct` script creates a report that is based on data stored in the `sum` directory that has been updated daily by the `runacct` script. After creating the report, the `monacct` script "cleans up" the `sum` directory to prepare the directory's files for the new `runacct` data.

What Happens if the System Shuts Down

If the system is shut down by using the `shutdown` command, the `shutacct` script is executed automatically. The `shutacct` script writes a *reason record* into the `/var/adm/wtmpx` file and turns off process accounting.

System Accounting (Task Map)

Task	Description	For Instructions
Set up system accounting.	Set up system accounting by performing the following tasks: ■ Create the `/etc/rc0.d/K22acct` and `/etc/rc2.d/S22acct` files. ■ Modify the `/var/spool/cron/crontabs/adm` and `/var/spool/cron/crontabs/root` `crontab` files.	"How to Set Up System Accounting" on page 302
Bill users.	Run the `/usr/lib/acct/chargefee` *username amount* command.	"How to Bill Users" on page 304
Fix a corrupted `wtmpx` file.	Convert the `wtmpx` file from binary to ASCII format.	"How to Fix a Corrupted `wtmpx` File" on page 305

Task	Description	For Instructions
Fix tacct errors.	Run the prtacct script to check the /var/adm/acct/sum/tacctprev file. Then, patch the latest/var/adm/acct/sum/tacct*MMDD* file. You will need to re-create the /var/adm/acct/sum/tacct file.	"How to Fix tacct Errors" on page 306
Restart the runacct script.	Remove the lastdate file and any lock files. Then, manually restart the runacct script.	"How to Restart the runacct Script" on page 307
Disable system accounting temporarily.	Edit theadm crontab file to stop the ckpacct, runacct, and monacct programs from running.	"How to Temporarily Stop System Accounting" on page 308
Disable system accounting permanently.	Delete the entries for the ckpacct, runacct, and monacct programs in the adm and crontab files.	"How to Permanently Disable System Accounting" on page 308

Setting Up System Accounting

You can set up system accounting to run while the system is in multiuser mode (Run Level 2). Generally, this task involves these steps:

1. Creating the /etc/rc0.d/K22acct and /etc/rc2.d/S22acct startup scripts
2. Modifying the /var/spool/cron/crontabs/adm and /var/spool/cron/crontabs/root crontab files

The following table describes the default accounting scripts.

TABLE 16–1 Default Accounting Scripts

Purpose	Accounting Script	Man Page	Run Frequency
Checks the size of the /usr/adm/pacct log file and makes sure that it does not get too large.	ckpacct	acctsh(1M)	Periodically

TABLE 16–1 Default Accounting Scripts *(Continued)*

Purpose	Accounting Script	Man Page	Run Frequency
Processes connect, disk, and fee accounting information. You can remove from this script the commands for the accounting features you do not want processed.	`runacct`	`runacct`(1M)	Daily
Generates fiscal accounting summary reports on a monthly basis. You can determine how often this script is run. You can remove from this script the commands for the accounting features you do not want to use.	`monacct`	`acctsh`(1M)	On a fiscal basis

You can choose which accounting scripts run by default. After these entries have been added to the `crontab` files, system accounting should run automatically.

▼ How to Set Up System Accounting

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **If necessary, install the `SUNWaccr` and `SUNWaccu` packages on your system by using the `pkgadd` command.**

3. **Install `/etc/init.d/acct` as the startup script for Run Level 2.**

   ```
   # ln /etc/init.d/acct /etc/rc2.d/S22acct
   ```

4. **Install `/etc/init.d/acct` as the stop script for Run Level 0.**

   ```
   # ln /etc/init.d/acct /etc/rc0.d/K22acct
   ```

5. **Add the following lines to the `adm crontab` file to start the `ckpacct`, `runacct`, and `monacct` scripts automatically.**

   ```
   # EDITOR=vi; export EDITOR
   # crontab -e adm
   ```

```
0 * * * * /usr/lib/acct/ckpacct
30 2 * * * /usr/lib/acct/runacct 2> /var/adm/acct/nite/fd2log
30 7 1 * * /usr/lib/acct/monacct
```

6. **Add the following line to the `root crontab` file to start the `dodisk` script automatically.**

   ```
   # crontab -e
   30 22 * * 4 /usr/lib/acct/dodisk
   ```

7. **Edit `/etc/acct/holidays` to include national holidays and local holidays.**

   For more information, see the `holidays`(4) man page and the example that follows.

8. **Reboot the system, or start system accounting manually by typing:**

   ```
   # /etc/init.d/acct start
   ```

Example 16–1 Setting Up Accounting (`adm crontab`)

This modified `adm crontab` contains entries for the `ckpacct`, `runacct`, and `monacct` scripts.

```
#ident  "@(#)adm        1.5     92/07/14 SMI"    /* SVr4.0 1.2   */
#
# The adm crontab file should contain startup of performance
# collection if the profiling and performance feature has been
# installed.
0 * * * * /usr/lib/acct/ckpacct
30 2 * * * /usr/lib/acct/runacct 2> /var/adm/acct/nite/fd2log
30 7 1 * * /usr/lib/acct/monacct
```

Example 16–2 Setting Up Accounting (`root crontab`)

This modified `root crontab` contains entries for the `dodisk` program.

```
#ident  "@(#)root       1.19    98/07/06 SMI"   /* SVr4.0 1.1.3.1       */
#
# The root crontab should be used to perform accounting data collection.
#
#
10 3 * * * /usr/sbin/logadm
15 3 * * 0 /usr/lib/fs/nfs/nfsfind
30 3 * * * [ -x /usr/lib/gss/gsscred_clean ] && /usr/lib/gss/gsscred_clean
30 22 * * 4 /usr/lib/acct/dodisk
```

Example 16–3 Setting Up Accounting (`/etc/acct/holidays`)

The following example shows a sample `/etc/acct/holidays` file.

```
* @(#)holidays    January 1, 2004
*
* Prime/Nonprime Table for UNIX Accounting System
```

```
*
* Curr    Prime    Non-Prime
* Year    Start    Start
*
  2004    0800    1800
*
* only the first column (month/day) is significant.
*
* month/day    Company
*         Holiday
*
1/1        New Years Day
7/4        Indep. Day
12/25        Christmas
```

Billing Users

If you provide special user services by request. Special services include restoring files or remote printing. You might want to bill users by running the `chargefee` utility. The `chargefee` utility records charges in the `/var/adm/fee` file. Each time the `runacct` utility is executed, new entries are merged into the total accounting records.

See the `acctsh`(1M) man page for more information.

▼ How to Bill Users

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Charge a user for special services.**

   ```
   # /usr/lib/acct/chargefee username amount
   ```

   username Is the user account you want to bill.

   amount Specifies the number of units to bill the user. This value is an arbitrary unit that you set to charge users based on some task such as printing or restoring a file. You would have to write a script that invokes the `chargefee` utility and charges a user for a specific task.

Example 16–4 Billing Users

In the following example, the user `print_customer` is charged 10 units.

```
# /usr/lib/acct/chargefee print_customer 10
```

Maintaining Accounting Information

This section describes how to fix corrupted system accounting files and how to restart the `runacct` script.

Fixing Corrupted Files and `wtmpx` Errors

Unfortunately, system accounting is not foolproof. Occasionally, a file becomes corrupted or lost. Some files can simply be ignored or restored from backup. However, certain files must be fixed to maintain the integrity of system accounting.

The `wtmpx` files seem to cause the most problems in the daily operation of system accounting. When the date is changed manually and the system is in multiuser mode, a set of date change records is written to the `/var/adm/wtmpx` file. The `wtmpfix` utility is designed to adjust the time stamps in the `wtmp` records when a date change is encountered. However, some combinations of date changes and reboots slip through the `wtmpfix` utility and cause the `acctcon` program to fail.

▼ How to Fix a Corrupted `wtmpx` File

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Change to the `/var/adm` directory.**

3. **Convert the `wtmpx` file from binary format to ASCII format.**

   ```
   # /usr/lib/acct/fwtmp < wtmpx > wtmpx.ascii
   ```

4. **Edit the `wtmpx.ascii` file to delete the corrupted records.**

5. **Convert the `wtmpx.ascii` file back to a binary file.**

```
# /usr/lib/acct/fwtmp -ic < wtmpx.ascii > wtmpx
```

See the fwtmp(1M) man page for more information.

Fixing tacct Errors

The integrity of the /var/adm/acct/sum/tacct file is important if you are charging users for system resources. Occasionally, unusual tacct records appear with negative numbers, duplicate user IDs, or a user ID of 65535. First, check the /var/adm/acct/sum/tacctprev file by using the prtacct script to print the file. If the contents look all right, patch the latest /var/adm/acct/sum/tacct*MMDD* file. Then, re-create the /var/adm/acct/sum/tacct file. The following steps outline a simple patch procedure.

▼ How to Fix tacct Errors

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Change to the /var/adm/acct/sum directory.**

3. **Convert the tacct*MMDD* file from binary format to ASCII format.**

   ```
   # /usr/lib/acct/acctmerg -v < tacctMMDD > xtacct
   ```

   MMDD is pair of two-digit numbers that represent the month and day.

4. **Edit the xtacct file, removing corrupted records and writing duplicate records to another file.**

5. **Convert the xtacct file from ASCII format to binary format.**

   ```
   # /usr/lib/acct/acctmerg -i < xtacct > tacctMMDD
   ```

6. **Merge the files tacctprev and tacct.*MMDD* into the tacct file.**

   ```
   # /usr/lib/acct/acctmerg < tacctprev tacctMMDD > tacct
   ```

Restarting the runacct Script

The runacct script can fail for several reasons. The following are the most common reasons:

- A system crash
- The /var directory is running out of space
- A corrupted wtmpx file

If the active.*MMDD* file exists, check it first for error messages. If the active and lock files exist, check the fd2log file for any relevant messages.

Run without arguments, the runacct script assumes that this invocation is the first invocation of the day. The argument *MMDD* is necessary if the runacct script is being restarted and specifies the month and day for which the runacct script reruns the accounting. The entry point for processing is based on the contents of the statefile file. To override the statefile file, include the desired state on the command line. For a description of the available states, see the runacct(1M) man page.

Caution – When you run the runacct program manually, be sure to run it as user adm.

▼ How to Restart the runacct Script

Steps

1. **Change directories to the /var/adm/acct/nite directory.**

```
$ cd /var/adm/acct/nite
```

2. **Remove the lastdate file and any lock* files, if any.**

```
$ rm lastdate lock*
```

The lastdate file contains the date that the runacct program was last run. Restarting the runacct script in the next step re-creates this file.

3. **Restart the runacct script.**

```
$ /usr/lib/acct/runacct MMDD [state] 2> /var/adm/acct/nite/fd2log &
```

MMDD Is the month and day specified by two-digit numbers.

state Specifies a state, or starting point, where the runacct script processing should begin.

Stopping and Disabling System Accounting

You can temporarily stop system accounting or permanently disable it.

▼ How to Temporarily Stop System Accounting

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Edit the `adm` `crontab` file to stop the `ckpacct`, `runacct`, and `monacct` programs from running by commenting out the appropriate lines.**

   ```
   # EDITOR=vi; export EDITOR
   # crontab -e adm
   #0 * * * * /usr/lib/acct/ckpacct
   #30 2 * * * /usr/lib/acct/runacct 2> /var/adm/acct/nite/fd2log
   #30 7 1 * * /usr/lib/acct/monacct
   ```

3. **Edit the `root` `crontab` file to stop the `dodisk` program from running by commenting out the appropriate line.**

   ```
   # crontab -e
   #30 22 * * 4 /usr/lib/acct/dodisk
   ```

4. **Stop the system accounting program.**

   ```
   # /etc/init.d/acct stop
   ```

5. **(Optional) Remove the newly added comment symbols from the `crontab` files.**

6. **Restart the system accounting program to re-enable system accounting.**

   ```
   # /etc/init.d/acct start
   ```

▼ How to Permanently Disable System Accounting

Steps

1. **Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Edit the `adm` `crontab` file and delete the entries for the `ckpacct`, `runacct`, and `monacct` programs.**

```
# EDITOR=vi; export EDITOR
# crontab -e adm
```

3. **Edit the `root` `crontab` file and delete the entries for the `dodisk` program.**

```
# crontab -e
```

4. **Remove the startup script for Run Level 2.**

```
# unlink /etc/rc2.d/S22acct
```

5. **Remove the stop script for Run Level 0.**

```
# unlink /etc/rc0.d/K22acct
```

6. **Stop the system accounting program.**

```
# /etc/init.d/acct stop
```

CHAPTER 17

System Accounting (Reference)

This chapter provides reference information about system accounting.

This is a list of the reference information in this chapter.

For more information about system accounting tasks, see Chapter 16.

runacct Script

The main daily accounting script, `runacct`, is normally invoked by the `cron` command outside of normal business hours. The `runacct` script processes connect, fee, disk, and process accounting files. This script also prepares daily and cumulative summary files for use by the `prdaily` and `monacct` scripts for billing purposes.

The `runacct` script takes care not to damage files if errors occur. A series of protection mechanisms that are used to perform the following tasks:

- Recognize an error
- Provide intelligent diagnostics
- Complete processing in such a way that the `runacct` script can be restarted with minimal intervention

This script records its progress by writing descriptive messages to the `active` file. Files used by the `runacct` script are assumed to be in the `/var/adm/acct/nite` directory, unless otherwise noted. All diagnostic output during the execution of the `runacct` script is written to the `fd2log` file.

When the runacct script is invoked, it creates the lock and lock1 files. These files are used to prevent simultaneous execution of the runacct script. The runacct program prints an error message if these files exist when it is invoked. The lastdate file contains the month and day the runacct script was last invoked, and is used to prevent more than one execution per day. If the runacct script detects an error, the following occurs:

- A message is written to the console
- Email is sent to root and adm
- Locks might be removed
- Diagnostics are saved
- Execution is ended

For instructions on how to restart the runacct script, see "How to Restart the runacct Script" on page 307.

To allow the runacct script to be restarted, processing is broken down into separate re-entrant states. The statefile file is used to track the last state completed. When each state is completed, the statefile file is updated to reflect the next state. After processing for the state is complete, the statefile file is read and the next state is processed. When the runacct script reaches the CLEANUP state, it removes the locks and ends. States are executed as shown in the following table.

TABLE 17–1 States of the runacct Script

State	Description
SETUP	The turnacct switch command is executed to create a new pacct file. The /var/adm/pacct*n* process accounting files (except for the pacct file) are moved to the /var/adm/Spacct*n*.*MMDD* files. The /var/adm/wtmpx file is moved to the /var/adm/acct/nite/wtmp.*MMDD* file (with the current time record added on the end) and a new /var/adm/wtmp file is created. The closewtmp and utmp2wtmp programs add records to the wtmp.*MMDD* file and the new wtmpx file to account for users who are currently logged in.
WTMPFIX	The wtmpfix program checks the wtmp.*MMDD* file in the nite directory for accuracy. Because some date changes cause the acctcon program to fail, the wtmpfix program attempts to adjust the time stamps in the wtmpx file if a record of a date change appears. This program also deletes any corrupted entries from the wtmpx file. The fixed version of the wtmp.*MMDD* file is written to the tmpwtmp file.
CONNECT	The acctcon program is used to record connect accounting records in the file ctacct.*MMDD*. These records are in tacct.h format. In addition, the acctcon program creates the lineuse and reboots files. The reboots file records all the boot records found in the wtmpx file.

TABLE 17–1 States of the runacct Script *(Continued)*

State	Description
PROCESS	The acctprc program is used to convert the /var/adm/Spacct*n.MMDD* process accounting files into complete accounting records in the ptacct*n.MMDD* files. The Spacct and ptacct files are correlated by number so that if the runacct script fails, the Spacct files are not processed.
MERGE	The acctmerg program merges the process accounting records with the connect accounting records to form the daytacct file.
FEES	The acctmerg program merges ASCII tacct records from the fee file into the daytacct file.
DISK	The dodisk script produces the disktacct file. If the dodisk script has been run, which produces the disktacct file, the DISK program merges the file into the daytacct file and moves the disktacct file to the /tmp/disktacct.*MMDD* file.
MERGETACCT	The acctmerg program merges the daytacct file with the sum/tacct file, the cumulative total accounting file. Each day, the daytacct file is saved in the sum/tacct.*MMDD* file so that the sum/tacct file can be re-created if it is corrupted or lost.
CMS	The acctcms program is run several times. This program is first run to generate the command summary by using the Spacct*n* files and write the data to the sum/daycms file. The acctcms program is then run to merge the sum/daycms file with the sum/cms cumulative command summary file. Finally, the acctcms program is run to produce nite/daycms and nite/cms, the ASCII command summary files from the sum/daycms and sum/cms files, respectively. The lastlogin program is used to create the /var/adm/acct/sum/loginlog log file. This file reports when each user last logged in. If the runacct script is run after midnight, the dates showing the time last logged in by some users will be incorrect by one day.
USEREXIT	Any installation-dependent (local) accounting program can be run at this point. The runacct script expects this program to be called the /usr/lib/acct/runacct.local program.
CLEANUP	This state cleans up temporary files, runs the prdaily script and saves its output in the sum/rpt.*MMDD* file, removes the locks, and then exits.

Caution – When restarting the runacct script in the CLEANUP state, remove the last ptacct file because this file will not be complete.

Daily Accounting Reports

The `runacct` shell script generates five basic reports upon each invocation. The following table describes these reports.

TABLE 17–2 Daily Accounting Reports

Report Type	Description
"Daily Report" on page 314	Shows terminal line utilization by `tty` number.
"Daily Usage Report" on page 315	Indicates usage of system resources by users (listed in order of user ID).
"Daily Command Summary" on page 316	Indicates usage of system resources by commands, listed in descending order of memory use. In other words, the command that used the most memory is listed first. This same information is reported for the month in the monthly command summary.
"Monthly Command Summary" on page 318	A cumulative summary that reflects the data accumulated since the last invocation of the `monacct` program.
"Last Login Report" on page 318	Shows the last time each user logged in (listed in chronological order).

Daily Report

This report gives information about each terminal line used. The following is a sample Daily Report.

```
Jan 16 02:30 2004  DAILY REPORT FOR venus Page 1


from Mon Jan 15 02:30:02 2004
to   Tue Oan 16 02:30:01 2004
1       runacct
1       acctcon

TOTAL DURATION IS 1440 MINUTES
LINE         MINUTES  PERCENT  # SESS   # ON   # OFF
console      868      60       1        1      2
TOTALS       868      --       1        1      2
```

The `from` and `to` lines specify the time period reflected in the report. This time period covers the time the last Daily Report was generated to the time the current Daily Report was generated. Then, the report presents a log of system reboots, shutdowns, power failure recoveries, and any other record written to the `/var/adm/wtmpx` file by the `acctwtmp` program. For more information, see the `acct`(1M) man page.

The second part of the report is a breakdown of terminal line utilization. The TOTAL DURATION tells how long the system was in multiuser mode (accessible through the terminal lines). The following table describes the data provided by the Daily Report.

TABLE 17–3 Daily Report Data

Column	Description
LINE	The terminal line or access port.
MINUTES	The number of minutes that the line was in use during the accounting period.
PERCENT	The TOTAL DURATION divided by the number of MINUTES.
# SESS	The number of times this line or port was accessed for a login session.
# ON	Same as SESS. (This column no longer has meaning. Previously, this column listed the number of times that a line or port was used to log in a user.)
# OFF T	The number of times a user logs out and any interrupts that occur on that line. Generally, interrupts occur on a port when ttymon is first invoked after the system is brought to multiuser mode. If the # OFF exceeds the # SESS by a large factor, the multiplexer, modem, or cable is probably going bad. Or, a bad connection exists somewhere. The most common cause is an unconnected cable dangling from the multiplexer.

During real time, you should monitor the /var/adm/wtmpx file because it is the file from which the connect accounting is derived. If the wtmpx file grows rapidly, execute the following command to see which tty line is the noisiest.

```
# /usr/lib/acct/acctcon -l file < /var/adm/wtmpx
```

If interruption is occurring frequently, general system performance will be affected. Additionally, the wtmp file might become corrupted. To correct this problem, see "How to Fix a Corrupted wtmpx File" on page 305.

Daily Usage Report

The Daily Usage Report breaks down system resource utilization by user. A sample of this report follows.

```
Jan 16 02:30 2004  DAILY USAGE REPORT FOR skisun Page 1


     LOGIN   CPU   (MINS)   KCORE-    MINS     CONNECT   (MINS) DISK   # OF   # OF   # DISK   FEE
UID  NAME    PRIME NPRIME PRIME     NPRIME   PRIME     NPRIME BLOCKS PROCS  SESS   SAMPLES
0    TOTAL   72     148     11006173 51168    26230634 57792  539    330    0      2150     1
0    root    32     76      11006164 33664    26230616 22784  0      0      0      127      0
```

```
4     adm     0      0       22       51       0        0      0       420     0      0        0
101   rimmer 39     72      894385   1766020 539       330    0       1603    1      0        0
```

The following table describes the data provided by the Daily Usage Report.

TABLE 17–4 Daily Usage Report Data

Column	Description
`UID`	User ID number.
`LOGIN NAME`	Login (or user) name of the user. Identifies a user who has multiple login names.
`CPU (MINS)`	Amount of time, in minutes, that the user's process used the central processing unit. Divided into `PRIME` and `NPRIME` (nonprime) utilization. The accounting system's version of this data is located in the `/etc/acct/holidays` file.
`KCORE-MINS`	A cumulative measure of the amount of memory in Kbyte segments per minute that a process uses while running. Divided into `PRIME` and `NPRIME` utilization.
`CONNECT (MINS)`	Amount of time, in minutes, that the a user was logged in to the system, or "real time." Divided into `PRIME` and `NPRIME` utilization. If these numbers are high while the `# OF PROCS` is low, you can conclude that the user logs in first thing in the morning and hardly touches the terminal the rest of the day.
`DISK BLOCKS`	Output from the `acctdusg` program, which runs the disk accounting programs and merges the accounting records (`daytacct`). For accounting purposes, a block is 512 bytes.
`# OF PROCS`	Number of processes invoked by the user. If large numbers appear, a user might have a shell procedure that has run out of control.
`# OF SESS`	Number of times that a user logged in to the system.
`# DISK SAMPLES`	Number of times that disk accounting was run to obtain the average number of `DISK BLOCKS`.
`FEE`	Often unused field that represents the total accumulation of units charged against the user by the `chargefee` script.

Daily Command Summary

The Daily Command Summary report shows the system resource utilization by command. With this report, you can identify the most heavily used commands. Based on how those commands use system resources, you can then gain insight on how best to tune the system.

These reports are sorted by `TOTAL KCOREMIN`, which is an arbitrary gauge but often useful for calculating drain on a system.

A sample Daily Command Summary follows.

```
                                TOTAL COMMAND SUMMARY
COMMAND    NUMBER      TOTAL    TOTAL      TOTAL    MEAN     MEAN     HOG   CHARS    BLOCKS
NAME         CMDS    KCOREMIN CPU-MIN REAL-MIN  SIZE-K  CPU-MIN  FACTOR TRNSFD   READ

TOTALS       2150  1334999.75  219.59 724258.50 6079.48   0.10   0.00   397338982 419448

netscape       43  2456898.50   92.03  54503.12 26695.51  2.14   0.00   947774912 225568
adeptedi        7    88328.22    4.03    404.12 21914.95  0.58   0.01    93155160   8774
dtmail          1    54919.17    5.33  17716.57 10308.94  5.33   0.00   213843968  40192
acroread        8    31218.02    2.67  17744.57 11682.66  0.33   0.00   331454464  11260
dtwm            1    16252.93    2.53  17716.57 6416.05   2.53   0.00   158662656  12848
dtterm          5     4762.71    1.30  76300.29 3658.93   0.26   0.00    33828352  11604
dtaction       23     1389.72    0.33       0.60 4196.43  0.01   0.55    18653184    539
dtsessio        1     1174.87    0.24  17716.57 4932.97   0.24   0.00    23535616   5421
dtcm            1      866.30    0.18  17716.57 4826.21   0.18   0.00     3012096   6490
```

The following table describes the data provided by the Daily Command Summary.

TABLE 17–5 Daily Command Summary Data

Column	Description
COMMAND NAME	Name of the command. All shell procedures are lumped together under the name sh because only object modules are reported by the process accounting system. You should monitor the frequency of programs called a.out or core, or any other unexpected name. You can use the acctcom program to determine who executed an oddly named command and if superuser privileges were used.
NUMBER CMDS	Total number of times this command was run during prime time.
TOTAL KCOREMIN	Total cumulative measurement of the Kbyte segments of memory used by a process per minute of run time.
TOTAL CPU-MIN	Total processing time this program accumulated during prime time.
TOTAL REAL-MIN	Total real-time (wall-clock) minutes this program accumulated.
MEAN SIZE-K	Mean (average) of the TOTAL KCOREMIN over the number of invocations reflected by the NUMBER CMDS.
MEAN CPU-MIN	Mean (average) derived from the NUMBER CMDS and the TOTAL CPU-MIN.
HOG FACTOR	Total CPU time divided by elapsed time. Shows the ratio of system availability to system utilization, providing a relative measure of total available CPU time consumed by the process during its execution.

TABLE 17–5 Daily Command Summary Data *(Continued)*

Column	Description
`CHARS TRNSFD`	Total number of characters transferred by the read and write system calls. Might be negative due to overflow.
`BLOCKS READ`	Total number of the physical block reads and writes that a process performed.

Monthly Command Summary

The format of the Daily Command Summary and the Monthly Command Summary reports are virtually the same. However, the daily summary reports only on the current accounting period while the monthly summary reports on the start of the fiscal period to the current date. In other words, the monthly report is a cumulative summary that reflects the data accumulated since the last invocation of the `monacct` program.

A sample Monthly Command Summary follows.

```
Jan 16 02:30 2004  MONTHLY TOTAL COMMAND SUMMARY Page 1


                                     TOTAL COMMAND SUMMARY
COMMAND    NUMBER      TOTAL   TOTAL     TOTAL   MEAN      MEAN    HOG      CHARS     BLOCKS
NAME         CMDS   KCOREMIN CPU-MIN  REAL-MIN  SIZE-K    CPU-MIN FACTOR  TRNSFD     READ

TOTALS      42718  4398793.50  361.92  956039.00 12154.09 0.01     0.00  16100942848 825171

netscape      789  3110437.25  121.03   79101.12 25699.58 0.15     0.00   3930527232 302486
adeptedi       84  1214419.00   50.20    4174.65 24193.62 0.60     0.01    890216640 107237
acroread      145   165297.78    7.01   18180.74 23566.84 0.05     0.00   1900504064  26053
dtmail          2    64208.90    6.35   20557.14 10112.43 3.17     0.00    250445824  43280
dtaction      800    47602.28   11.26      15.37  4226.93 0.01     0.73    640057536   8095
soffice.       13    35506.79    0.97       9.23 36510.84 0.07     0.11    134754320   5712
dtwm            2    20350.98    3.17   20557.14  6419.87 1.59     0.00    190636032  14049
```

For a description of the data provided by the Monthly Command Summary, see "Daily Command Summary" on page 316.

Last Login Report

This report gives the date when a particular login was last used. You can use this information to find unused logins and login directories that can be archived and deleted. A Last Login Report follows.

```
Jan 16 02:30 2004  LAST LOGIN Page 1
```

```
01-06-12  kryten        01-09-08  protoA      01-10-14  ripley
01-07-14  lister        01-09-08  protoB      01-10-15  scutter1
01-08-16  pmorph        01-10-12  rimmer      01-10-16  scutter2
```

Examining the pacct File With acctcom

At any time, you can examine the contents of the /var/adm/pacct*n* files, or any file with records in the acct.h format, by using the acctcom program. If you do not specify any files and do not provide any standard input when you run this command, the acctcom command reads the pacct file. Each record read by the acctcom command represents information about a terminated process. Active processes can be examined by running the ps command.

The default output of the acctcom command provides the following information:

```
# acctcom
COMMAND                            START    END          REAL     CPU    MEAN
NAME      USER     TTYNAME         TIME     TIME        (SECS)  (SECS) SIZE(K)
#accton   root      ?              02:30:01 02:30:01     0.03     0.01  304.00
turnacct  adm       ?              02:30:01 02:30:01     0.42     0.01  320.00
mv        adm       ?              02:30:01 02:30:01     0.07     0.01  504.00
utmp_upd  adm       ?              02:30:01 02:30:01     0.03     0.01  712.00
utmp_upd  adm       ?              02:30:01 02:30:01     0.01     0.01  824.00
utmp_upd  adm       ?              02:30:01 02:30:01     0.01     0.01  912.00
utmp_upd  adm       ?              02:30:01 02:30:01     0.01     0.01  920.00
utmp_upd  adm       ?              02:30:01 02:30:01     0.01     0.01 1136.00
utmp_upd  adm       ?              02:30:01 02:30:01     0.01     0.01  576.00
closewtm  adm       ?              02:30:01 02:30:01     0.10     0.01  664.00
```

Field	Explanation
COMMAND NAME	Command name (pound (#) sign if the command was executed with superuser privileges)
USER	User name
TTYNAME	tty name (listed as ? if unknown)
START TIME	Command execution starting time
END TIME	Command execution ending time
REAL (SECS)	Real time (in seconds)
CPU (SECS)	CPU time (in seconds)
MEAN SIZE (K)	Mean size (in Kbytes)

You can obtain the following information by using acctcom command options:

- State of the fork/exec flag (1 for fork without exec)

- System exit status
- Hog factor
- Total kcore minutes
- CPU factor
- Characters transferred
- Blocks read

The following table describes the acctcom command options.

TABLE 17–6 Options for the acctcom Command

Option	Description
-a	Shows average statistics about the processes selected. The statistics are printed after the output is recorded.
-b	Reads the files backward, showing latest commands first. This option has no effect if reading standard input.
-f	Prints the fork/exec flag and system exit status columns. The output is an octal number.
-h	Instead of mean memory size, shows the hog factor, which is the fraction of total available CPU time consumed by the process during its execution. Hog factor = *total-CPU-time/elapsed-time*.
-i	Prints columns that contains the I/O counts in the output.
-k	Shows total kcore minutes instead of memory size.
-m	Shows mean core size. This size is the default.
-q	Prints average statistics, not output records.
-r	Shows CPU factor: *user-time*/(*system-time* + *user-time)*.
-t	Shows separate system and user CPU times.
-v	Excludes column headings from the output.
-C *sec*	Shows only processes with total CPU time (system plus user) that exceeds *sec* seconds.
-e *time*	Shows processes existing at or before *time*, given in the format *hr[:min[:sec]]*
-E *time*	Shows processes starting at or before *time*, given in the format *hr[:min[:sec]]*. Using the same time for both -S and -E, shows processes that existed at the time.
-g *group*	Shows only processes that belong to *group*.
-H *factor*	Shows only processes that exceed *factor*, where *factor* is the "hog factor" (see the -h option).

TABLE 17–6 Options for the `acctcom` Command *(Continued)*

Option	Description
`-I` *chars*	Shows only processes that transferred more characters than the cutoff number specified by *chars*.
`-l` *line*	Show only processes that belong to the terminal `/dev/`*line*.
`-n` *pattern*	Shows only commands that match *pattern* (a regular expression except that "+" means one or more occurrences).
`-o` *ofile*	Instead of printing the records, copies them in `acct.h` format to *ofile*.
`-O` *sec*	Shows only processes with CPU system time that exceeds *sec* seconds.
`-s` *time*	Show processes existing at or after *time*, given in the format *hr*[:*min*[:*sec*]].
`-S` *time*	Show processes starting at or after *time*, given in the format *hr*[:*min*[:*sec*]].
`-u` *user*	Shows only processes that belong to *user*.

System Accounting Files

The `/var/adm` directory contains the active data collection files. The following table describes the accounting files in this directory.

TABLE 17–7 Files in the `/var/adm` Directory

File	Description
`dtmp`	Output from the `acctdusg` program
`fee`	Output from the `chargefee` program, which are the ASCII `tacct` records
`pacct`	Active process accounting file
`pacct`*n*	Process accounting files that are switched by running the `turnacct` script
`Spacct`*n.MMDD*	Process accounting files for *MMDD* during execution of the `runacct` script

The `/var/adm/acct` directory contains the `nite`, `sum`, and `fiscal` directories. These directories contain the actual data collection files. For example, the `nite` directory contains files that are reused daily by the `runacct` script. A brief summary of the files in the `/var/adm/acct/nite` directory follows.

TABLE 17–8 Files in the /var/adm/acct/nite Directory

File	Description
active	Used by the runacct script to record progress and print warning and error messages
active.*MMDD*	Same as the active file after the runacct script detects an error
cms	ASCII total command summary used by the prdaily script
ctacct.*MMDD*	Connect accounting records in tacct.h format
ctmp	Output of acctcon1 program, which consists of connect session records in ctmp.h format (acctcon1 and acctcon2 are provided for compatibility purposes)
daycms	ASCII daily command summary used by the prdaily script
daytacct	Total accounting records for one day in tacct.h format
disktacct	Disk accounting records in tacct.h format, created by the dodisk script
fd2log	Diagnostic output during execution of the runacct script
lastdate	Last day the runacct script executed (in date +%m%d format)
lineuse	tty line usage report used by the prdaily script
lock	Used to control serial use of the runacct script
log	Diagnostic output from the acctcon program
log.*MMDD*	Same as the log file after the runacct script detects an error
owtmpx	Previous day's wtmpx file
reboots	Beginning and ending dates from the wtmpx file, and a listing of reboots
statefile	Used to record current state during execution of the runacct script
tmpwtmp	wtmpx file corrected by the wtmpfix program
wtmperror	Contains wtmpfix error messages
wtmperror *MMDD*	Same as the wtmperror file after the runacct script detects an error
wtmp*MMDD*	The runacct script's copy of the wtmpx file

The sum directory contains the cumulative summary files updated by the runacct script and used by the monacct script. The following table summarizes the files in the /var/adm/acct/sum directory.

TABLE 17–9 Files in the /var/adm/acct/sum Directory

File	Description
cms	Total command summary file for current fiscal period in binary format
cmsprev	Command summary file without latest update
daycms	Command summary file for the day's usage in internal summary format
loginlog	Record of last date each user logged in; created by the lastlogin script and used in the prdaily script
rprt.*MMDD*	Saved output of prdaily script
tacct	Cumulative total accounting file for current fiscal period
tacctprev	Same as the tacct file without latest update
tacct.*MMDD*	Total accounting file for *MMDD*

The fiscal directory contains periodic summary files that are created by the monacct script. The following table summarizes the files in the /var/adm/acct/fiscal directory.

TABLE 17–10 Files in the /var/adm/acct/fiscal Directory

File	Description
cms*n*	Total command summary file for fiscal period *n* in internal summary format
fiscrpt*n*	Report similar to rprt*n* for fiscal period *n*
tacct*n*	Total accounting file for fiscal period *n*

Files Produced by the runacct Script

The following table summarizes the most useful files produced by the runacct script. These files are found in the /var/adm/acct directory.

TABLE 17–11 Files Created by the runacct Script

File	Description
nite/daytacct	The total accounting file for the day in tacct.h format.

TABLE 17–11 Files Created by the `runacct` Script *(Continued)*

File	Description
`nite/lineuse`	The `runacct` script calls the `acctcon` program to gather data on terminal line usage from the `/var/adm/acct/nite/tmpwtmp` file and writes the data to the `/var/adm/acct/nite/lineuse` file. The `prdaily` script uses this data to report line usage. This report is especially useful for detecting bad lines. If the ratio between the number of logouts to logins is greater than three to one, the line is very likely failing.
`sum/cms`	This file is the accumulation of each day's command summaries. The accumulation restarts when the `monacct` script is executed. The ASCII version is the `nite/cms` file.
`sum/daycms`	The `runacct` script calls the `acctcms` program to process the commands used during the day to create the Daily Command Summary report and stores the data in the `/var/adm/acct/sum/daycms` file. The ASCII version is the `/var/adm/acct/nite/daycms` file.
`sum/loginlog`	The `runacct` script calls the `lastlogin` script to update the last date logged in for the logins in the `/var/adm/acct/sum/loginlog` file. The `lastlogin` command also removes from this file any logins that are no longer valid.
`sum/rprt.`*MMDD*	Each execution of the `runacct` script saves a copy of the daily report that was printed by the `prdaily` script.
`sum/tacct`	Contains the accumulation of each day's `nite/daytacct` data and is used for billing purposes. The `monacct` script restarts accumulating this data each month or fiscal period.

CHAPTER 18

Managing System Performance (Overview)

Achieving good performance from a computer or network is an important part of system administration. This chapter provides an overview of some factors that contribute to managing the performance of the computer systems in your care.

This is a list of the overview information in this chapter.

What's New in Managing System Performance?

This section describes new features for managing system performance.

Enhanced `pfiles` Tool

The `/proc` file system has been enhanced to include file name information in the `/proc/`*pic*`/path` directory. This information is used by `pfiles` to display file names for each file in the process. This change provides new insight into process behavior. For more information, see "How to Display Information About Processes" on page 337 and the `proc`(1) man page.

CPU Performance Counters

The CPU Performance Counter (CPC) system is updated to give better access to the performance analysis features available in the SPARC and x86 platforms that run the Solaris Operating System.

The CPC commands `cpustat` and `cputrack` have enhanced, command-line syntax for specifying CPU information. For example, in previous versions of the Solaris OS, you were required to specify two counters. The configuration of both commands now allows you to specify only one counter, as shown in the following example:

```
# cputrack -c pic0=Cycle_cnt ls -d .
time lwp      event      pic0      pic1
.
0.034   1       exit    841167
```

For simple measurements, you can even omit the counter configuration, as shown in the following example:

```
# cputrack -c Cycle_cnt ls -d .
   time lwp      event      pic0      pic1
  .
  0.016   1       exit    850736
```

For more information on using the `cpustat` command, see the `cpustat`(1M) man page. For more information on using the `cputrack` command, see the `cputrack`(1) man page.

Where to Find System Performance Tasks

System Performance Task	For More Information
Manage processes	Chapter 19
Monitor system performance	Chapter 20
Change Solaris tunable parameters	*Solaris Tunable Parameters Reference Manual*
Manage System Performance Tasks	Chapter 2, "Projects and Tasks (Overview)," in *System Administration Guide: N1 Grid Containers, Resource Management, and Solaris Zones*

System Performance Task	For More Information
Manage Processes With FX and FS Schedulers	Chapter 8, "Fair Share Scheduler (Overview)," in *System Administration Guide: N1 Grid Containers, Resource Management, and Solaris Zones*

System Performance and System Resources

The performance of a computer system depends upon how the system uses and allocates its resources. Monitor your system's performance regularly so that you know how it behaves under normal conditions. You should have a good idea of what to expect, and be able to recognize a problem when it occurs.

System resources that affect performance are described in the following table.

System Resource	Description
Central processing unit (CPU)	The CPU processes instructions by fetching instructions from memory and executing them.
Input/output (I/O) devices	I/O devices transfer information into and out of the computer. Such a device could be a terminal and keyboard, a disk drive, or a printer.
Memory	Physical (or main) memory is the amount of random access memory (RAM) on the system.

Chapter 20 describes the tools that display statistics about the system's activity and performance.

Processes and System Performance

The following table describes terms that are related to processes.

TABLE 18–1 Process Terminology

Term	Description
Process	Any system activity or job. Each time you boot a system, execute a command, or start an application, the system activates one or more processes.
Lightweight process (LWP)	A virtual CPU or execution resource. LWPs are scheduled by the kernel to use available CPU resources based on their scheduling class and priority. LWPs include a kernel thread and an LWP. A kernel thread contains information that has to be in memory all the time. An LWP contains information that is swappable.
Application thread	A series of instructions with a separate stack that can execute independently in a user's address space. Application threads can be multiplexed on top of LWPs.

A process can consist of multiple LWPs and multiple application threads. The kernel schedules a kernel-thread structure, which is the scheduling entity in the SunOS environment. Various process structures are described in the following table.

TABLE 18–2 Process Structures

Structure	Description
`proc`	Contains information that pertains to the whole process and must be in main memory all the time
`kthread`	Contains information that pertains to one LWP and must be in main memory all the time
`user`	Contains the "per process" information that is swappable
`klwp`	Contains the "per LWP process" information that is swappable

The following figure illustrates the relationships among these process structures.

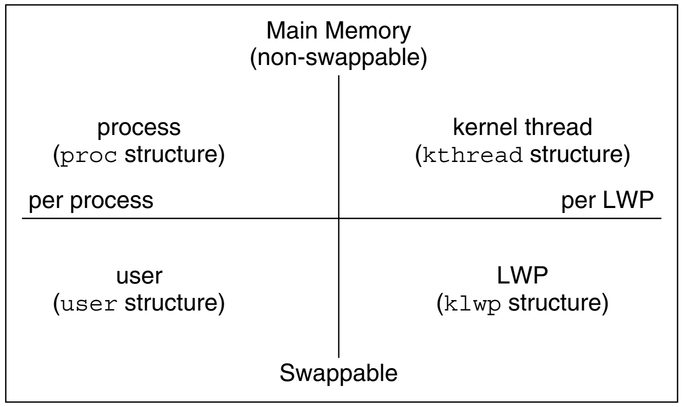

FIGURE 18–1 Relationships Among Process Structures

Most process resources are accessible to all the threads in the process. Almost all process virtual memory is shared. A change in shared data by one thread is available to the other threads in the process.

About Monitoring System Performance

While your computer is running, counters in the operating system are incremented to track various system activities. System activities that are tracked are as follows:

- Central processing unit (CPU) utilization
- Buffer usage
- Disk and tape input/output (I/O) activity
- Terminal device activity
- System call activity
- Context switching
- File access
- Queue activity
- Kernel tables
- Interprocess communication
- Paging
- Free memory and swap space
- Kernel memory allocation (KMA)

Monitoring Tools

The Solaris software provides several tools to help you track how your system is performing. The following table describes these tools.

TABLE 18–3 Performance Monitoring Tools

Command	Description	For More Information
cpustat and cputrack commands	Monitors performance of a system or a process using CPU performance counters.	cpustat(1M) and cputrack(1)
netstat and nfsstat commands	Displays information about network performance	netstat(1M) and nfsstat(1M)
ps and prstat commands	Displays information about active processes	Chapter 19
sar and sadc commands	Collects and reports on system activity data	Chapter 20
Sun Enterprise SyMON	Collects system activity data on Sun's enterprise-level systems	*Sun Enterprise SyMON 2.0.1 Software User's Guide*
swap command	Displays information about available swap space on your system	Chapter 20, "Configuring Additional Swap Space (Tasks)," in *System Administration Guide: Devices and File Systems*
vmstat and iostat commands	Summarizes system activity data, such as virtual memory statistics, disk usage, and CPU activity	Chapter 20
cputrack and cpustat commands	Assists in accessing hardware performance counter facilities provided by microprocessors	cputrack(1) and cpustat(1M) man pages
kstat and mpstat commands	Examines the available kernel statistics, or kstats, on the system and reports those statistics which match the criteria specified on the command line. The mpstat command reports processor statistics in tabular form.	kstat(1M) and mpstat(1M) man pages.

CHAPTER 19

Managing System Processes (Tasks)

This chapter describes the procedures for managing system processes.

For information on the procedures associated with managing system processes, see the following:

- "Managing System Processes (Task Map)" on page 331
- "Managing Process Class Information (Task Map)" on page 343

For overview information about managing system processes, see the following:

- "Commands for Managing System Processes" on page 332
- "Managing Process Class Information" on page 343

Managing System Processes (Task Map)

Task	Description	For Instructions
List processes.	Use the `ps` command to list all the processes on a system.	"How to List Processes" on page 336
Display information about processes.	Use the `pgrep` command to obtain the process IDs for processes that you want to display more information about.	"How to Display Information About Processes" on page 337

Task	Description	For Instructions
Control processes.	Locate processes by using the `pgrep` command. Then, use the appropriate `pcommand` (`/proc`) to control the process.) command that you want to run. See Table 19–3 for a description of the (`/proc`) commands.	"How to Control Processes" on page 338
Kill a process.	Locate a process, either by process name or process ID. You can use either the `pkill` or `kill` commands to terminate the process.	"How to Terminate a Process (`pkill`)" on page 339 "How to Terminate a Process (`kill`)" on page 340

Commands for Managing System Processes

The following table describes the commands for managing system processes.

TABLE 19–1 Commands for Managing Processes

Command	Description	Man Page
`ps`, `pgrep`, `prstat`, `pkill`	Checks the status of active processes on a system, as well as displays detailed information about the processes	`ps`(1), `pgrep`(1), and`prstat`(1M)
`pkill`	Functions identically to `pgrep` but finds or signals processes by name or other attribute and terminates the process. Each matching process is signaled as if by instead of having its process ID printed.	`pgrep`(1), and `pkill`(1)
`pargs, preap`	Assists with processes debugging	`pargs`(1), and `preap`(1)

TABLE 19–1 Commands for Managing Processes (Continued)

Command	Description	Man Page
`dispadmin`	Lists default process scheduling policies	`dispadmin`(1M)
`priocntl`	Assigns processes to a priority class and manages process priortities	`priocntl`(1)
`nice`	Changes the priority of a timesharing process	`nice`(1)
`psrset`	Binds specific process groups to a group of processors rather than to just a single processor	`psrset`(1M)

The Solaris Management Console's Processes tool enables you to manage processes with a user-friendly interface. For information on using and starting the Solaris Management Console, see Chapter 2, "Working With the Solaris Management Console (Tasks)," in *System Administration Guide: Basic Administration*.

Using the `ps` Command

The `ps` command enables you to check the status of active processes on a system, as well as display technical information about the processes. This data is useful for administrative tasks such as determining how to set process priorities.

Depending on which options you use, the `ps` command reports the following information:

- Current status of the process
- Process ID
- Parent process ID
- User ID
- Scheduling class
- Priority
- Address of the process
- Memory used
- CPU time used

The following table describes some fields that are reported by the `ps` command. Which fields are displayed depend on which option you choose. For a description of all available options, see the `ps`(1) man page.

TABLE 19–2 Summary of Fields in ps Reports

Field	Description
UID	The effective user ID of the process's owner.
PID	The process ID.
PPID	The parent process ID.
C	The processor xutilization for scheduling. This field is not displayed when the -c option is used.
CLS	The scheduling class to which the process belongs such as real-time, system, or timesharing. This field is included only with the -c option.
PRI	The kernel thread's scheduling priority. Higher numbers indicate a higher priority.
NI	The process's nice number, which contributes to its scheduling priority. Making a process "nicer" means lowering its priority.
ADDR	The address of the proc structure.
SZ	The virtual address size of the process.
WCHAN	The address of an event or lock for which the process is sleeping.
STIME	The starting time of the process in hours, minutes, and seconds.
TTY	The terminal from which the process, or its parent, was started. A question mark indicates that there is no controlling terminal.
TIME	The total amount of CPU time used by the process since it began.
CMD	The command that generated the process.

Using the /proc File System and Commands

You can display detailed information about the processes that are listed in the /proc directory by using process commands. The following table lists the /proc process commands. The /proc directory is also known as the process file system (PROCFS). Images of active processes are stored here by their process ID number.

TABLE 19–3 Process Commands (/proc)

Process Command	Description
pcred	Displays process credential information
pfiles	Reports fstat and fcntl information for open files in a process

TABLE 19–3 Process Commands (/proc) *(Continued)*

Process Command	Description
pflags	Prints /proc tracing flags, pending signals and held signals, and other status information
pldd	Lists the dynamic libraries that are linked into a process
pmap	Prints the address space map of each process
psig	Lists the signal actions and handlers of each process
prun	Starts each process
pstack	Prints a hex+symbolic stack trace for each lwp in each process
pstop	Stops each process
ptime	Times a process by using microstate accounting
ptree	Displays the process trees that contain the process
pwait	Displays status information after a process terminates
pwdx	Displays the current working directory for a process

For more information, see proc(1).

The process tools are similar to some options of the ps command, except that the output that is provided by these commands is more detailed. In general, the process commands do the following:

- Display more information about processes, such as fstat and fcntl, working directories, and trees of parent and child processes
- Provide control over processes by allowing users to stop or resume them

Managing Processes With Process Commands (/proc)

You can display detailed, technical information about processes or control active processes by using some of the process commands. Table 19–3 lists some of the /proc commands.

If a process becomes trapped in an endless loop, or if the process takes too long to execute, you might want to stop (kill) the process. For more information about stopping processes using the kill or the pkill command, see Chapter 19.

The /proc file system is a directory hierarchy that contains additional subdirectories for state information and control functions.

The /proc file system also provides an xwatchpoint facility that is used to remap read-and-write permissions on the individual pages of a process's address space. This facility has no restrictions and is MT-safe.

Debugging tools have been modified to use /proc's xwatchpoint facility, which means that the entire xwatchpoint process is faster.

The following restrictions have been removed when you set xwatchpoints by using the dbx debugging tool:

- Setting xwatchpoints on local variables on the stack due to SPARC based system register windows
- Setting xwatchpoints on multithreaded processes

For more information, see the proc(4), and mdb(1) man pages.

▼ How to List Processes

Step ● **Use the ps command to list all the processes on a system.**

```
$ ps [-efc]
```

ps — Displays only the processes that are associated with your login session.

-ef — Displays full information about all the processes that are being executed on the system.

-c — Displays process scheduler information.

Example 19–1 Listing Processes

The following example shows output from the ps command when no options are used.

```
$ ps
   PID TTY      TIME COMD
  1664 pts/4    0:06 csh
  2081 pts/4    0:00 ps
```

The following example shows output from the ps -ef command. This output shows that the first process that is executed when the system boots is sched (the swapper) followed by the init process, pageout, and so on.

```
$ ps -ef
     UID   PID  PPID  C    STIME TTY      TIME CMD
    root     0     0  0   Dec 20 ?        0:17 sched
    root     1     0  0   Dec 20 ?        0:00 /etc/init -
    root     2     0  0   Dec 20 ?        0:00 pageout
    root     3     0  0   Dec 20 ?        4:20 fsflush
    root   374   367  0   Dec 20 ?        0:00 /usr/lib/saf/ttymon
```

```
    root   367      1  0   Dec 20 ?        0:00 /usr/lib/saf/sac -t 300
    root   126      1  0   Dec 20 ?        0:00 /usr/sbin/rpcbind
    root    54      1  0   Dec 20 ?        0:00 /usr/lib/sysevent/syseventd
    root    59      1  0   Dec 20 ?        0:00 /usr/lib/picl/picld
    root   178      1  0   Dec 20 ?        0:03 /usr/lib/autofs/automountd
    root   129      1  0   Dec 20 ?        0:00 /usr/sbin/keyserv
    root   213      1  0   Dec 20 ?        0:00 /usr/lib/lpsched
    root   154      1  0   Dec 20 ?        0:00 /usr/sbin/inetd -s
    root   139      1  0   Dec 20 ?        0:00 /usr/lib/netsvc/yp/ypbind ...
    root   191      1  0   Dec 20 ?        0:00 /usr/sbin/syslogd
    root   208      1  0   Dec 20 ?        0:02 /usr/sbin/nscd
    root   193      1  0   Dec 20 ?        0:00 /usr/sbin/cron
    root   174      1  0   Dec 20 ?        0:00 /usr/lib/nfs/lockd
  daemon   175      1  0   Dec 20 ?        0:00 /usr/lib/nfs/statd
    root   376      1  0   Dec 20 ?        0:00 /usr/lib/ssh/sshd
    root   226      1  0   Dec 20 ?        0:00 /usr/lib/power/powerd
    root   315      1  0   Dec 20 ?        0:00 /usr/lib/nfs/mountd
    root   237      1  0   Dec 20 ?        0:00 /usr/lib/utmpd
    .
    .
    .
```

▼ How to Display Information About Processes

Steps

1. **Obtain the process ID of the process that you want to display more information about.**

   `# pgrep` *process*

   where *process* is the name of the process you want to display more information about.

   The process ID is displayed in the first column of the output.

2. **Display the process information that you need.**

   `# /usr/bin/`*pcommand pid*

   pcommand — Is the (`/proc`) command that you want to run. Table 19–3 lists and describes these commands.

   pid — Identifies the process ID.

Example 19–2 Displaying Information About Processes

The following example shows how to use process commands to display more information about a `cron` process.

```
# pgrep cron      1
4780
# pwdx 4780      2
```

```
4780:   /var/spool/cron/atjobs
# ptree 4780     3
4780  /usr/sbin/cron
# pfiles 4780     4
4780:   /usr/sbin/cron
  Current rlimit: 256 file descriptors
   0: S_IFCHR mode:0666 dev:290,0 ino:6815752 uid:0 gid:3 rdev:13,2
      O_RDONLY|O_LARGEFILE
      /devices/pseudo/mm@0:null
   1: S_IFREG mode:0600 dev:32,128 ino:42054 uid:0 gid:0 size:9771
      O_WRONLY|O_APPEND|O_CREAT|O_LARGEFILE
      /var/cron/log
   2: S_IFREG mode:0600 dev:32,128 ino:42054 uid:0 gid:0 size:9771
      O_WRONLY|O_APPEND|O_CREAT|O_LARGEFILE
      /var/cron/log
   3: S_IFIFO mode:0600 dev:32,128 ino:42049 uid:0 gid:0 size:0
      O_RDWR|O_LARGEFILE
      /etc/cron.d/FIFO
   4: S_IFIFO mode:0000 dev:293,0 ino:4630 uid:0 gid:0 size:0
      O_RDWR|O_NONBLOCK
   5: S_IFIFO mode:0000 dev:293,0 ino:4630 uid:0 gid:0 size:0
      O_RDWR
```

1. Obtains the process ID for the `lpsched` process
2. Displays the current working directory for the `lpsched` process
3. Displays the process tree that contains the `lpsched` process
4. Displays `fstat` and `fcntl` information

▼ How to Control Processes

Steps

1. **Obtain the process ID of the process that you want to control.**

   `# pgrep` *process*

   where *process* is the name of the process you want to control.

   The process ID displayed in the first column of the output.

2. **Use the appropriate process command to control the process.**

   `# /usr/bin/`*pcommand pid*

   pcommand — Is the process (`/proc`) command that you want to run. Table 19–3 lists and describes these commands.

   pid — Identifies the process ID.

3. **Verify the process status.**

   `# ps -ef | grep` *pid*

Example 19–3 Controlling Processes

The following example shows how to use process command to stop and restart the `dtpad` process.

```
# pgrep dtpad       Obtains the process ID for the dtpad process.
2921
# pstop 2921      Stops the dtpad process.
# prun 2921       3
```

1. Obtains the process ID for the `dtpad` process
2. Stops the `dtpad` process
3. Restarts the `dtpad` process

Terminating a Process (`pkill, kill`)

Sometimes, you might need to stop (kill) a process. The process might be in an endless loop. Or, you might have started a large job that you want to stop before it is completed. You can kill any process that you own. Superuser can kill any process in the system except for those processes with process IDs of `0`, `1`, `2`, `3`, and `4`. Killing these processes most likely will crash the system.

For more information, see the `pgrep`(1) and `pkill`(1) and `kill`(1) man pages.

▼ How to Terminate a Process (`pkill`)

Steps

1. **(Optional) Become superuser or assume an equivalent role to terminate the process of another user.**
2. **Obtain the process ID for the process that you want to terminate.**

   ```
   $ pgrep process
   ```

   where *process* is the name of the process that you want to terminate.

   For example:

   ```
   $ pgrep netscape
   587
   566
   ```

   The process ID is displayed in the output.

Note – To obtain process information on a Sun Ray™, use the following commands:

```
# ps -fu user
```

This command lists all user processes.

```
# ps -fu user | grep process
```

This command locates a specific process for a user.

3. **Terminate the process.**

   ```
   $ pkill [signal] process
   ```

   signal When no signal is included in the `pkill` command-line syntax, the default signal that is used is –15 (`SIGTERM`). Using the –9 signal (`SIGKILL`) with the `pkill` command ensures that the process terminates promptly. However, the –9 signal should not be used to kill certain processes, such as a database process, or an LDAP server process. The result is that data might be lost.

   process Is the name of the process to stop.

   Tip – When using the `pkill` command to terminate a process, first try using the command by itself, without including a signal option. Wait a few minutes to see if the process terminates before using the `pkill` command with the -9 signal.

4. **Verify that the process has been terminated.**

   ```
   $ pgrep process
   ```

   The process you terminated should no longer be listed in the output of the `pgrep` command.

▼ How to Terminate a Process (`kill`)

Steps

1. **(Optional) Become superuser or assume an equivalent role to terminate the process of another user.**

2. **Obtain the process ID of the process that you want to terminate.**

   ```
   $ ps -fu user
   ```

   where *user* is the user that you want to display processes for.

For example:

```
$ ps -fu userabc
userabc   328   323   2   Mar 12 ?        10:18 /usr/openwin/bin/Xsun
:0 -nobanner -auth /var/dt/A:0-WmayOa
userabc   366   349   0   Mar 12 ?        0:00 /usr/openwin/bin/fbconsole
userabc   496   485   0   Mar 12 ?        0:09 /usr/dt/bin/sdtperfmeter
 -f -H -t cpu -t disk -s 1 -name fpperfmeter
userabc   349   332   0   Mar 12 ?        0:00 /bin/ksh /usr/dt/bin/Xsession
userabc   440   438   0   Mar 12 pts/3    0:00 -csh -c unsetenv _ PWD;
unsetenv DT;     setenv DISPLAY :0;
userabc   372     1   0   Mar 12 ?        0:00 /usr/openwin/bin/speckeysd
userabc   438   349   0   Mar 12 pts/3    0:00 /usr/dt/bin/sdt_shell -c
unset
.
.
.
```

The process ID is displayed in the first column of the output.

3. **Terminate the process.**

   ```
   $ kill [signal-number] pid
   ```

   signal When no signal is included in the `kill` command-line syntax, the default signal that is used is –15 (`SIGKILL`). Using the –9 signal (`SIGTERM`) with the `kill` command ensures that the process terminates promptly. However, the –9 signal should not be used to kill certain processes, such as a database process, or an LDAP server process. The result is that data might be lost.

   pid Is the process ID of the process that you want to terminate.

   Tip – When using the `kill` command to stop a process, first try using the command by itself, without including a signal option. Wait a few minutes to see if the process terminates before using the `kill` command with the -9 signal.

   ---

4. **Verify that the process has been terminated.**

   ```
   $ pgrep pid
   ```

   The process you terminated should no longer be listed in the output of the `pgrep` command.

Debugging a Process (pargs, preap)

The pargs command and the preap command improve process debugging. The pargs command prints the arguments and environment variables associated with a live process or core file. The preap command removes defunct (zombie) processes. A zombie process has not yet had its exit status claimed by its parent. These processes are generally harmless but can consume system resources if they are numerous. You can use the pargs and preap commands to examine any process that you have the privileges to examine. As superuser, you can examine any process.

For information on using the preap command, see the preap(1) man page. For information on the using the pargs command, see the pargs(1) man page. See also, the proc(1) man page.

EXAMPLE 19–4 Debugging a Process (pargs)

The pargs command solves a long-standing problem of being unable to display with the ps command all the arguments that are passed to a process. The following example shows how to use the pargs command in combination with the pgrep command to display the arguments that are passed to a process.

```
# pargs `pgrep ttymon`
579:    /usr/lib/saf/ttymon -g -h -p system-name console login:
-T sun -d /dev/console -l
argv[0]: /usr/lib/saf/ttymon
argv[1]: -g
argv[2]: -h
argv[3]: -p
argv[4]: system-name console login:
argv[5]: -T
argv[6]: sun
argv[7]: -d
argv[8]: /dev/console
argv[9]: -l
argv[10]: console
argv[11]: -m
argv[12]: ldterm,ttcompat
548:    /usr/lib/saf/ttymon
argv[0]: /usr/lib/saf/ttymon
```

The following example shows how to use the pargs -e command to display the environment variables that are associated with a process.

```
$ pargs -e 6763
6763: tcsh
envp[0]: DISPLAY=:0.0
```

Managing Process Class Information (Task Map)

Task	Description	For Instructions
Display basic information about process classes.	Use the `priocntl -l` command. to Display process scheduling classes and priority ranges.	"How to Display Basic Information About Process Classes (`priocntl`)" on page 344
Display the global priority of a process.	Use the `ps -ecl` command to display the global priority of a process.	"How to Display the Global Priority of a Process" on page 345
Designate a process priority.	Start a process with a designated priority by using the `priocntl -e -c` `command.`	"How to Designate a Process Priority (`priocntl`)" on page 345
Change scheduling parameters of a timesharing process.	Use the `priocntl -s -m` command to change scheduling parameters in a timesharing process.	"How to Change Scheduling Parameters of a Timesharing Process (`priocntl`)" on page 346
Change the class of a process.	Use the `priocntl -s -c` command to change the class of a process.	"How to Change the Class of a Process (`priocntl`)" on page 347
Change the priority of a process.	Use the `/usr/bin/nice` command with the appropriate options to lower or raise the priority of a process.	"How to Change the Priority of a Process (`nice`)" on page 348

Managing Process Class Information

The following list identifies the process scheduling classes that can be configured on your system. Also included is the user priority range for the timesharing class. The possible process scheduling classes are as follows:

- Fair share (`FSS`)
- Fixed (`FX`)
- System (`SYS`)

- Interactive (IA)
- Real-time (RT)
- Timesharing (TS)
  - The user-supplied priority ranges from -60 to +60.
  - The priority of a process is inherited from the parent process. This priority is referred to as the *user-mode priority*.
  - The system looks up the user-mode priority in the timesharing dispatch parameter table. Then, the system adds in any `nice` or `priocntl` (user-supplied) priority and ensures a 0–59 range to create a *global priority*.

Changing the Scheduling Priority of Processes (`priocntl`)

The scheduling priority of a process is the priority assigned by the process scheduler, according to scheduling policies. The `dispadmin` command lists the default scheduling policies. For more information, see the `dispadmin`(1M) man page.

You can use the `priocntl` command to assign processes to a priority class and to manage process priorities. For instructions on using the `priocntl` command to manage processes, see "How to Designate a Process Priority (`priocntl`)" on page 345.

▼ How to Display Basic Information About Process Classes (`priocntl`)

Step ● **Display process scheduling classes and priority ranges with the `priocntl -l` command.**

```
$ priocntl -l
```

Example 19–5 Displaying Basic Information About Process Classes (`priocntl`)

The following example shows output from the `priocntl -l` command.

```
# priocntl -l
CONFIGURED CLASSES
==================

SYS (System Class)

TS (Time Sharing)
        Configured TS User Priority Range: -60 through 60
```

```
FX (Fixed priority)
        Configured FX User Priority Range: 0 through 60

IA (Interactive)
        Configured IA User Priority Range: -60 through 60
```

▼ How to Display the Global Priority of a Process

Step ● **Display the global priority of a process by using the `ps` command.**

```
$ ps -ecl
```

The global priority is listed under the `PRI` column.

Example 19–6 Displaying the Global Priority of a Process

The following example shows `ps -ecl` command output. The values in the `PRI` column show that the `pageout` process has the highest priority, while the `sh` process has the lowest priority.

```
$ ps -ecl
 F S UID PID  PPID CLS PRI  ADDR      SZ  WCHAN    TTY       TIME   COMD
19 T 0    0    0    SYS 96   f00d05a8   0           ?        0:03  sched
 8 S 0    1    0    TS  50   ff0f4678 185  ff0f4848 ?       36:51   init
19 S 0    2    0    SYS 98   ff0f4018   0  f00c645c ?        0:01 pageout
19 S 0    3    0    SYS 60   ff0f5998   0  f00d0c68 ?      241:01 fsflush
 8 S 0   269   1    TS  58   ff0f5338 303  ff49837e ?        0:07    sac
 8 S 0   204   1    TS  43   ff2f6008  50  ff2f606e console  0:02     sh
```

▼ How to Designate a Process Priority (`priocntl`)

Steps

1. **(Optional) Assume the Primary Administrator role, or become superuser.**

   The Primary Administrator role includes the Primary Administrator profile. To create the role and assign the role to a user, see Chapter 2, "Working With the Solaris Management Console (Tasks)," in *System Administration Guide: Basic Administration*.

2. **Start a process with a designated priority.**

   `# `**`priocntl -e -c`** *class* **`-m`** *user-limit* **`-p`** *pri command-name*

`-e`	Executes the command.
`-c` *class*	Specifies the class within which to run the process. The valid classes are `TS` (timesharing), `RT` (real time), IA (interactive), FSS (fair share), and FX (fixed priority).

`-m` *user-limit*	When you use the `-p` option, specifies the maximum amount you can raise or lower your priority,
`-p` *pri command-name*	Lets you specify the relative priority in the RT class for a real-time thread. For a timesharing process, the `-p` option lets you specify the user-supplied priority, which ranges from -60 to +60.

3. **Verify the process status.**

```
# ps -ecl | grep command-name
```

Example 19–7 Designating a Process Priority (`priocntl`)

The following example shows how to start the `find` command with the highest possible user-supplied priority.

```
# priocntl -e -c TS -m 60 -p 60 find . -name core -print
# ps -ecl | grep find
```

▼ How to Change Scheduling Parameters of a Timesharing Process (`priocntl`)

Steps

1. **(Optional) Assume the Primary Administrator role, or become superuser.**

   The Primary Administrator role includes the Primary Administrator profile. To create the role and assign the role to a user, see Chapter 2, "Working With the Solaris Management Console (Tasks)," in *System Administration Guide: Basic Administration*.

2. **Change the scheduling parameters of a running timesharing process.**

```
# priocntl -s -m user-limit [-p user-priority] -i idtype idlist
```

`-s`	Lets you set the upper limit on the user priority range and change the current priority.
`-m` *user-limit*	When you use the `-p` option, specifies the maximum amount you can raise or lower the priority.
`-p` *user-priority*	Allows you to designate a priority.
`-i` *xidtype xidlist*	Uses a combination of *xidtype* and *xidlist* to identify the process or processes. The *xidtype* specifies the type of ID, such as the process ID or the user ID. Use *xidlist* to identify a list of process IDs or user IDs.

3. **Verify the process status.**

```
# ps -ecl | grep idlist
```

Example 19–8 Changing Scheduling Parameters of a Timesharing Process (`priocntl`)

The following example shows how to execute a command with a 500-millisecond time slice, a priority of 20 in the RT class, and a global priority of 120.

```
# priocntl -e -c RT -m 500 -p 20 myprog
# ps -ecl | grep myprog
```

▼ How to Change the Class of a Process (`priocntl`)

Steps

1. **(Optional) Become superuser or assume an equivalent role.**
2. **Change the class of a process.**

   ```
   # priocntl -s -c class -i idtype idlist
   ```

   `-s` Lets you set the upper limit on the user priority range and change the current priority.

   `-c` *class* Specifies the class, `TS` for time-sharing or `RT for real-time`, to which you are changing the process.

   `-i` *idtype idlist* Uses a combination of *xidtype* and *xidlist* to identify the process or processes. The *xidtype* specifies the type of ID, such as the process ID or user ID. Use *xidlist* to identify a list of process IDs or user IDs.

   Note – You must be superuser or working in a real-time shell to change a process from, or to, a real-time process. If, as superuser, you change a user process to the real-time class, the user cannot subsequently change the real-time scheduling parameters by using the `priocntl -s` command.

   ---

3. **Verify the process status.**

   ```
   # ps -ecl | grep idlist
   ```

Example 19–9 Changing the Class of a Process (`priocntl`)

The following example shows how to change all the processes that belong to user `15249` to real-time processes.

```
# priocntl -s -c RT -i uid 15249
# ps -ecl | grep 15249
```

Changing the Priority of a Timesharing Process (`nice`)

The `nice` command is only supported for backward compatibility to previous Solaris releases. The `priocntl` command provides more flexibility in managing processes.

The priority of a process is determined by the policies of its scheduling class and by its *nice number*. Each timesharing process has a global priority. The global priority is calculated by adding the user-supplied priority, which can be influenced by the `nice` or `priocntl` commands, and the system-calculated priority.

The execution priority number of a process is assigned by the operating system. The priority number is determined by several factors, including the process's scheduling class, how much CPU time it has used, and in the case of a timesharing process, its `nice` number.

Each timesharing process starts with a default `nice` number, which it inherits from its parent process. The `nice` number is shown in the `NI` column of the `ps` report.

A user can lower the priority of a process by increasing its user-supplied priority. However, only superuser can lower a `nice` number to increase the priority of a process. This restriction prevents users from increasing the priorities of their own processes, thereby monopolizing a greater share of the CPU.

The `nice` numbers range from 0 to +39, with 0 representing the highest priority. The default `nice` value for each timesharing process is 20. Two versions of the command are available: the standard version, `/usr/bin/nice`, and the C shell built-in command.

▼ How to Change the Priority of a Process (`nice`)

Using this procedure, a user can lower the priority of a process. However, superuser can raise or lower the priority of a process.

Note – This section describes the syntax of the `/usr/bin/nice` command and not the C-shell built-in `nice`command. For information about the C-shell `nice` command, see the `csh`(1) man page.

Steps

1. **Determine whether you want to change the priority of a process, either as a user or as superuser. Then, select one of the following:**
   - As a user, follow the examples in Step 2 to lower the priority of a command.
   - As a superuser, follow the examples in Step 3 to raise or lower priorities of a command.

2. **As a user, lower the priority of a command by increasing the `nice` number.**

   The following `nice` command executes *command-name* with a lower priority by raising the nice number by 5 units.

   ```
   $ /usr/bin/nice -5 command-name
   ```

   In the preceding command, the minus sign designates that what follows is an option. This command could also be specified as follows:

   ```
   % /usr/bin/nice -n 5 command-name
   ```

   The following `nice` command lowers the priority of *command-name* by raising the nice number by the default increment of 10 units, but not beyond the maximum value of 39.

   ```
   % /usr/bin/nice command-name
   ```

3. **As superuser or assuming an equivalent role, raise or lower the priority of a command by changing the `nice` number.**

   The following `nice` command raises the priority of *command-name* by lowering the nice number by 10 units, but not below the minimum value of 0.

   ```
   # /usr/bin/nice --10 command-name
   ```

   In the preceding command, the first minus sign designates that what follows is an option. The second minus sign indicates a negative number.

   The following `nice` command lowers the priority of *command-name* by raising the nice number by 5 units, but not beyond the maximum value of 39.

   ```
   # /usr/bin/nice -5 command-name
   ```

See Also For more information, see the `nice`(1) man page.

Troubleshooting Problems With System Processes

Here are some tips on obvious problems you might encounter:

- Look for several identical jobs that are owned by the same user. This problem might occur because of a running script that starts a lot of background jobs without waiting for any of the jobs to finish.
- Look for a process that has accumulated a large amount of CPU time. You can identify this problem by checking the `TIME` field in the `ps` output. Possibly, the process is in an endless loop.
- Look for a process that is running with a priority that is too high. Use the `ps -c` command to check the `CLS` field, which displays the scheduling class of each process. A process executing as a real-time (`RT`) process can monopolize the CPU. Or, look for a timesharing (`TS`) process with a high `nice` number. A user with superuser privileges might have increased the priority of a process. The system administrator can lower the priority by using the `nice` command.
- Look for a runaway process. A runaway process progressively uses more and more CPU time. You can identify this problem by looking at the time when the process started (`STIME`) and by watching the cumulation of CPU time (`TIME`) for a while.

CHAPTER **20**

Monitoring System Performance (Tasks)

This chapter describes procedures for monitoring system performance by using the `vmstat`, `iostat`, `df`, and `sar` commands.

For information about the procedures associated with monitoring system performance, see the following:

- "Displaying System Information (Task Map)" on page 351
- "Monitoring System Activities (Task Map)" on page 360

Displaying System Information (Task Map)

Task	Description	For Instructions
Display virtual memory Statistics.	Collect virtual memory statistics by using the `vmstat` command.	"How to Display Virtual Memory Statistics (`vmstat`)" on page 353
Display system event information.	Display system event information by using the `vmstat` command with the `-s` option	"How to Display System Event Information (`vmstat -s`)" on page 354
Display swapping statistics.	Use the `vmstat` command with the `-S` option to display swapping statistics.	"How to Display Swapping Statistics (`vmstat -S`)" on page 355

Task	Description	For Instructions
Display interrupts per device.	Use the `vmstat` command with the `-i` option to show the number of interrupts per device.	"How to Display Interrupts Per Device (`vmstat -i`)" on page 355
Display disk utilization.	Use the `iostat` command to report disk input and output statistics.	"How to Display Disk Utilization Information (`iostat`)" on page 356
Display extended disk statistics.	Use the `iostat` command with the `-xtc` option to display extended disk statistics.	"How to Display Extended Disk Statistics (`iostat -xtc`)" on page 357
Display disk space information.	The `df -k` command displays disk space information in Kbytes.	"How to Display Disk Space Information (`df -k`)" on page 359

Displaying Virtual Memory Statistics (`vmstat`)

You can use the `vmstat` command to report virtual memory statistics and information about system events such as CPU load, paging, number of context switches, device interrupts, and system calls. The `vmstat` command can also display statistics on swapping, cache flushing, and interrupts.

The following table describes the fields in the `vmstat` command output.

TABLE 20–1 Output From the `vmstat` Command

Category	Field Name	Description
`procs`		Reports on the following:
	`r`	The number of kernel threads in the dispatch queue
	`b`	The number of blocked kernel threads that are waiting for resources
	`w`	The number of swapped out LWPs that are waiting for processing resources to finish
`memory`		Reports on usage of real memory and virtual memory:
	`swap`	Available swap space

TABLE 20–1 Output From the `vmstat` Command *(Continued)*

Category	Field Name	Description
	`free`	Size of the free list
`page`		Reports on page faults and paging activity, in units per second:
	`re`	Pages reclaimed
	`mf`	Minor faults and major faults
	`pi`	Kbytes paged in
	`po`	Kbytes paged out
	`fr`	Kbytes freed
	`de`	Anticipated memory that is needed by recently swapped-in processes
	`sr`	Pages scanned by the `page` daemon not currently in use. If `sr` does not equal zero, the `page` daemon has been running.
`disk`		Reports the number of disk operations per second, showing data on up to four disks
`faults`		Reports the trap/interrupt rates per second:
	`in`	Interrupts per second
	`sy`	System calls per second
	`cs`	CPU context switch rate
`cpu`		Reports on the use of CPU time:
	`us`	User time
	`sy`	System time
	`id`	Idle time

For a more detailed description of this command, see the `vmstat`(1M) man page.

▼ How to Display Virtual Memory Statistics (`vmstat`)

Step ● **Collect virtual memory statistics by using the `vmstat` command with a time interval in seconds.**

```
$ vmstat n
```

where *n* is the interval in seconds between reports.

Example 20–1 Displaying Virtual Memory Statistics

The following example shows the `vmstat` display of statistics that were gathered at five-second intervals.

```
$ vmstat 5
kthr      memory            page            disk          faults      cpu
 r b w   swap  free  re  mf pi po fr de sr dd f0 s1 --   in   sy   cs us sy id
 0 0 0 863160 365680  0   3  1  0  0  0  0  0  0  0  0  406  378  209  1  0 99
 0 0 0 765640 208568  0  36  0  0  0  0  0  0  0  0  0  479 4445 1378  3  3 94
 0 0 0 765640 208568  0   0  0  0  0  0  0  0  0  0  0  423  214  235  0  0 100
 0 0 0 765712 208640  0   0  0  0  0  0  0  3  0  0  0  412  158  181  0  0 100
 0 0 0 765832 208760  0   0  0  0  0  0  0  0  0  0  0  402  157  179  0  0 100
 0 0 0 765832 208760  0   0  0  0  0  0  0  0  0  0  0  403  153  182  0  0 100
 0 0 0 765832 208760  0   0  0  0  0  0  0  0  0  0  0  402  168  177  0  0 100
 0 0 0 765832 208760  0   0  0  0  0  0  0  0  0  0  0  402  153  178  0  0 100
 0 0 0 765832 208760  0  18  0  0  0  0  0  0  0  0  0  407  165  186  0  0 100
```

▼ How to Display System Event Information (`vmstat -s`)

Step ● **Run the `vmstat -s` command to show how many system events have taken place since the last time the system was booted.**

```
$ vmstat -s
        0 swap ins
        0 swap outs
        0 pages swapped in
        0 pages swapped out
   522586 total address trans. faults taken
    17006 page ins
       25 page outs
    23361 pages paged in
       28 pages paged out
    45594 total reclaims
    45592 reclaims from free list
        0 micro (hat) faults
   522586 minor (as) faults
    16189 major faults
    98241 copy-on-write faults
   137280 zero fill page faults
    45052 pages examined by the clock daemon
        0 revolutions of the clock hand
       26 pages freed by the clock daemon
     2857 forks
       78 vforks
     1647 execs
```

```
34673885 cpu context switches
65943468 device interrupts
  711250 traps
63957605 system calls
 3523925 total name lookups (cache hits 99%)
   92590 user   cpu
   65952 system cpu
16085832 idle   cpu
    7450 wait   cpu
```

How to Display Swapping Statistics (`vmstat -S`)

Step ● **Run `vmstat -S` to show swapping statistics.**

```
$ vmstat -S
kthr      memory            page            disk          faults      cpu
 r b w   swap  free  si   so pi po fr de sr dd f0 s1 --   in   sy   cs us sy id
 0 0 0 862608 364792  0    0  1  0  0  0  0  0  0  0  0  406  394  213  1  0 99
```

The swapping statistics fields are described in the following list. For a description of the other fields, see Table 20–1.

`si` Average number of LWPs that are swapped in per second

`so` Number of whole processes that are swapped out

Note – The `vmstat` command truncates the output of `si` and `so` fields. Use the `sar` command to display a more accurate accounting of swap statistics.

How to Display Interrupts Per Device (`vmstat -i`)

Step ● **Run the `vmstat -i` command to show the number of interrupts per device.**

Example 20–2 Displaying Interrupts Per Device

The following example shows output from the `vmstat -i` command.

```
$ vmstat -i
interrupt         total     rate
--------------------------------
clock          52163269      100
esp0            2600077        4
zsc0              25341        0
```

```
zsc1              48917        0
cgsixc0             459        0
lec0             400882        0
fdc0                 14        0
bppc0                 0        0
audiocs0              0        0
--------------------------------
Total          55238959      105
```

Displaying Disk Utilization Information (iostat)

Use the `iostat` command to report statistics about disk input and output, and to produce measures of throughput, utilization, queue lengths, transaction rates, and service time. For a detailed description of this command, refer to the `iostat`(1M) man page.

▼ How to Display Disk Utilization Information (iostat)

Step ● **You can display disk utilization information by using the `iostat` command with a time interval in seconds.**

```
$ iostat 5
     tty          fd0           sd3           nfs1          nfs31           cpu
tin tout kps tps serv  kps tps serv  kps tps serv  kps tps serv  us sy wt id
  0    1   0   0  410    3   0   29    0   0    9    3   0   47   4  2  0 94
```

The first line of output shows the statistics since the last time the system was booted. Each subsequent line shows the interval statistics. The default is to show statistics for the terminal (`tty`), disks (`fd` and `sd`), and CPU (`cpu`).

Example 20–3 Displaying Disk Utilization Information

The following example shows disk statistics that were gathered every five seconds.

```
$ iostat 5
tty        sd0           sd6           nfs1          nfs49           cpu
tin tout kps tps serv  kps tps serv  kps tps serv  kps tps serv  us sy wt id
 0    0   1   0   49    0   0    0    0   0    0    0   0   15   0  0  0 100
 0   47   0   0    0    0   0    0    0   0    0    0   0    0   0  0  0 100
```

```
 0   16   0   0    0    0   0    0    0   0    0    0   0    0   0   0   0 100
 0   16   0   0    0    0   0    0    0   0    0    0   0    0   0   0   0 100
 0   16  44   6  132    0   0    0    0   0    0    0   0    0   0   0   1 99
 0   16   0   0    0    0   0    0    0   0    0    0   0    0   0   0   0 100
 0   16   0   0    0    0   0    0    0   0    0    0   0    0   0   0   0 100
 0   16   0   0    0    0   0    0    0   0    0    0   0    0   0   0   0 100
 0   16   0   0    0    0   0    0    0   0    0    0   0    0   0   0   0 100
 0   16   0   0    0    0   0    0    0   0    0    0   0    0   0   0   0 100
 0   16   3   1   23    0   0    0    0   0    0    0   0    0   0   0   1 99
 0   16   0   0    0    0   0    0    0   0    0    0   0    0   0   0   0 100
 0   16   0   0    0    0   0    0    0   0    0    0   0    0   0   0   0 100
 0   16   0   0    0    0   0    0    0   0    0    0   0    0   0   0   0 100
```

The following table describes the fields in the output of the `iostat` *n* command.

Device Type	Field Name	Description
Terminal	Device Type	
	`tin`	Number of characters in the terminal input queue
	`tout`	Number of characters in the terminal output queue
Disk	Device Type	
	`bps`	Blocks per second
	`tps`	Transactions per second
	`serv`	Average service time, in milliseconds
CPU	Device Type	
	`us`	In user mode
	`sy`	In system mode
	`wt`	Waiting for I/O
	`id`	Idle

▼ How to Display Extended Disk Statistics (`iostat -xtc`)

Step ● **Run the `iostat -xtc` command to display extended disk statistics.**

```
$ iostat -xtc
                 extended device statistics                      tty         cpu
device       r/s    w/s   kr/s    kw/s wait actv  svc_t  %w  %b  tin tout  us sy wt id
```

```
fd0           0.0    0.0    0.0    0.0  0.0  0.0    0.0   0   0    0    0    0   0   0 100
sd0           0.0    0.0    0.4    0.4  0.0  0.0   49.5   0   0
sd6           0.0    0.0    0.0    0.0  0.0  0.0    0.0   0   0
nfs1          0.0    0.0    0.0    0.0  0.0  0.0    0.0   0   0
nfs49         0.0    0.0    0.0    0.0  0.0  0.0   15.1   0   0
nfs53         0.0    0.0    0.4    0.0  0.0  0.0   24.5   0   0
nfs54         0.0    0.0    0.0    0.0  0.0  0.0    6.3   0   0
nfs55         0.0    0.0    0.0    0.0  0.0  0.0    4.9   0   0
```

The `iostat -xtc` command displays a line of output for each disk. The output fields are described in the following list.

`r/s` Reads per second

`w/s` Writes per second

`kr/s` Kbytes read per second

`kw/s` Kbytes written per second

`wait` Average number of transactions that are waiting for service (queue length)

`actv` Average number of transactions that are actively being serviced

`svc_t` Average service time, in milliseconds

`%w` Percentage of time that the queue is not empty

`%b` Percentage of time that the disk is busy

Displaying Disk Space Statistics (`df`)

Use the `df` command to show the amount of free disk space on each mounted disk. The *usable* disk space that is reported by `df` reflects only 90 percent of full capacity, as the reporting statistics allows for 10 percent above the total available space. This *head room* normally stays empty for better performance.

The percentage of disk space actually reported by the `df` command is used space divided by usable space.

If the file system exceeds 90 percent capacity, you could transfer files to a disk that is not as full by using the `cp` command. Alternately, you could transfer files to a tape by using the `tar` or `cpio` commands. Or, you could remove the files.

For a detailed description of this command, see the `df`(1M) man page.

▼ How to Display Disk Space Information (`df -k`)

Step ● **Use the `df -k` command to display disk space information in Kbytes.**

```
$ df -k
Filesystem            kbytes    used   avail capacity  Mounted on
/dev/dsk/c0t3d0s0     192807   40231  133296    24%    /
```

Example 20–4 Displaying File System Information

The following example shows the output from the `df -k` command.

```
$ df -k
Filesystem            kbytes    used   avail capacity  Mounted on
/dev/dsk/c0t0d0s0     384120  131596  214112    39%    /
/dev/dsk/c0t0d0s6    1388419 1050390  282493    79%    /usr
/devices                   0       0       0     0%    /devices
/proc                      0       0       0     0%    /proc
mnttab                     0       0       0     0%    /etc/mnttab
fd                         0       0       0     0%    /dev/fd
swap                  467152      40  467112     1%    /var/run
swap                  467160      48  467112     1%    /tmp
/dev/dsk/c0t0d0s4    1784644 1525360  205745    89%    /export
venus:/usr/dist     20612581 13237316 6963015    66%    /usr/dist
```

The following table describes the output of the `df -k` command.

Field Name	Description
`kbytes`	Total size of usable space in the file system
`used`	Amount of space used
`avail`	Amount of space available for use
`capacity`	Amount of space used, as a percentage of the total capacity
`mounted on`	Mount point

Monitoring System Activities (Task Map)

Task	Description	For Instructions
Check file access.	Display file access operation status by using the `sar` command with the `-a` option.	"How to Check File Access (`sar -a`)" on page 362
Check buffer activity.	Display buffer activity statistics by using the `sar` command with the `-b` option.	"How to Check Buffer Activity (`sar -b`)" on page 363
Check system call statistics.	Display system call statistics by using the `sar` command with the `-c` option.	"How to Check System Call Statistics (`sar -c`)" on page 364
Check disk activity.	Check disk activity by using the `sar` command with the `-d` option.	"How to Check Disk Activity (`sar -d`)" on page 365
Check page-out and memory.	Use the `sar` command with the `-g` option to display page-out memory freeing activities.	"How to Check Page-Out and Memory (`sar -g`)" on page 367
Check kernel memory allocation.	The kernel memory allocation (KMA) allows a kernel subsystem to allocate and free memory, as needed. Use the `sar` command with the `-k` option to check KMA.	"How to Check Kernel Memory Allocation (`sar -k`)" on page 369
Check interprocess communication.	Use the `sar` command with the `-m` option to report interprocess communication activities.	"How to Check Interprocess Communication (`sar -m`)" on page 371
Check page-in activity.	Use the `sar` command with the `-p` option to report page-in activity.	"How to Check Page-In Activity (`sar -p`)" on page 371
Check queue activity.	Use the `sar` command with the `-q` option to check the following: ■ Average queue length while queue is occupied ■ Percentage of time that the queue is occupied	"How to Check Queue Activity (`sar -q`)" on page 373

Task	Description	For Instructions
Check unused memory.	Use the sar command with the -r option to report the number of memory pages and swap file disk blocks that are currently used.	"How to Check Unused Memory (sar -r)" on page 374
Check CPU utilization.	Use the sar command with the -v option to display CPU utilization statistics.	"How to Check CPU Utilization (sar -u)" on page 375
Check system table status.	Use the sar command with the -v option to report status on the following system tables: ■ Process ■ Inode ■ File ■ Shared memory record	"How to Check System Table Status (sar -v)" on page 376
Check swapping activity.	Use the sar command with the -w option to check swapping activity.	"How to Check Swapping Activity (sar -w)" on page 377
Check terminal activity.	Use the sar command with the -y option to monitor terminal device activity.	"How to Check Terminal Activity (sar -y)" on page 378
Check overall system performance.	The sar -A command displays statistics from all options to provide overall system performance information.	"How to Check Overall System Performance (sar -A)" on page 380
Set up automatic data collection.	To set up your system to collect data automatically and to run the sar commands, do the following: ■ Run the svcadm enable system/sar:default command ■ Edit the /var/spool/cron/crontabs/sys file	"How to Set Up Automatic Data Collection" on page 383

Monitoring System Activities (sar)

Use the sar command to perform the following tasks:

- Organize and view data about system activity.
- Access system activity data on a special request basis.

- Generate automatic reports to measure and monitor system performance, as well as special request reports to pinpoint specific performance problems. For information on how to set up the sar command to run on your system, as well as a description of these tools, see "Collecting System Activity Data Automatically (sar)" on page 380.

For a detailed description of this command, see the sar(1) man page.

▼ How to Check File Access (sar -a)

Step ● **Display file access operation statistics with the sar -a command.**

```
$ sar -a

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004

00:00:00  iget/s namei/s dirbk/s
01:00:00       0       3       0
02:00:00       0       3       0
03:00:00       0       3       0
04:00:00       0       3       0
05:00:00       0       3       0
06:00:00       0       3       0
07:00:00       0       3       0
08:00:00       0       3       0
08:20:01       0       3       0
08:40:00       0       3       0
09:00:00       0       3       0
09:20:01       0      10       0
09:40:01       0       1       0
10:00:02       0       5       0

Average        0       4       0
```

The following list describes the field names and description of operating system routines that are reported by the sar -a command.

iget/s
: The number of requests made for inodes that were not in the directory name look-up cache (DNLC).

namei/s
: The number of file system path searches per second. If namei does not find a directory name in the DNLC, it calls iget to get the inode for either a file or directory. Hence, most igets are the result of DNLC misses.

dirbk/s
: The number of directory block reads issued per second.

The larger the reported values for these operating system routines, the more time the kernel is spending to access user files. The amount of time reflects how heavily programs and applications are using the file systems. The -a option is helpful for viewing how disk-dependent an application is.

▼ How to Check Buffer Activity (`sar -b`)

Step ● **Display buffer activity statistics with the `sar -b` command.**

The buffer is used to cache metadata. Metadata includes inodes, cylinder group blocks, and indirect blocks.

```
$ sar -b
00:00:00 bread/s lread/s %rcache bwrit/s lwrit/s %wcache pread/s pwrit/s
01:00:00       0       0     100       0       0      55       0       0
```

Example 20–5 Checking Buffer Activity (`sar -b`)

The following example of `sar -b` command output shows that the `%rcache` and `%wcache` buffers are not causing any slowdowns. All the data is within acceptable limits.

```
$ sar -b

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004

00:00:04 bread/s lread/s %rcache bwrit/s lwrit/s %wcache pread/s pwrit/s
01:00:00       0       0     100       0       0      94       0       0
02:00:01       0       0     100       0       0      94       0       0
03:00:00       0       0     100       0       0      92       0       0
04:00:00       0       1     100       0       1      94       0       0
05:00:00       0       0     100       0       0      93       0       0
06:00:00       0       0     100       0       0      93       0       0
07:00:00       0       0     100       0       0      93       0       0
08:00:00       0       0     100       0       0      93       0       0
08:20:00       0       1     100       0       1      94       0       0
08:40:01       0       1     100       0       1      93       0       0
09:00:00       0       1     100       0       1      93       0       0
09:20:00       0       1     100       0       1      93       0       0
09:40:00       0       2     100       0       1      89       0       0
10:00:00       0       9     100       0       5      92       0       0
10:20:00       0       0     100       0       0      68       0       0
10:40:00       0       1      98       0       1      70       0       0
11:00:00       0       1     100       0       1      75       0       0

Average        0       1     100       0       1      91       0       0
```

The following table describes the buffer activities that are displayed by the `-b` option.

Field Name	Description
`bread/s`	Average number of reads per second that are submitted to the buffer cache from the disk

Field Name	Description
`lread/s`	Average number of logical reads per second from the buffer cache
`%rcache`	Fraction of logical reads that are found in the buffer cache (100 % minus the ratio of `bread/s` to `lread/s`)
`bwrit/s`	Average number of physical blocks (512 blocks) that are written from the buffer cache to disk, per second
`lwrit/s`	Average number of logical writes to the buffer cache, per second
`%wcache`	Fraction of logical writes that are found in the buffer cache (100 % minus the ratio of `bwrit/s` to `lwrit/s`)
`pread/s`	Average number of physical reads, per second, that use character device interfaces
`pwrit/s`	Average number of physical write requests, per second, that use character device interfaces

The most important entries are the cache hit ratios `%rcache` and `%wcache`. These entries measure the effectiveness of system buffering. If `%rcache` falls below 90 percent, or if `%wcache` falls below 65 percent, it might be possible to improve performance by increasing the buffer space.

▼ How to Check System Call Statistics (`sar -c`)

Step ● **Display system call statistics by using the `sar -c` command.**

```
$ sar -c
00:00:00 scall/s sread/s swrit/s  fork/s  exec/s rchar/s wchar/s
01:00:00      38       2       2    0.00    0.00     149     120
```

Example 20–6 Checking System Call Statistics (`sar -c`)

The following example shows output from the `sar -c` command.

```
$ sar -c

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004

00:00:04 scall/s sread/s swrit/s  fork/s  exec/s rchar/s wchar/s
01:00:00      89      14       9    0.01    0.00    2906    2394
```

```
02:00:01      89      14       9    0.01    0.00    2905    2393
03:00:00      89      14       9    0.01    0.00    2908    2393
04:00:00      90      14       9    0.01    0.00    2912    2393
05:00:00      89      14       9    0.01    0.00    2905    2393
06:00:00      89      14       9    0.01    0.00    2905    2393
07:00:00      89      14       9    0.01    0.00    2905    2393
08:00:00      89      14       9    0.01    0.00    2906    2393
08:20:00      90      14       9    0.01    0.01    2914    2395
08:40:01      90      14       9    0.01    0.00    2914    2396
09:00:00      90      14       9    0.01    0.01    2915    2396
09:20:00      90      14       9    0.01    0.01    2915    2396
09:40:00     880     207     156    0.08    0.08   26671    9290
10:00:00    2020     530     322    0.14    0.13   57675   36393
10:20:00     853     129      75    0.02    0.01   10500    8594
10:40:00    2061     524     450    0.08    0.08  579217  567072
11:00:00    1658     404     350    0.07    0.06 1152916 1144203

Average      302      66      49    0.02    0.01   57842   55544
```

The following table describes the system call categories that are reported by the `-c` option. Typically, reads and writes account for about half of the total system calls. However, the percentage varies greatly with the activities that are being performed by the system.

Field Name	Description
`scall/s`	The number of all types of system calls per second, which is generally about 30 per second on a system with 4 to 6 users.
`sread/s`	The number of `read` system calls per second.
`swrit/s`	The number of `write` system calls per second.
`fork/s`	The number of `fork` system calls per second, which is about 0.5 per second on a system with 4 to 6 users. This number increases if shell scripts are running.
`exec/s`	The number of `exec` system calls per second. If `exec/s` divided by `fork/s` is greater than 3, look for inefficient `PATH` variables.
`rchar/s`	The number of characters (bytes) transferred by `read` system calls per second.
`wchar/s`	The number of characters (bytes) transferred by `write` system calls per second.

▼ How to Check Disk Activity (`sar -d`)

Step ● **Display disk activity statistics with the `sar -d` command.**

```
$ sar -d

00:00:00   device        %busy   avque   r+w/s  blks/s  avwait  avserv
```

Example 20–7 Checking Disk Activity

This abbreviated example illustrates the output from the `sar -d` command.

```
$ sar -d

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004

12:36:32   device        %busy   avque   r+w/s  blks/s  avwait  avserv

12:40:01   dad1             15     0.7      26     399    18.1    10.0
           dad1,a           15     0.7      26     398    18.1    10.0
           dad1,b            0     0.0       0       1     1.0     3.0
           dad1,c            0     0.0       0       0     0.0     0.0
           dad1,h            0     0.0       0       0     0.0     6.0
           fd0               0     0.0       0       0     0.0     0.0
           nfs1              0     0.0       0       0     0.0     0.0
           nfs2              1     0.0       1      12     0.0    13.2
           nfs3              0     0.0       0       2     0.0     1.9
           nfs4              0     0.0       0       0     0.0     7.0
           nfs5              0     0.0       0       0     0.0    57.1
           nfs6              1     0.0       6     125     4.3     3.2
           nfs7              0     0.0       0       0     0.0     6.0
           sd1               0     0.0       0       0     0.0     5.4
           ohci0,bu          0     0.0       0       0     0.0     0.0
           ohci0,ct          0     0.0       0       0     0.0     0.0
           ohci0,in          0     0.0       7       0     0.0     0.0
           ohci0,is          0     0.0       0       0     0.0     0.0
           ohci0,to          0     0.0       7       0     0.0     0.0
```

The following table describes the disk device activities that are reported by the `-d` option.

Field Name	Description
`device`	Name of the disk device that is being monitored.
`%busy`	Percentage of time the device spent servicing a transfer request.
`avque`	The sum of the average wait time plus the average service time.
`r+w/s`	Number of read-and-write transfers to the device, per second.

Field Name	Description
blks/s	Number of 512-byte blocks that are transferred to the device, per second.
avwait	Average time, in milliseconds, that transfer requests wait idly in the queue. This time is measured only when the queue is occupied.
avserv	Average time, in milliseconds, for a transfer request to be completed by the device. For disks, this value includes seek times, rotational latency times, and data transfer times.

Note that queue lengths and wait times are measured when something is in the queue. If %busy is small, large queues and service times probably represent the periodic efforts by the system to ensure that altered blocks are promptly written to the disk.

▼ How to Check Page-Out and Memory (sar -g)

Step ● **Use the sar -g command to display page-out and memory freeing activities in averages.**

```
$ sar -g
00:00:00  pgout/s ppgout/s pgfree/s pgscan/s %ufs_ipf
01:00:00     0.00     0.00     0.00     0.00     0.00
```

The output displayed by the sar -g command is a good indicator of whether more memory might be needed. Use the ps -elf command to show the number of cycles that are used by the page daemon. A high number of cycles, combined with high values for the pgfree/s and pgscan/s fields, indicates a memory shortage.

The sar -g command also shows whether inodes are being recycled too quickly and causing a loss of reusable pages.

Example 20–8 Checking Page-Out and Memory (sar -g)

The following example shows output from the sar -g command.

```
$ sar -g

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004

00:00:00  pgout/s ppgout/s pgfree/s pgscan/s %ufs_ipf
01:00:00     0.00     0.00     0.00     0.00     0.00
02:00:00     0.01     0.01     0.01     0.00     0.00
03:00:00     0.00     0.00     0.00     0.00     0.00
04:00:00     0.00     0.00     0.00     0.00     0.00
05:00:00     0.00     0.00     0.00     0.00     0.00
```

```
06:00:00     0.00     0.00     0.00      0.00     0.00
07:00:00     0.00     0.00     0.00      0.00     0.00
08:00:00     0.00     0.00     0.00      0.00     0.00
08:20:01     0.00     0.00     0.00      0.00     0.00
08:40:00     0.00     0.00     0.00      0.00     0.00
09:00:00     0.00     0.00     0.00      0.00     0.00
09:20:01     0.05     0.52     1.62     10.16     0.00
09:40:01     0.03     0.44     1.47      4.77     0.00
10:00:02     0.13     2.00     4.38     12.28     0.00
10:20:03     0.37     4.68    12.26     33.80     0.00

Average      0.02     0.25     0.64      1.97     0.00
```

The following table describes the output from the `-g` option.

Field Name	Description
`pgout/s`	The number of page-out requests per second.
`ppgout/s`	The actual number of pages that are paged-out, per second. A single page-out request might involve paging-out multiple pages.
`pgfree/s`	The number of pages, per second, that are placed on the free list.
`pgscan/s`	The number of pages, per second, that are scanned by the `page` daemon. If this value is high, the `page` daemon is spending a lot of time checking for free memory. This situation implies that more memory might be needed.
`%ufs_ipf`	The percentage of `ufs` inodes taken off the free list by `iget` that had reusable pages associated with them. These pages are flushed and cannot be reclaimed by processes. Thus, this field represents the percentage of `igets` with page flushes. A high value indicates that the free list of inodes is page-bound, and that the number of `ufs` inodes might need to be increased.

Checking Kernel Memory Allocation

The KMA allows a kernel subsystem to allocate and free memory, as needed. Rather than statically allocating the maximum amount of memory it is expected to require under peak load, the KMA divides requests for memory into three categories:

- Small (less than 256 bytes)
- Large (512 bytes to 4 Kbytes)

- Oversized (greater than 4 Kbytes)

The KMA keeps two pools of memory to satisfy small requests and large requests. The oversized requests are satisfied by allocating memory from the system page allocator.

If you are checking a system that is being used to write drivers or STREAMS that use KMA resources, then the `sar -k` command will likely prove useful. Otherwise, you will probably not need the information it provides. Any driver or module that uses KMA resources, but does not specifically return the resources before it exits, can create a memory leak. A memory leak causes the amount of memory that is allocated by KMA to increase over time. Thus, if the `alloc` fields of the `sar -k` command increase steadily over time, there might be a memory leak. Another indication of a memory leak is failed requests. If this problem occurs, a memory leak has probably caused KMA to be unable to reserve and allocate memory.

If it appears that a memory leak has occurred, you should check any drivers or STREAMS that might have requested memory from KMA and not returned it.

▼ How to Check Kernel Memory Allocation (`sar -k`)

Step ● **Use the `sar -k` command to report on the following activities of the Kernel Memory Allocator (KMA).**

```
$ sar -k
00:00:00 sml_mem   alloc  fail  lg_mem   alloc   fail  ovsz_alloc  fail
01:00:00 2523136 1866512     0 18939904 14762364    0      360448     0
02:00:02 2523136 1861724     0 18939904 14778748    0      360448     0
```

Example 20–9 Checking Kernel Memory Allocation (`sar -k`)

The following is an abbreviated example of `sar -k` output.

```
$ sar -k

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004

00:00:04 sml_mem   alloc  fail  lg_mem   alloc  fail  ovsz_alloc  fail
01:00:00 6119744 4852865     0 60243968 54334808   156     9666560     0
02:00:01 6119744 4853057     0 60243968 54336088   156     9666560     0
03:00:00 6119744 4853297     0 60243968 54335760   156     9666560     0
04:00:00 6119744 4857673     0 60252160 54375280   156     9666560     0
05:00:00 6119744 4858097     0 60252160 54376240   156     9666560     0
06:00:00 6119744 4858289     0 60252160 54375608   156     9666560     0
07:00:00 6119744 4858793     0 60252160 54442424   156     9666560     0
08:00:00 6119744 4858985     0 60252160 54474552   156     9666560     0
08:20:00 6119744 4858169     0 60252160 54377400   156     9666560     0
08:40:01 6119744 4857345     0 60252160 54376880   156     9666560     0
```

```
09:00:00 6119744 4859433     0 60252160 54539752   156     9666560     0
09:20:00 6119744 4858633     0 60252160 54410920   156     9666560     0
09:40:00 6127936 5262064     0 60530688 55619816   156     9666560     0
10:00:00 6545728 5823137     0 62996480 58391136   156     9666560     0
10:20:00 6545728 5758997     0 62996480 57907400   156     9666560     0
10:40:00 6734144 6035759     0 64389120 59743064   156    10493952     0
11:00:00 6996288 6394872     0 65437696 60935936   156    10493952     0

Average  6258044 5150556     0 61138340 55609004   156     9763900     0
```

The following table describes the output from the `-k` option.

Field Name	Description
`sml_mem`	The amount of memory, in bytes, that the KMA has available in the small memory request pool. In this pool, here a small request is less than 256 bytes.
`alloc`	The amount of memory, in bytes, that the KMA has allocated from its small memory request pool to small memory requests.
`fail`	The number of requests for small amounts of memory that failed.
`lg_mem`	The amount of memory, in bytes, that the KMA has available in the large memory request pool. In this pool, a large request is from 512 bytes to 4 Kbytes.
`alloc`	The amount of memory, in bytes, that the KMA has allocated from its large memory request pool to large memory requests.
`fail`	The number of failed requests for large amounts of memory.
`ovsz_alloc`	The amount of memory that is allocated for oversized requests, which are requests that are greater than 4 Kbytes. These requests are satisfied by the page allocator. Thus, there is no pool.
`fail`	The number of failed requests for oversized amounts of memory.

▼ How to Check Interprocess Communication (`sar -m`)

Step ● **Use the `sar -m` command to report interprocess communication activities.**

```
$ sar -m
00:00:00   msg/s  sema/s
01:00:00    0.00    0.00
```

These figures are usually zero (0.00), unless you are running applications that use messages or semaphores.

The following list describes the output from the `-m` option.

`msg/s` The number of message operations (sends and receives) per second

`sema/s` The number of semaphore operations per second

Example 20–10 Checking Interprocess Communication (`sar -m`)

The following abbreviated example shows output from the `sar -m` command.

```
$ sar -m

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004

00:00:00   msg/s  sema/s
01:00:00    0.00    0.00
02:00:02    0.00    0.00
03:00:00    0.00    0.00
04:00:00    0.00    0.00
05:00:01    0.00    0.00
06:00:00    0.00    0.00

Average     0.00    0.00
```

▼ How to Check Page-In Activity (`sar -p`)

Step ● **Use the `sar -p` command to report page-in activity, which includes protection and translation faults.**

```
$ sar -p
00:00:00  atch/s  pgin/s ppgin/s  pflt/s  vflt/s slock/s
01:00:00    0.07    0.00    0.00    0.21    0.39    0.00
```

Example 20–11 Checking Page-In Activity (`sar -p`)

The following example shows output from the `sar -p` command.

```
$ sar -p

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004

00:00:04  atch/s  pgin/s ppgin/s  pflt/s  vflt/s slock/s
01:00:00    0.09    0.00    0.00    0.78    2.02    0.00
02:00:01    0.08    0.00    0.00    0.78    2.02    0.00
03:00:00    0.09    0.00    0.00    0.81    2.07    0.00
04:00:00    0.11    0.01    0.01    0.86    2.18    0.00
05:00:00    0.08    0.00    0.00    0.78    2.02    0.00
06:00:00    0.09    0.00    0.00    0.78    2.02    0.00
07:00:00    0.08    0.00    0.00    0.78    2.02    0.00
08:00:00    0.09    0.00    0.00    0.78    2.02    0.00
08:20:00    0.11    0.00    0.00    0.87    2.24    0.00
08:40:01    0.13    0.00    0.00    0.90    2.29    0.00
09:00:00    0.11    0.00    0.00    0.88    2.24    0.00
09:20:00    0.10    0.00    0.00    0.88    2.24    0.00
09:40:00    2.91    1.80    2.38    4.61   17.62    0.00
10:00:00    2.74    2.03    3.08    8.17   21.76    0.00
10:20:00    0.16    0.04    0.04    1.92    2.96    0.00
10:40:00    2.10    2.50    3.42    6.62   16.51    0.00
11:00:00    3.36    0.87    1.35    3.92   15.12    0.00

Average     0.42    0.22    0.31    1.45    4.00    0.00
```

The following table describes the reported statistics from the `-p` option.

Field Name	Description
`atch/s`	The number of page faults, per second, that are satisfied by reclaiming a page currently in memory (attaches per second). Instances include reclaiming an invalid page from the free list and sharing a page of text that is currently being used by another process. An example is two or more processes that are accessing the same program text.
`pgin/s`	The number of times, per second, that file systems receive page-in requests.
`ppgin/s`	The number of pages paged in, per second. A single page-in request, such as a soft-lock request (see `slock/s`) or a large block size, might involve paging-in multiple pages.
`pflt/s`	The number of page faults from protection errors. Instances of protection faults indicate illegal access to a page and "copy-on-writes." Generally, this number consists primarily of "copy-on-writes."

Field Name	Description
vflt/s	The number of address translation page faults, per second. These faults are known as validity faults. Validity faults occur when a valid process table entry does not exist for a given virtual address.
slock/s	The number of faults, per second, caused by software lock requests that require physical I/O. An example of the occurrence of a soft-lock request is the transfer of data from a disk to memory. The system locks the page that is to receive the data so that the page cannot be claimed and used by another process.

▼ How to Check Queue Activity (sar -q)

Step ● **Use the sar -q command to report the following information:**

- The Average queue length while the queue is occupied.
- The percentage of time that the queue is occupied.

```
$ sar -q
00:00:00 runq-sz %runocc swpq-sz %swpocc
```

The following list describes the output from the -q option.

runq-sz — The number of kernel threads in memory that are waiting for a CPU to run. Typically, this value should be less than 2. Consistently higher values mean that the system might be CPU-bound.

%runocc — The percentage of time that the dispatch queues are occupied.

swpq-sz — No longer reported by the sar command.

%swpocc — No longer reported by the sar command.

Example 20–12 Checking Queue Activity

The following example shows output from the sar -q command. If the %runocc value is high (greater than 90 percent) and the runq-sz value is greater than 2, the CPU is heavily loaded and response is degraded. In this case, additional CPU capacity might be required to obtain acceptable system response.

```
$ sar -q

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004

00:00:04 runq-sz %runocc swpq-sz %swpocc
```

```
01:00:00     1.0       0     0.0       0
02:00:01     1.3       0     0.0       0
03:00:00     1.0       0     0.0       0
04:00:00     1.0       0     0.0       0
05:00:00     1.0       0     0.0       0
06:00:00     2.0       0     0.0       0
07:00:00     0.0       0     0.0       0
08:00:00     1.0       0     0.0       0
08:20:00     1.0       0     0.0       0
08:40:01     2.0       0     0.0       0
09:00:00     0.0       0     0.0       0
09:20:00     1.0       0     0.0       0
09:40:00     1.2       2     0.0       0
10:00:00     1.2       2     0.0       0
10:20:00     1.0       1     0.0       0
10:40:00     1.3       9     0.0       0
11:00:00     1.2       7     0.0       0

Average      1.2       1     0.0       0
```

▼ How to Check Unused Memory (sar -r)

Step ● **Use the `sar -r` command to report the number of memory pages and swap-file disk blocks that are currently unused.**

```
$ sar -r
00:00:00 freemem freeswap
01:00:00    2135   401922
```

The following list describes the output from the `-r` option.

`freemem` The average number of memory pages that are available to user processes over the intervals sampled by the command. Page size is machine-dependent.

`freeswap` The number of 512-byte disk blocks that are available for page swapping.

Example 20–13 Checking Unused Memory (sar -r)

The following example shows output from the `sar -r` command.

```
$ sar -r

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004

00:00:04 freemem freeswap
01:00:00   44717  1715062
02:00:01   44733  1715496
03:00:00   44715  1714746
```

```
04:00:00   44751  1715403
05:00:00   44784  1714743
06:00:00   44794  1715186
07:00:00   44793  1715159
08:00:00   44786  1714914
08:20:00   44805  1715576
08:40:01   44797  1715347
09:00:00   44761  1713948
09:20:00   44802  1715478
09:40:00   41770  1682239
10:00:00   35401  1610833
10:20:00   34295  1599141
10:40:00   33943  1598425
11:00:00   30500  1561959

Average    43312  1699242
```

▼ How to Check CPU Utilization (`sar -u`)

Step ● **Use the `sar -u` command to display CPU utilization statistics.**

```
$ sar -u
00:00:00    %usr    %sys    %wio   %idle
01:00:00       0       0       0     100
```

The `sar` command without any options is equivalent to the `sar -u` command. At any given moment, the processor is either busy or idle. When busy, the processor is in either user mode or system mode. When idle, the processor is either waiting for I/O completion or "sitting still" with no work to do.

The following list describes output from the `-u` option.

`%usr` Lists the percentage of time that the processor is in user mode

`%sys` Lists the percentage of time that the processor is in system mode

`%wio` Lists the percentage of time that the processor is idle and waiting for I/O completion

`%idle` Lists the percentage of time that the processor is idle and not waiting for I/O

A high `%wio` value generally means that a disk slowdown has occurred.

Example 20–14 Checking CPU Utilization (`sar -u`)

The following example shows output from the `sar -u` command.

```
$ sar -u

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004
```

```
00:00:04    %usr    %sys    %wio   %idle
01:00:00       0       0       0     100
02:00:01       0       0       0     100
03:00:00       0       0       0     100
04:00:00       0       0       0     100
05:00:00       0       0       0     100
06:00:00       0       0       0     100
07:00:00       0       0       0     100
08:00:00       0       0       0     100
08:20:00       0       0       0      99
08:40:01       0       0       0      99
09:00:00       0       0       0      99
09:20:00       0       0       0      99
09:40:00       4       1       0      95
10:00:00       4       2       0      94
10:20:00       1       1       0      98
10:40:00      18       3       0      79
11:00:00      25       3       0      72

Average        2       0       0      98
```

▼ How to Check System Table Status (`sar -v`)

Step ● **Use the `sar -v` command to report the status of the process table, inode table, file table, and shared memory record table.**

```
$ sar -v
00:00:00  proc-sz    ov  inod-sz    ov  file-sz    ov   lock-sz
01:00:00   43/922     0 2984/4236    0  322/322     0    0/0
```

Example 20–15 Checking System Table Status (`sar -v`)

The following abbreviated example shows output from the `sar -v` command. This example shows that all tables are large enough to have no overflows. These tables are all dynamically allocated based on the amount of physical memory.

```
$ sar -v

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004

00:00:04  proc-sz    ov  inod-sz    ov  file-sz    ov   lock-sz
01:00:00   69/8010    0 3476/34703    0    0/0       0    0/0
02:00:01   69/8010    0 3476/34703    0    0/0       0    0/0
03:00:00   69/8010    0 3476/34703    0    0/0       0    0/0
04:00:00   69/8010    0 3494/34703    0    0/0       0    0/0
05:00:00   69/8010    0 3494/34703    0    0/0       0    0/0
06:00:00   69/8010    0 3494/34703    0    0/0       0    0/0
07:00:00   69/8010    0 3494/34703    0    0/0       0    0/0
```

```
08:00:00   69/8010    0 3494/34703    0    0/0        0    0/0
08:20:00   69/8010    0 3494/34703    0    0/0        0    0/0
08:40:01   69/8010    0 3494/34703    0    0/0        0    0/0
09:00:00   69/8010    0 3494/34703    0    0/0        0    0/0
09:20:00   69/8010    0 3494/34703    0    0/0        0    0/0
09:40:00   74/8010    0 3494/34703    0    0/0        0    0/0
10:00:00   75/8010    0 4918/34703    0    0/0        0    0/0
10:20:00   72/8010    0 4918/34703    0    0/0        0    0/0
10:40:00   71/8010    0 5018/34703    0    0/0        0    0/0
11:00:00   77/8010    0 5018/34703    0    0/0        0    0/0
```

Output from the -v option is described in the following table.

Field Name	Description
proc-sz	The number of process entries (proc structures) that are currently being used, or allocated, in the kernel.
inod-sz	The total number of inodes in memory compared to the maximum number of inodes that are allocated in the kernel. This number is not a strict high watermark. The number can overflow.
file-sz	The size of the open system file table. The sz is given as 0, because space is allocated dynamically for the file table.
ov	The overflows that occur between sampling points for each table.
lock-sz	The number of shared memory record table entries that are currently being used, or allocated, in the kernel. The sz is given as 0 because space is allocated dynamically for the shared memory record table.

▼ How to Check Swapping Activity (sar -w)

Step ● **Use the sar -w command to report swapping and switching activity.**

```
$ sar -w
00:00:00 swpin/s bswin/s swpot/s bswot/s pswch/s
01:00:00    0.00     0.0    0.00     0.0      22
```

The following list describes target values and observations related to the sar -w command output.

swpin/s The number of LWP transfers into memory per second.

`bswin/s`	The number of blocks transferred for swap-ins per second. `/* (float)PGTOBLK(xx->cvmi.pgswapin) / sec_diff */`
`swpot/s`	The average number of processes that are swapped out of memory per second. If the number is greater than 1, you might need to increase memory.
`bswot/s`	The number of blocks that are transferred for swap-outs per second.
`pswch/s`	The number of kernel thread switches, per second.

Note – All process swap-ins include process initialization.

Example 20–16 Checking Swap Activity (`sar -w`)

The following example shows output from the `sar -w` command.

```
$ sar -w

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004

00:00:04 swpin/s bswin/s swpot/s bswot/s pswch/s
01:00:00    0.00     0.0    0.00     0.0     132
02:00:01    0.00     0.0    0.00     0.0     133
03:00:00    0.00     0.0    0.00     0.0     133
04:00:00    0.00     0.0    0.00     0.0     134
05:00:00    0.00     0.0    0.00     0.0     133
06:00:00    0.00     0.0    0.00     0.0     133
07:00:00    0.00     0.0    0.00     0.0     132
08:00:00    0.00     0.0    0.00     0.0     131
08:20:00    0.00     0.0    0.00     0.0     133
08:40:01    0.00     0.0    0.00     0.0     132
09:00:00    0.00     0.0    0.00     0.0     132
09:20:00    0.00     0.0    0.00     0.0     132
09:40:00    0.00     0.0    0.00     0.0     335
10:00:00    0.00     0.0    0.00     0.0     601
10:20:00    0.00     0.0    0.00     0.0     353
10:40:00    0.00     0.0    0.00     0.0     747
11:00:00    0.00     0.0    0.00     0.0     804

Average     0.00     0.0    0.00     0.0     198
```

▼ How to Check Terminal Activity (`sar -y`)

Step ● **Use the `sar -y` command to monitor terminal device activities.**

```
$ sar -y
00:00:00 rawch/s canch/s outch/s rcvin/s xmtin/s mdmin/s
01:00:00       0       0       0       0       0       0
```

If you have a lot of terminal I/O, you can use this report to determine if any bad lines exist. The activities recorded are defined in the following list.

`rawch/s` Input characters (raw queue) per second

`canch/s` Input characters that are processed by canon (canonical queue) per second

`outch/s` Output characters (output queue) per second

`rcvin/s` Receiver hardware interrupts per second

`xmtin/s` Transmitter hardware interrupts per second

`mdmin/s` Modem interrupts per second

The number of modem interrupts per second (`mdmin/s`) should be close to zero. The receive and transmit interrupts per second (`xmtin/s` and `rcvin/s`) should be less than or equal to the number of incoming or outgoing characters, respectively. If not, check for bad lines.

Example 20–17 Checking Terminal Activity (`sar -y`)

The following example shows output from the `sar -y` command.

```
$ sar -y

SunOS balmyday 5.10 s10_51 sun4u    03/18/2004

00:00:04 rawch/s canch/s outch/s rcvin/s xmtin/s mdmin/s
01:00:00       0       0       0       0       0       0
02:00:01       0       0       0       0       0       0
03:00:00       0       0       0       0       0       0
04:00:00       0       0       0       0       0       0
05:00:00       0       0       0       0       0       0
06:00:00       0       0       0       0       0       0
07:00:00       0       0       0       0       0       0
08:00:00       0       0       0       0       0       0
08:20:00       0       0       0       0       0       0
08:40:01       0       0       0       0       0       0
09:00:00       0       0       0       0       0       0
09:20:00       0       0       0       0       0       0
09:40:00       0       0       1       0       0       0
10:00:00       0       0      37       0       0       0
10:20:00       0       0       0       0       0       0
10:40:00       0       0       3       0       0       0
11:00:00       0       0       3       0       0       0

Average        0       0       1       0       0       0
```

▼ How to Check Overall System Performance (sar -A)

Step ● **Use the sar -A command to display statistics from all options to provide a view of overall system performance.**

This command provides a more global perspective. If data from more than a single time segment is shown, the report includes averages.

Collecting System Activity Data Automatically (sar)

Three commands are involved in the automatic collection of system activity data: sadc, sa1, and sa2.

The sadc data collection utility periodically collects data on system activity and saves the data in a file in binary format, one file for each 24-hour period. You can set up the sadc command to run periodically (usually once each hour), and whenever the system boots to multiuser mode. The data files are placed in the /var/adm/sa directory. Each file is named sa*dd*, where *dd* is the current date. The format of the command is as follows:

```
/usr/lib/sa/sadc [t n] [ofile]
```

The command samples *n* times with an interval of *t* seconds, which should be greater than five seconds between samples. This command then writes to the binary *ofile* file, or to standard output.

Running the sadc Command When Booting

The sadc command should be run at system boot time to record the statistics from when the counters are reset to zero. To make sure that the sadc command is run at boot time, the svcadm enable system/sar:default command writes a record to the daily data file.

The command entry has the following format:

```
/usr/bin/su sys -c "/usr/lib/sa/sadc /var/adm/sa/sa`date +%d`"
```

Running the `sadc` Command Periodically With the `sa1` Script

To generate periodic records, you need to run the `sadc` command regularly. The simplest way to do so is to uncomment the following lines in the `/var/spool/cron/crontabs/sys` file:

```
# 0 * * * 0-6 /usr/lib/sa/sa1
# 20,40 8-17 * * 1-5 /usr/lib/sa/sa1
# 5 18 * * 1-5 /usr/lib/sa/sa2 -s 8:00 -e 18:01 -i 1200 -A
```

The `sys crontab` entries do the following:

- The first two `crontab` entries cause a record to be written to the `/var/adm/sa/sa`*dd* file every 20 minutes from 8 a.m. to 5 p.m., Monday through Friday, and every hour on the hour otherwise.
- The third entry writes a record to the `/var/adm/sa/sar`*dd* file hourly, Monday through Friday, and includes all `sar` options.

You can change these defaults to meet your needs.

Producing Reports With the `sa2` Shell Script

Another shell script, `sa2`, produces reports rather than binary data files. The `sa2` command invokes the `sar` command and writes the ASCII output to a report file.

Setting Up Automatic Data Collection (`sar`)

The `sar` command can be used either to gather system activity data itself or to report what has been collected in the daily activity files that are created by the `sadc` command.

The `sar` command has the following formats:

`sar [-aAbcdgkmpqruvwy] [-o` *file*`]` *t [n]*

`sar [-aAbcdgkmpqruvwy] [-s` *time*`] [-e` *time*`] [-i` *sec*`] [-f` *file*`]`

The following `sar` command samples cumulative activity counters in the operating system every *t* seconds, *n* times. The *t* should be five seconds or greater. Otherwise, the command itself might affect the sample. You must specify a time interval in which to take the samples. Otherwise, the command operates according to the second format. The default value of *n* is 1. The following example takes two samples separated by 10 seconds. If the `-o` option were specified, samples are saved in binary format.

```
$ sar -u 10 2
```

Other important information about the `sar` command includes the following:

- With no sampling interval or number of samples specified, the `sar` command extracts data from a previously recorded file. This file is either the file specified by the `-f` option or, by default, the standard daily activity file, `/var/adm/sa/sa`*dd*, for the most recent day.
- The `-s` and `-e` options define the starting time and the ending time for the report. Starting and ending times are of the form *hh*[:*mm*[:*ss*]], where *hh*, *mm*, and *ss* represent hours, minutes, and seconds.
- The `-i` option specifies, in seconds, the intervals between record selection. If the `-i` option is not included, all intervals that are found in the daily activity file are reported.

The following table lists the `sar` options and their actions.

TABLE 20–2 Options for the `sar` Command

Option	Actions
`-a`	Checks file access operations
`-b`	Checks buffer activity
`-c`	Checks system calls
`-d`	Checks activity for each block device
`-g`	Checks page-out and memory freeing
`-k`	Checks kernel memory allocation
`-m`	Checks interprocess communication
`-nv`	Checks system table status
`-p`	Checks swap and dispatch activity
`-q`	Checks queue activity
`-r`	Checks unused memory
`-u`	Checks CPU utilization
`-w`	Checks swapping and switching volume
`-y`	Checks terminal activity
`-A`	Reports overall system performance, which is the same as entering all options.

Using no option is equivalent to calling the `sar` command with the `-u` option.

▼ How to Set Up Automatic Data Collection

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2. **Run the `svcadm enable system/sar:default` command.**

   This version of the `sadc` command writes a special record that marks the time when the counters are reset to zero (boot time).

3. **Edit the `/var/spool/cron/crontabs/sys` `crontab` file and uncomment the following lines:**

   ```
   # 0 * * * 0-6 /usr/lib/sa/sa1
   # 20,40 8-17 * * 1-5 /usr/lib/sa/sa1
   # 5 18 * * 1-5 /usr/lib/sa/sa2 -s 8:00 -e 18:01 -i 1200 -A
   ```

CHAPTER 21

Troubleshooting Software Problems (Overview)

This chapter provides a general overview of troubleshooting software problems, including information on troubleshooting system crashes and viewing system messages.

This is a list of information in this chapter.

What's New in Troubleshooting?

Dynamic Tracing Facility

The Solaris DTrace facility is a comprehensive dynamic tracking facility that gives you a new level of observerability into the Solaris kernel and user processes. DTrace helps you understand your system by permitting you to dynamically instrument the OS kernel and user processes to record data that you specify at locations of interest, called, *probes*. Each probe can be associated with custom programs that are written in the new D programming language. All of DTrace's instrumentation is entirely dynamic and available for use on your production system. For more information, see the `dtrace`(1M) man page and the *Solaris Dynamic Tracing Guide*.

kmdb Replaces kadb as Standard Solaris Kernel Debugger

In this release, kmdb has replaced kadb as the standard "in situ" Solaris kernel debugger.

kmdb brings all the power and flexibility of mdb to live kernel debugging. kmdb supports the following:

- Debugger commands (dcmds)
- Debugger modules (dmods)
- Access to kernel type data
- Kernel execution control
- Inspection
- Modification

For more information, see the kmdb(1) man page. For step-by-step instructions on using kmdb to troubleshoot a system, see the "How to Boot the System With the Kernel Debugger (kmdb)" in *System Administration Guide: Basic Administration*.

Where to Find Software Troubleshooting Tasks

Troubleshooting Task	For More Information
Manage system crash information	Chapter 24
Manage core files	Chapter 23
Troubleshoot software problems such as reboot failures and backup problems	Chapter 25
Troubleshoot file access problems	Chapter 26
Troubleshoot printing problems	Chapter 27
Resolve UFS file system inconsistencies	Chapter 28
Troubleshoot software package problems	Chapter 29

Troubleshooting a System Crash

If a system running the Solaris Operating System crashes, provide your service provider with as much information as possible, including crash dump files.

What to Do if the System Crashes

The most important things to remember are:

1. Write down the system console messages.

   If a system crashes, making it run again might seem like your most pressing concern. However, before you reboot the system, examine the console screen for messages. These messages can provide some insight about what caused the crash. Even if the system reboots automatically and the console messages have disappeared from the screen, you might be able to check these messages by viewing the system error log, the `/var/adm/messages` file. For more information about viewing system error log files, see "How to View System Messages" on page 392.

   If you have frequent crashes and can't determine their cause, gather all the information you can from the system console or the `/var/adm/messages` files, and have it ready for a customer service representative to examine. For a complete list of troubleshooting information to gather for your service provider, see "Troubleshooting a System Crash" on page 387.

   If the system fails to reboot successfully after a system crash, see Chapter 25.

2. Synchronize the disks and reboot.

   ```
   ok sync
   ```

   If the system fails to reboot successfully after a system crash, see Chapter 25.

Check to see if a system crash dump was generated after the system crash. System crash dumps are saved by default. For information about crash dumps, see Chapter 24.

Gathering Troubleshooting Data

Answer the following questions to help isolate the system problem. Use "Troubleshooting a System Crash Checklist" on page 388 for gathering troubleshooting data for a crashed system.

TABLE 21–1 Identifying System Crash Data

Question	Description
Can you reproduce the problem?	This is important because a reproducible test case is often essential for debugging really hard problems. By reproducing the problem, the service provider can build kernels with special instrumentation to trigger, diagnose, and fix the bug.
Are you using any third-party drivers?	Drivers run in the same address space as the kernel, with all the same privileges, so they can cause system crashes if they have bugs.
What was the system doing just before it crashed?	If the system was doing anything unusual like running a new stress test or experiencing higher-than-usual load, that might have led to the crash.
Were there any unusual console messages right before the crash?	Sometimes the system will show signs of distress before it actually crashes; this information is often useful.
Did you add any tuning parameters to the `/etc/system` *file?*	Sometimes tuning parameters, such as increasing shared memory segments so that the system tries to allocate more than it has, can cause the system to crash.
Did the problem start recently?	If so, did the onset of problems coincide with any changes to the system, for example, new drivers, new software, different workload, CPU upgrade, or a memory upgrade.

Troubleshooting a System Crash Checklist

Use this checklist when gathering system data for a crashed system.

Item	Your Data
Is a system crash dump available?	
Identify the operating system release and appropriate software application release levels.	
Identify system hardware.	
Include `prtdiag` output for sun4u systems. Include Explorer output for other systems.	

Item	Your Data
Are patches installed? If so, include `showrev -p` output.	
Is the problem reproducible?	
Does the system have any third-party drivers?	
What was the system doing before it crashed?	
Were there any unusual console messages right before the system crashed?	
Did you add any parameters to the `/etc/system` file?	
Did the problem start recently?	

CHAPTER 22

Managing System Messages

This chapter describes system messaging features in the Solaris Operating System.

Viewing System Messages

System messages display on the console device. The text of most system messages look like this:

[ID *msgid facility*.*priority*]

For example:

```
[ID 672855 kern.notice] syncing file systems...
```

If the message originated in the kernel, the kernel module name is displayed. For example:

```
Oct 1 14:07:24 mars ufs: [ID 845546 kern.notice] alloc: /: file system full
```

When a system crashes, it might display a message on the system console like this:

`panic:` *error message*

Less frequently, this message might be displayed instead of the panic message:

```
Watchdog reset !
```

The error logging daemon, `syslogd`, automatically records various system warnings and errors in message files. By default, many of these system messages are displayed on the system console and are stored in the `/var/adm` directory. You can direct where these messages are stored by setting up system message logging. For more information, see "Customizing System Message Logging" on page 394. These messages can alert you to system problems, such as a device that is about to fail.

The /var/adm directory contains several message files. The most recent messages are in /var/adm/messages file (and in messages.*), and the oldest are in the messages.3 file. After a period of time (usually every ten days), a new messages file is created. The messages.0 file is renamed messages.1, messages.1 is renamed messages.2, and messages.2 is renamed messages.3. The current /var/adm/messages.3 file is deleted.

Because the /var/adm directory stores large files containing messages, crash dumps, and other data, this directory can consume lots of disk space. To keep the /var/adm directory from growing too large, and to ensure that future crash dumps can be saved, you should remove unneeded files periodically. You can automate this task by using the crontab file. For more information on automating this task, see "How to Delete Crash Dump Files" on page 258 and Chapter 15.

▼ How to View System Messages

Step ● **Display recent messages generated by a system crash or reboot by using the dmesg command.**

```
$ dmesg
```

Or, use the more command to display one screen of messages at a time.

```
$ more /var/adm/messages
```

Example 22–1 Viewing System Messages

The following example shows output from the dmesg command.

```
$ dmesg
Jan  3 08:44:41 starbug genunix: [ID 540533 kern.notice] SunOS Release 5.10 ...
Jan  3 08:44:41 starbug genunix: [ID 913631 kern.notice] Copyright 1983-2003 ...
Jan  3 08:44:41 starbug genunix: [ID 678236 kern.info] Ethernet address ...
Jan  3 08:44:41 starbug unix: [ID 389951 kern.info] mem = 131072K (0x8000000)
Jan  3 08:44:41 starbug unix: [ID 930857 kern.info] avail mem = 121888768
Jan  3 08:44:41 starbug rootnex: [ID 466748 kern.info] root nexus = Sun Ultra 5/
10 UPA/PCI (UltraSPARC-IIi 333MHz)
Jan  3 08:44:41 starbug rootnex: [ID 349649 kern.info] pcipsy0 at root: UPA 0x1f0x0
Jan  3 08:44:41 starbug genunix: [ID 936769 kern.info] pcipsy0 is /pci@1f,0
Jan  3 08:44:41 starbug pcipsy: [ID 370704 kern.info] PCI-device: pci@1,1, simba0
Jan  3 08:44:41 starbug genunix: [ID 936769 kern.info] simba0 is /pci@1f,0/pci@1,1
Jan  3 08:44:41 starbug pcipsy: [ID 370704 kern.info] PCI-device: pci@1, simba1
Jan  3 08:44:41 starbug genunix: [ID 936769 kern.info] simba1 is /pci@1f,0/pci@1
Jan  3 08:44:57 starbug simba: [ID 370704 kern.info] PCI-device: ide@3, uata0
Jan  3 08:44:57 starbug genunix: [ID 936769 kern.info] uata0 is /pci@1f,0/pci@1,
1/ide@3
Jan  3 08:44:57 starbug uata: [ID 114370 kern.info] dad0 at pci1095,6460
.
.
.
```

See Also For more information, see the dmesg(1M) man page.

System Log Rotation

System log files are rotated by the logadm command from an entry in the root crontab file. The /usr/lib/newsyslog script is no longer used.

The system log rotation is defined in the /etc/logadm.conf file. This file includes log rotation entries for processes such as syslogd. For example, one entry in the /etc/logadm.conf file specifies that the /var/log/syslog file is rotated weekly unless the file is empty. The most recent syslog file becomes syslog.0, the next most recent becomes syslog.1, and so on. Eight previous syslog log files are kept.

The /etc/logadm.conf file also contains time stamps of when the last log rotation occurred.

You can use the logadm command to customize system logging and to add additional logging in the /etc/logadm.conf file as needed.

For example, to rotate the Apache access and error logs, use the following commands:

```
# logadm -w /var/apache/logs/access_log -s 100m
# logadm -w /var/apache/logs/error_log -s 10m
```

In this example, the Apache access_log file is rotated when it reaches 100 Mbytes in size, with a .0, .1, (and so on) suffix, keeping 10 copies of the old access_log file. The error_log is rotated when it reaches 10 Mbytes in size with the same suffixes and number of copies as the access_log file.

The /etc/logadm.conf entries for the preceding Apache log rotation examples look similar to the following:

```
# cat /etc/logadm.conf
.
.
.
/var/apache/logs/error_log -s 10m
/var/apache/logs/access_log -s 100m
```

For more information, see logadm(1M).

You can use the logadm command as superuser or by assuming an equivalent role (with Log Management rights). With role-based access control (RBAC), you can grant non-root users the privilege of maintaining log files by providing access to the logadm command.

For example, add the following entry to the /etc/user_attr file to grant user andy the ability to use the logadm command:

```
andy::::profiles=Log Management
```

Or, you can set up a role for log management by using the Solaris Management Console. For more information about setting up a role, see "Role-Based Access Control (Overview)" in *System Administration Guide: Security Services*.

Customizing System Message Logging

You can capture additional error messages that are generated by various system processes by modifying the `/etc/syslog.conf` file. By default, the `/etc/syslog.conf` file directs many system process messages to the `/var/adm/messages` files. Crash and boot messages are stored here as well. To view `/var/adm` messages, see "How to View System Messages" on page 392.

The `/etc/syslog.conf` file has two columns separated by tabs:

facility.level ... action

facility.level — A *facility* or system source of the message or condition. May be a comma-separated listed of facilities. Facility values are listed in Table 22–1. A *level*, indicates the severity or priority of the condition being logged. Priority levels are listed in Table 22–2.

> **Note –** Two entries for the same facility should not be put on the same line even if the entries are for different priorities. Putting a priority in the syslog file indicates that all messages of that priority and higher are logged. The last message takes precedence. For a given facility and level, `syslogd` matches all messages for that level and all higher levels.

action — The action field indicates where the messages are forwarded.

The following example shows sample lines from a default `/etc/syslog.conf` file.

```
user.err                                        /dev/sysmsg
user.err                                        /var/adm/messages
user.alert                                      'root, operator'
user.emerg                                      *
```

This means the following user messages are automatically logged:

- User errors are printed to the console and also are logged to the `/var/adm/messages` file.
- User messages requiring immediate action (`alert`) are sent to the root and operator users.

- User emergency messages are sent to individual users.

The most common error condition sources are shown in the following table. The most common priorities are shown in Table 22–2 in order of severity.

TABLE 22–1 Source Facilities for `syslog.conf` Messages

Source	Description
`kern`	The kernel
`auth`	Authentication
`daemon`	All daemons
`mail`	Mail system
`lp`	Spooling system
`user`	User processes

Note – The number of `syslog` facilities that can be activated in the `/etc/syslog.conf` file is unlimited.

TABLE 22–2 Priority Levels for `syslog.conf` Messages

Priority	Description
`emerg`	System emergencies
`alert`	Errors requiring immediate correction
`crit`	Critical errors
`err`	Other errors
`info`	Informational messages
`debug`	Output used for debugging
`none`	This setting doesn't log output

▼ How to Customize System Message Logging

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Edit the `/etc/syslog.conf` file, adding or changing message sources, priorities, and message locations according to the syntax described in `syslog.conf`(4).**

3. **Exit the file, saving the changes.**

Example 22–2 Customizing System Message Logging

This sample `/etc/syslog.conf` `user.emerg` facility sends user emergency messages to root *and* individual users.

```
user.emerg                                      'root, *'
```

Enabling Remote Console Messaging

The following new console features improve your ability to troubleshoot remote systems:

- The `consadm` command enables you to select a serial device as an *auxiliary* (or remote) console. Using the `consadm` command, a system administrator can configure one or more serial ports to display redirected console messages and to host `sulogin` sessions when the system transitions between run levels. This feature enables you to dial in to a serial port with a modem to monitor console messages and participate in `init` state transitions. (For more information, see `sulogin`(1M) and the step-by-step procedures that follow.)

  While you can log in to a system using a port configured as an auxiliary console, it is primarily an output device displaying information that is also displayed on the default console. If boot scripts or other applications read and write to and from the default console, the write output displays on all the auxiliary consoles, but the input is only read from the default console. (For more information on using the `consadm` command during an interactive login session, see "Using the `consadm` Command During an Interactive Login Session" on page 398.)

- Console output now consists of kernel and `syslog` messages written to a new pseudo device, `/dev/sysmsg`. In addition, `rc` script startup messages are written to `/dev/msglog`. Previously, all of these messages were written to `/dev/console`.

  Scripts that direct console output to `/dev/console` need to be changed to `/dev/msglog` if you want to see script messages displayed on the auxiliary consoles. Programs referencing `/dev/console` should be explicitly modified to use `syslog()` or `strlog()` if you want messages to be redirected to an auxiliary device.

- The `consadm` command runs a daemon to monitor auxiliary console devices. Any display device designated as an auxiliary console that disconnects, hangs up or loses carrier, is removed from the auxiliary console device list and is no longer active. Enabling one or more auxiliary consoles does not disable message display on the default console; messages continue to display on /dev/console.

Using Auxiliary Console Messaging During Run Level Transitions

Keep the following in mind when using auxiliary console messaging during run level transitions:

- Input cannot come from an auxiliary console if user input is expected for an rc script that is run when a system is booting. The input must come from the default console.
- The sulogin program, invoked by init to prompt for the superuser password when transitioning between run levels, has been modified to send the superuser password prompt to each auxiliary device in addition to the default console device.
- When the system is in single-user mode and one or more auxiliary consoles are enabled using the consadm command, a console login session runs on the first device to supply the correct superuser password to the sulogin prompt. When the correct password is received from a console device, sulogin disables input from all other console devices.
- A message is displayed on the default console and the other auxiliary consoles when one of the consoles assumes single-user privileges. This message indicates which device has become the console by accepting a correct superuser password. If there is a loss of carrier on the auxiliary console running the single-user shell, one of two actions might occur:
  - If the auxiliary console represents a system at run level 1, the system proceeds to the default run level.
  - If the auxiliary console represents a system at run level S, the system displays the ENTER RUN LEVEL (0-6, s or S): message on the device where the init s or shutdown command had been entered from the shell. If there isn't any carrier on that device either, you will have to reestablish carrier and enter the correct run level. The init or shutdown command will not redisplay the run-level prompt.
- If you are logged in to a system using a serial port, and an init or shutdown command is issued to transition to another run level, the login session is lost whether this device is the auxiliary console or not. This situation is identical to Solaris releases without auxiliary console capabilities.
- Once a device is selected as an auxiliary console using the consadm command, it remains the auxiliary console until the system is rebooted or the auxiliary console is unselected. However, the consadm command includes an option to set a device as the auxiliary console across system reboots. (See the following procedure for step-by-step instructions.)

Using the `consadm` Command During an Interactive Login Session

If you want to run an interactive login session by logging in to a system using a terminal that is connected to a serial port, and then using the `consadm` command to see the console messages from the terminal, note the following behavior.

- If you use the terminal for an interactive login session while the auxiliary console is active, the console messages are sent to the `/dev/sysmsg` or `/dev/msglog` devices.
- While you issue commands on the terminal, input goes to your interactive session and not to the default console (`/dev/console`).
- If you run the `init` command to change run levels, the remote console software kills your interactive session and runs the `sulogin` program. At this point, input is accepted only from the terminal and is treated like it's coming from a console device. This allows you to enter your password to the `sulogin` program as described in "Using Auxiliary Console Messaging During Run Level Transitions" on page 397.

  Then, if you enter the correct password on the (auxiliary) terminal, the auxiliary console runs an interactive `sulogin` session, locks out the default console and any competing auxiliary console. This means the terminal essentially functions as the system console.
- From here you can change to run level 3 or go to another run level. If you change run levels, `sulogin` runs again on all console devices. If you exit or specify that the system should come up to run level 3, then all auxiliary consoles lose their ability to provide input. They revert to being display devices for console messages.

  As the system is coming up, you must provide information to `rc` scripts on the default console device. After the system comes back up, the `login` program runs on the serial ports and you can log back into another interactive session. If you've designated the device to be an auxiliary console, you will continue to get console messages on your terminal, but all input from the terminal goes to your interactive session.

▼ How to Enable an Auxiliary (Remote) Console

The `consadm` daemon does not start monitoring the port until after you add the auxiliary console with the `consadm` command. As a security feature, console messages are only redirected until carrier drops, or the auxiliary console device is unselected. This means carrier must be established on the port before you can successfully use the `consadm` command.

For more information on enabling an auxiliary console, see the `consadm`(1m) man page.

Steps **1. Log in to the system as superuser.**

2. **Enable the auxiliary console.**

```
# consadm -a devicename
```

3. **Verify that the current connection is the auxiliary console.**

```
# consadm
```

Example 22–3 Enabling an Auxiliary (Remote) Console

```
# consadm -a /dev/term/a
# consadm
 /dev/term/a
```

▼ How to Display a List of Auxiliary Consoles

Steps

1. **Log in to the system as superuser.**
2. **Select one of the following steps:**
   a. **Display the list of auxiliary consoles.**

   ```
   # consadm
   /dev/term/a
   ```

   b. **Display the list of persistent auxiliary consoles.**

   ```
   # consadm -p
   /dev/term/b
   ```

▼ How to Enable an Auxiliary (Remote) Console Across System Reboots

Steps

1. **Log in to the system as superuser.**
2. **Enable the auxiliary console across system reboots.**

   ```
   # consadm -a -p devicename
   ```

   This adds the device to the list of persistent auxiliary consoles.
3. **Verify that the device has been added to the list of persistent auxiliary consoles.**

   ```
   # consadm
   ```

Example 22–4 Enabling an Auxiliary (Remote) Console Across System Reboots

```
# consadm -a -p /dev/term/a
# consadm
```

```
/dev/term/a
```

▼ How to Disable an Auxiliary (Remote) Console

Steps

1. **Log in to the system as superuser.**
2. **Select one of the following steps:**
   a. **Disable the auxiliary console.**

   ```
   # consadm -d devicename
   ```

   or

   b. **Disable the auxiliary console and remove it from the list of persistent auxiliary consoles.**

   ```
   # consadm -p -d devicename
   ```

3. **Verify that the auxiliary console has been disabled.**

   ```
   # consadm
   ```

Example 22–5 Disabling an Auxiliary (Remote) Console

```
# consadm -d /dev/term/a
# consadm
```

CHAPTER 23

Managing Core Files (Tasks)

This chapter describes how to manage core files with the `coreadm` command.

For information on the procedures associated with managing core files, see "Managing Core Files (Task Map)" on page 401.

Managing Core Files (Task Map)

Task	Description	For Instructions
1. Display the current core dump configuration	Display the current core dump configuration by using the `coreadm` command.	"How to Display the Current Core Dump Configuration" on page 404
2. Modify the core dump configuration	Modify the core dump configuration to do one of the following: Set a core file name pattern. Enable a per-process core file path. Enable a global core file path.	"How to Set a Core File Name Pattern" on page 405 "How to Enable a Per-Process Core File Path" on page 405 "How to Enable a Global Core File Path" on page 405
3. Examine a Core Dump File	Use the `proc` tools to view a core dump file.	"Examining Core Files" on page 406

Managing Core Files Overview

Core files are generated when a process or application terminates abnormally. Core files are managed with the `coreadm` command.

For example, you can use the `coreadm` command to configure a system so that all process core files are placed in a single system directory. This means it is easier to track problems by examining the core files in a specific directory whenever a Solaris process or daemon terminates abnormally.

Configurable Core File Paths

Two new configurable `core` file paths that can be enabled or disabled independently of each other are:

- A per-process core file path, which defaults to `core` and is enabled by default. If enabled, the per-process core file path causes a `core` file to be produced when the process terminates abnormally. The per-process path is inherited by a new process from its parent process.

  When generated, a per-process core file is owned by the owner of the process with read/write permissions for the owner. Only the owning user can view this file.
- A global core file path, which defaults to `core` and is disabled by default. If enabled, an *additional* core file with the same content as the per-process core file is produced by using the global core file path.

  When generated, a global core file is owned by superuser with read/write permissions for superuser only. Non-privileged users cannot view this file.

When a process terminates abnormally, it produces a core file in the current directory by default. If the global core file path is enabled, each abnormally terminating process might produce two files, one in the curent working directory, and one in the global core file location.

By default, a `setuid` process does not produce core files using either the global or per-process path.

Expanded Core File Names

If a global core file directory is enabled, core files can be distinguished from one another by using the variables described in the following table.

Variable Name	Variable Definition
%d	Executable file directory name, up to a maximum of MAXPATHLEN characters
%f	Executable file name, up to a maximum of MAXCOMLEN characters
%g	Effective group ID
%m	Machine name (uname -m)
%n	System node name)uname -n)
%p	Process ID
%t	Decimal value of time(2)
%u	Effective user ID
%z	Name of the zone in which process is executed (zonename)
%%	Literal %

For example, if the global core file path is set to:

```
/var/core/core.%f.%p
```

and a sendmail process with PID 12345 terminates abnormally, it produces the following core file:

```
/var/core/core.sendmail.12345
```

Setting the Core File Name Pattern

You can set a core file name pattern on a global, zone, or per-process basis. In additin, you can set the per-process defaults that persist across a system reboot.

For example, the following coreadm command sets the default per-process core file pattern. This setting applies to all processes that have not explicitly overridden the default core file pattern. This setting persists across system reboots.

```
# coreadm -i /var/core/core.%f.%p
```

Global core values are stored in the /etc/coreadm.conf file, which means these settings are saved across a system reboot.

This coreadm command sets the per-process core file name pattern for any processes:

```
$ coreadm -p /var/core/core.%f.%p $$
```

The $$ symbols represent a placeholder for the process ID of the currently running shell. The per-process core file name pattern is inherited by all child processes.

Once a global or per-process core file name pattern is set, it must be enabled with the coreadm -e command. See the following procedures for more information.

You can set the core file name pattern for all processes run during a user's login session by putting the command in a user's $HOME/.profile or .login file.

Enabling setuid Programs to Produce Core Files

You can use the coreadm command to enable or disable setuid programs to produce core files for all system processes or on a per-process basis by setting the following paths:

- If the global setuid option is enabled, a global core file path allows all setuid programs on a system to produce core files.
- If the per-process setuid option is enable, a per-process core file path allows specific setuid processes to produce core files.

By default, both flags are disabled. For security reasons, the global core file path must be a full pathname, starting with a leading /. If superuser disables per-process core files, individual users cannot obtain core files.

The setuid core files are owned by superuser with read/write permissions for superuser only. Regular users cannot access them even if the process that produced the setuid core file was owned by an ordinary user.

For more information, see coreadm(1M).

How to Display the Current Core Dump Configuration

Use the coreadm command without any options to display the current core dump configuration.

```
$ coreadm
              global core file pattern:
     global core file content: default
       init core file pattern: core
       init core file content: default
            global core dumps: disabled
       per-process core dumps: enabled
```

```
     global setid core dumps: disabled
per-process setid core dumps: disabled
    global core dump logging: disabled
```

▼ How to Set a Core File Name Pattern

Step ● **Determine whether you want to set a per-process or global core file and select one of the following:**

a. **Set a per-process file name pattern.**

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

b. **Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

c. **Set a global file name pattern.**

```
# coreadm -g /var/corefiles/%f.%p
```

▼ How to Enable a Per-Process Core File Path

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Enable a per-process core file path.**

   ```
   # coreadm -e process
   ```

3. **Display the current process core file path to verify the configuration.**

   ```
   $ coreadm $$
   1180:   /home/kryten/corefiles/%f.%p
   ```

▼ How to Enable a Global Core File Path

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Enable a global core file path.**

```
# coreadm -e global -g /var/core/core.%f.%p
```

3. **Display the current process core file path to verify the configuration.**

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

Troubleshooting Core File Problems

Error Message

```
NOTICE: 'set allow_setid_core = 1' in /etc/system is obsolete
NOTICE: Use the coreadm command instead of 'allow_setid_core'
```

Cause
: You have an obsolete parameter that allows `setuid` core files in your `/etc/system` file.

Solution
: Remove `allow_setid_core=1` from the `/etc/system` file. Then use the `coreadm` command to enable global `setuid` core file paths.

Examining Core Files

Some of the proc tools have been enhanced to examine process core files as well as live processes. The proc tools are utilities that can manipulate features of the `/proc` file system.

The `/usr/proc/bin/pstack`, `pmap`, `pldd`, `pflags`, and `pcred` tools can now be applied to core files by specifying the name of the core file on the command line, similar to the way you specify a process ID to these commands.

For more information on using proc tools to examine core files, see `proc`(1).

EXAMPLE 23–1 Examining Core Files With proc Tools

```
$ ./a.out
Segmentation Fault(coredump)
$ /usr/proc/bin/pstack ./core
core './core' of 19305: ./a.out
 000108c4 main     (1, ffbef5cc, ffbef5d4, 20800, 0, 0) + 1c
 00010880 _start   (0, 0, 0, 0, 0, 0) + b8
```

CHAPTER 24

Managing System Crash Information (Tasks)

This chapter describes how to manage system crash information in the Solaris Operating System.

For information on the procedures associated with managing system crash information, see "Managing System Crash Information (Task Map)" on page 409.

Managing System Crash Information (Task Map)

The following task map identifies the procedures needed to manage system crash information.

Task	Description	For Instructions
1. Display the current crash dump configuration	Display the current crash dump configuration by using the `dumpadm` command.	"How to Display the Current Crash Dump Configuration" on page 413
2. Modify the crash dump configuration	Use the `dumpadm` command to specify the type of data to dump, whether or not the system will use a dedicated dump device, the directory for saving crash dump files, and the amount of space that must remain available after crash dump files are written.	"How to Modify a Crash Dump Configuration" on page 413
3. Examine a crash dump file	Use the `mdb` command to view crash dump files.	"How to Examine a Crash Dump" on page 415

Task	Description	For Instructions
4. (Optional) Recover from a full crash dump directory	The system crashes but no room is available in the `savecore` directory, and you want to save some critical system crash dump information.	"How to Recover From a Full Crash Dump Directory (Optional)" on page 416
5. (Optional) Disable or enable the saving of crash dump files	Use the `dumpadm` command to disable or enable the saving the crash dump files. Saving crash dump files is enabled by default.	"How to Disable or Enable Saving Crash Dumps" on page 416

System Crashes (Overview)

System crashes can occur due to hardware malfunctions, I/O problems, and software errors. If the system crashes, it will display an error message on the console, and then write a copy of its physical memory to the dump device. The system will then reboot automatically. When the system reboots, the `savecore` command is executed to retrieve the data from the dump device and write the saved crash dump to your `savecore` directory. The saved crash dump files provide invaluable information to your support provider to aid in diagnosing the problem.

System Crash Dump Files

The `savecore` command runs automatically after a system crash to retrieve the crash dump information from the dump device and writes a pair of files called `unix.`*X* and `vmcore.`*X*, where *X* identifies the dump sequence number. Together, these files represent the saved system crash dump information.

Crash dump files are sometimes confused with *core* files, which are images of user applications that are written when the application terminates abnormally.

Crash dump files are saved in a predetermined directory, which by default, is `/var/crash/`*hostname*. In previous Solaris releases, crash dump files were overwritten when a system rebooted, unless you manually enabled the system to save the images of physical memory in a crash dump file. Now, the saving of crash dump files is enabled by default.

System crash information is managed with the `dumpadm` command. For more information, see "The `dumpadm` Command" on page 411.

Saving Crash Dumps

You can examine the control structures, active tables, memory images of a live or crashed system kernel, and other information about the operation of the kernel by using the `mdb` utility. Using `mdb` to its full potential requires a detailed knowledge of the kernel, and is beyond the scope of this manual. For information on using this utility, see the `mdb`(1) man page.

Additionally, crash dumps saved by `savecore` can be useful to send to a customer service representative for analysis of why the system is crashing.

The `dumpadm` Command

Use the `dumpadm` command to manage system crash dump information in the Solaris Operating System.

- The `dumpadm` command enables you to configure crash dumps of the operating system. The `dumpadm` configuration parameters include the dump content, dump device, and the directory in which crash dump files are saved.
- Dump data is stored in compressed format on the dump device. Kernel crash dump images can be as big as 4 Gbytes or more. Compressing the data means faster dumping and less disk space needed for the dump device.
- Saving crash dump files is run in the background when a dedicated dump device, not the swap area, is part of the dump configuration. This means a booting system does not wait for the `savecore` command to complete before going to the next step. On large memory systems, the system can be available before `savecore` completes.
- System crash dump files, generated by the `savecore` command, are saved by default.
- The `savecore -L` command is a new feature which enables you to get a crash dump of the live running the Solaris OS. This command is intended for troubleshooting a running system by taking a snapshot of memory during some bad state, such as a transient performance problem or service outage. If the system is up and you can still run some commands, you can execute the `savecore -L` command to save a snapshot of the system to the dump device, and then immediately write out the crash dump files to your `savecore` directory. Because the system is still running, you can only use the `savecore -L` command if you have configured a dedicated dump device.

The following table describes `dumpadm`'s configuration parameters.

Dump Parameter	Description
dump device	The device that stores dump data temporarily as the system crashes. When the dump device is not the swap area, savecore runs in the background, which speeds up the boot process.
savecore directory	The directory that stores system crash dump files.
dump content	Type of memory data to dump.
minimum free space	Minimum amount of free space required in the savecore directory after saving crash dump files. If no minimum free space has been configured, the default is one Mbyte.

For more information, see dumpadm(1M).

The dump configuration parameters managed by the dumpadm command are stored in the /etc/dumpadm.conf file.

Note – Do not edit the /etc/dumpadm.conf file manually. Editing this file manually could result in an inconsistent system dump configuration.

How the dumpadm Command Works

During system startup, the dumpadm command is invoked by the /etc/init.d/savecore script to configure crash dumps parameters based on information in the /etc/dumpadm.conf file.

Specifically, dumpadm initializes the dump device and the dump content through the /dev/dump interface.

After the dump configuration is complete, the savecore script looks for the location of the crash dump file directory by parsing the content of /etc/dumpadm.conf file. Then, savecore is invoked to check for crash dumps and check the content of the minfree file in the crash dump directory.

Dump Devices and Volume Managers

Do not configure a dedicated dump device that is under the control of volume management product such as Solaris Volume Manager for accessibility and performance reasons. You can keep your swap areas under the control of Solaris Volume Manager and this is a recommend practice, but keep your dump device separate.

Managing System Crash Dump Information

Keep the following key points in mind when you are working with system crash information:

- You must be superuser or assume an equivalent role to access and manage system crash information.
- Do not disable the option of saving system crash dumps. System crash dump files provide an invaluable way to determine what is causing the system to crash.
- Do not remove important system crash information until it has been sent to your customer service representative.

▼ How to Display the Current Crash Dump Configuration

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Display the current crash dump configuration.**

   ```
   # dumpadm      Dump content: kernel pages
          Dump device: /dev/dsk/c0t3d0s1 (swap)
   Savecore directory: /var/crash/venus
     Savecore enabled: yes
   ```

   The preceding example output means:

   - The dump content is kernel memory pages.
   - Kernel memory will be dumped on a swap device, `/dev/dsk/c0t3d0s1`. You can identify all your swap areas with the `swap -l` command.
   - System crash dump files will be written in the `/var/crash/venus` directory.
   - Saving crash dump files is enabled.

▼ How to Modify a Crash Dump Configuration

Steps

1. **Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Identify the current crash dump configuration.**

```
# dumpadm
      Dump content: kernel pages
       Dump device: /dev/dsk/c0t3d0s1 (swap)
Savecore directory: /var/crash/pluto
  Savecore enabled: yes
```

This output identifies the default dump configuration for a system running the Solaris 9 release.

3. **Modify the crash dump configuration.**

```
# dumpadm -c content -d dump-device -m nnnk | nnnm | nnn% -n -s savecore-dir
```

`-c` *content*	Specifies the type of data to dump. Use `kernel` to dump of all kernel memory, `all` to dump all of memory, or `curproc`, to dump kernel memory and the memory pages of the process whose thread was executing when the crash occurred. The default dump content is kernel memory.
`-d` *dump-device*	Specifies the device that stores dump data temporarily as the system crashes. The primary swap device is the default dump device.
`-m` *nnnk* | *nnnm* | *nnn%*	Specifies the minimum free disk space for saving crash dump files by creating a `minfree` file in the current `savecore` directory. This parameter can be specified in Kbytes (`nnnk`), Mbytes (`nnnm`) or file system size percentage (`nnn%`). The `savecore` command consults this file prior to writing the crash dump files. If writing the crash dump files, based on their size, would decrease the amount of free space below the `minfree` threshold, the dump files are not written and an error message is logged. For information on recovering from this scenario, see "How to Recover From a Full Crash Dump Directory (Optional)" on page 416.
`-n`	Specifies that `savecore` should not be run when the system reboots. This dump configuration is not recommended. If system crash information is written to the swap device, and `savecore` is not enabled, the crash dump information is overwritten when the system begins to swap.

`-s`	Specifies an alternate directory for storing crash dump files. The default directory is `/var/crash/hostname` where `hostname` is the output of the `uname -n` command.

Example 24–1 Modifying a Crash Dump Configuration

In this example, all of memory is dumped to the dedicated dump device, `/dev/dsk/c0t1d0s1`, and the minimum free space that must be available after the crash dump files are saved is 10% of the file system space.

```
# dumpadm
      Dump content: kernel pages
       Dump device: /dev/dsk/c0t3d0s1 (swap)
Savecore directory: /var/crash/pluto
  Savecore enabled: yes
 # dumpadm -c all -d /dev/dsk/c0t1d0s1 -m 10%
      Dump content: all pages
       Dump device: /dev/dsk/c0t1d0s1 (dedicated)
Savecore directory: /var/crash/pluto (minfree = 77071KB)
  Savecore enabled: yes
```

▼ How to Examine a Crash Dump

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Examine a crash dump by using the `mdb` utility.**

   ```
   # /usr/bin/mdb [-k] crashdump-file
   ```

`-k`	Specifies kernel debugging mode by assuming the file is an operating system crash dump file.
crashdump-file	Specifies the operating system crash dump file.

3. **Display crash status information.**

   ```
   # /usr/bin/mdb file-name
   > ::status
      .
      .
      .
   > ::system
      .
      .
   ```

.

Example 24–2 Examining a Crash Dump

The following example shows sample output from the `mdb` utility, which includes system information and identifies the tunables that are set in this system's `/etc/system` file.

```
# /usr/bin/mdb -k unix.0
Loading modules: [ unix krtld genunix ip nfs ipc ptm ]
> ::status
debugging crash dump /dev/mem (64-bit) from ozlo
operating system: 5.10 Generic (sun4u)
> ::system
set ufs_ninode=0x9c40 [0t40000]
set ncsize=0x4e20 [0t20000]
set pt_cnt=0x400 [0t1024]
```

▼ How to Recover From a Full Crash Dump Directory (Optional)

In this scenario, the system crashes but no room is left in the `savecore` directory, and you want to save some critical system crash dump information.

Steps

1. **Log in as superuser or assume an equivalent role after the system reboots.**
2. **Clear out the savecore directory, usually `/var/crash/`*hostname*, by removing existing crash dump files that have already been sent to your service provider. Or, run the `savecore` command and specify an alternate directory that has sufficient disk space. See the next step.**
3. **Manually run the `savecore` command and if necessary, specify an alternate savecore directory.**

   ```
   # savecore [ directory ]
   ```

▼ How to Disable or Enable Saving Crash Dumps

Steps

1. **Become superuser or assume an equivalent role.**

   Roles contain authorizations and privileged commands. For more information about roles, see "Configuring RBAC (Task Map)" in *System Administration Guide: Security Services*.

2. **Disable or enable the saving of crash dumps on your system.**

```
# dumpadm -n | -y
```

Example 24–3 Disabling the Saving of Crash Dumps

This example illustrates how to disable the saving of crash dumps on your system.

```
# dumpadm -n
      Dump content: all pages
       Dump device: /dev/dsk/c0t1d0s1 (dedicated)
Savecore directory: /var/crash/pluto (minfree = 77071KB)
  Savecore enabled: no
```

Example 24–4 Enabling the Saving of Crash Dumps

This example illustrates how to enable the saving of crash dump on your system.

```
# dumpadm -y
      Dump content: all pages
       Dump device: /dev/dsk/c0t1d0s1 (dedicated)
Savecore directory: /var/crash/pluto (minfree = 77071KB)
  Savecore enabled: yes
```

CHAPTER 25

Troubleshooting Miscellaneous Software Problems (Tasks)

This chapter describes miscellaneous software problems that may occur occasionally and are relatively easy to fix. Troubleshooting miscellaneous software problems includes solving problems that aren't related to a specific software application or topic, such as unsuccessful reboots and full file systems. Resolving these problems are described in the following sections.

This is a list of the information in this chapter.

What to Do If Rebooting Fails

If the system does not reboot completely, or if it reboots and then crashes again, there may be a software or hardware problem that is preventing the system from booting successfully.

Cause of System Not Booting	How to Fix the Problem
The system can't find `/platform/‘uname -m‘/kernel/unix`.	You may need to change the `boot-device` setting in the PROM on a SPARC based system. For information on changing the default boot device, see Chapter 11, "Booting a System (Tasks)," in *System Administration Guide: Basic Administration*.

Cause of System Not Booting	How to Fix the Problem
There is no default boot device on an x86 based system. The message displayed is: `Not a UFS filesystem.`	Boot the system using the Configuration Assistant/boot diskette and select the disk from which to boot.
There's an invalid entry in the `/etc/passwd` file.	For information on recovering from an invalid `passwd` file, see Chapter 11, "Booting a System (Tasks)," in *System Administration Guide: Basic Administration* or Chapter 12, "Booting a System (Tasks)," in *System Administration Guide: Basic Administration*.
There's a hardware problem with a disk or another device.	Check the hardware connections: ■ Make sure the equipment is plugged in. ■ Make sure all the switches are set properly. ■ Look at all the connectors and cables, including the Ethernet cables. ■ If all this fails, turn off the power to the system, wait 10 to 20 seconds, and then turn on the power again.

If none of the above suggestions solve the problem, contact your local service provider.

64-bit SPARC: Troubleshooting 64–bit Solaris Boot Problems

The current Solaris release only ships a 64–bit kernel on SPARC based systems. The 32–bit kernel is no longer shipped on SPARC based systems. If the flash PROM is out of date, the Solaris OS will not run. You will have to download the SMCC flash prom CD-ROM or use a previous release of the Solaris OS, in 32–bit mode, in order to upgrade the flash PROM.

What to Do if You Forgot Root Password

If you forget the root password and you cannot log into the system, you will have to:

- Stop the system by using the keyboard stop sequence.
- Boot from a boot or install server or from a local CD-ROM.
- Mount the root (/) file system.

- Remove the root password from the /etc/shadow file.
- Reboot the system.
- Log in and set root's password.

This procedure is fully described in "Booting a System" in *System Administration Guide: Basic Administration* and "Booting a System" in *System Administration Guide: Basic Administration*.

The following examples describe how to recover from a forgotten root password on both SPARC and x86 based systems.

EXAMPLE 25–1 SPARC: What to Do if You Forgot Root Password

The following example shows how to recover when you forget the root password by booting from the network. This example assumes that the boot server is already available. Be sure to apply a new root password after the system has rebooted.

```
(Use keyboard abort sequence--Press Stop A keys to stop the system)
ok boot net -s
# mount /dev/dsk/c0t3d0s0 /a
# cd /a/etc
# TERM=vt100
# export TERM
# vi shadow
(Remove root's encrypted password string)
# cd /
# umount /a
# init 6
```

EXAMPLE 25–2 x86: What to Do if You Forgot Root Password

The following example shows how to recover when you forget root's password by booting from the network. This example assumes that the boot server is already available. Be sure to apply a new root password after the system has rebooted.

```
(Use keyboard abort sequence to stop the system
Type any key to continue

SunOS Secondary Boot version 3.00
Autobooting from bootpath: /pci@0,0/pci1028,10a@3/sd@0,0:a
Running Configuration Assistant...

If the system hardware has changed, or to boot from a different
device, interrupt the autoboot process by pressing ESC.

Initializing system
Please wait...
     <<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci1028,10a@3/sd@0,0:a
Boot args:

Type    b [file-name] [boot-flags] <ENTER>     to boot with options
or      i <ENTER>                              to enter boot interpreter
```

```
or      <ENTER>                              to boot with defaults

<<< timeout in 5 seconds >>>

Select (b)oot or (i)nterpreter: b -s
Loading kmdb...
SunOS Release 5.10 Version s10_62 32-bit
Copyright 1983-2004 Sun Microsystems, Inc.  All rights reserved.
Use is subject to license terms.
configuring IPv4 interfaces: iprb0.
add net default: gateway 172.20.26.248
Hostname: neptune
.
.
.
# mount /dev/dsk/c0t0d0s0 /a
      .
      .
      .
# cd /a/etc
# vi shadow
(Remove root's encrypted password string)
# cd /
# umount /a
# init 6
```

What to Do if a System Hangs

A system can freeze or hang rather than crash completely if some software process is stuck. Follow these steps to recover from a hung system.

1. Determine whether the system is running a window environment and follow these suggestions. If these suggestions don't solve the problem, go to step 2.
   - Make sure the pointer is in the window where you are typing the commands
   - Press Control-q in case the user accidentally pressed Control-s, which freezes the screen. Control-s freezes only the window, not the entire screen. If a window is frozen, try using another window.
   - If possible, log in remotely from another system on the network. Use the `pgrep` command to look for the hung process. If it looks like the window system is hung, identify the process and kill it.
2. Press Control-\ to force a "quit" in the running program and (probably) write out a `core` file.
3. Press Control-c to interrupt the program that might be running.

4. Log in remotely and attempt to identify and kill the process that is hanging the system.
5. Log in remotely, become superuser or assume an equivalent role and reboot the system.
6. If the system still does not respond, force a crash dump and reboot. For information on forcing a crash dump and booting, see "Forcing a Crash Dump and Rebooting the System" in *System Administration Guide: Basic Administration* or "Forcing a Crash Dump and Rebooting the System" in *System Administration Guide: Basic Administration*.
7. If the system still does not respond, turn the power off, wait a minute or so, then turn the power back on.
8. If you cannot get the system to respond at all, contact your local service provider for help.

What to Do if a File System Fills Up

When the root (/) file system or any other file system fills up, you will see the following message in the console window:

```
.... file system full
```

There are several reasons why a file system fills up. The following sections describe several scenarios for recovering from a full file system. For information on routinely cleaning out old and unused files to prevent full file systems, see Chapter 13.

File System Fills Up Because a Large File or Directory Was Created

Reason Error Occurred	How to Fix the Problem
Someone accidentally copied a file or directory to the wrong location. This also happens when an application crashes and writes a large `core` file into the file system.	Log in as superuser or assume an equivalent role and use the `ls -tl` command in the specific file system to identify which large file is newly created and remove it. For information on removing `core` files, see "How to Find and Delete `core` Files" on page 257.

A TMPFS File System is Full Because the System Ran Out of Memory

Reason Error Occurred	How to Fix the Problem
This can occur if TMPFS is trying to write more than it is allowed or some current processes are using a lot of memory.	For information on recovering from tmpfs-related error messages, see the tmpfs(7FS) man page.

What to Do if File ACLs Are Lost After Copy or Restore

Reason Error Occurred	How to Fix the Problem
If files or directories with ACLs are copied or restored into the /tmp directory, the ACL attributes are lost. The /tmp directory is usually mounted as a temporary file system, which doesn't support UFS file system attributes such as ACLs.	Copy or restore files into the /var/tmp directory instead.

Troubleshooting Backup Problems

This section describes some basic troubleshooting techniques to use when backing up and restoring data.

The root (/) File System Fills Up After You Back Up a File System

You back up a file system, and the root (/) file system fills up. Nothing is written to the media, and the ufsdump command prompts you to insert the second volume of media.

Reason Error Occurred	How to Fix the Problem
If you used an invalid destination device name with the -f option, the ufsdump command wrote to a file in the /dev directory of the root (/) file system, filling it up. For example, if you typed /dev/rmt/st0 instead of /dev/rmt/0, the backup file /dev/rmt/st0 was created on the disk rather than being sent to the tape drive.	Use the ls -tl command in the /dev directory to identify which file is newly created and abnormally large, and remove it.

Make Sure the Backup and Restore Commands Match

You can only use the ufsrestore command to restore files backed up with the ufsdump command. If you back up with the tar command, restore with the tar command. If you use the ufsrestore command to restore a tape that was written with another command, an error message tells you that the tape is not in ufsdump format.

Check to Make Sure You Have the Right Current Directory

It is easy to restore files to the wrong location. Because the ufsdump command always copies files with full path names relative to the root of the file system, you should usually change to the root directory of the file system before running the ufsrestore command. If you change to a lower-level directory, after you restore the files you will see a complete file tree created under that directory.

Interactive Commands

When you use the interactive command, a ufsrestore> prompt is displayed, as shown in this example:

```
# ufsrestore ivf /dev/rmt/0
Verify volume and initialize maps
Media block size is 126
Dump   date: Fri Jan 30 10:13:46 2004
Dumped from: the epoch
Level 0 dump of /export/home on starbug:/dev/dsk/c0t0d0s7
Label: none
Extract directories from tape
Initialize symbol table.
ufsrestore >
```

At the `ufsrestore>` prompt, you can use the commands listed on Chapter 27, "UFS Backup and Restore Commands (Reference)," in *System Administration Guide: Devices and File Systems* to find files, create a list of files to be restored, and restore them.

CHAPTER 26

Troubleshooting File Access Problems (Tasks)

This chapter provides information on resolving file access problems such as those related to incorrect permissions and search paths.

This is a list of troubleshooting topics in this chapter.

Users frequently experience problems, and call on a system administrator for help, because they cannot access a program, a file, or a directory that they could previously use. Whenever you encounter such a problem, investigate one of three areas:

- The user's search path may have been changed, or the directories in the search path may not be in the proper order.
- The file or directory may not have the proper permissions or ownership.
- The configuration of a system accessed over the network may have changed.

This chapter briefly describes how to recognize problems in each of these three areas and suggests possible solutions.

Solving Problems With Search Paths (`Command not found`)

A message of `Command not found` indicates one of the following:

- The command is not available on the system.
- The command directory is not in the search path.

To fix a search path problem, you need to know the pathname of the directory where the command is stored.

If the wrong version of the command is found, a directory that has a command of the same name is in the search path. In this case, the proper directory may be later in the search path or may not be present at all.

You can display your current search path by using the `echo $PATH` command. For example:

```
$ echo $PATH
/home/kryten/bin:/sbin:/usr/sbin:/usr/bin:/usr/dt:/usr/dist/exe
```

Use the `which` command to determine whether you are running the wrong version of the command. For example:

```
$ which acroread
/usr/doctools/bin/acroread
```

Note – The `which` command looks in the `.cshrc` file for path information. The `which` command might give misleading results if you execute it from the Bourne or Korn shell and you have a `.cshrc` file than contains aliases for the `which` command. To ensure accurate results, use the `which` command in a C shell, or, in the Korn shell, use the `whence` command.

▼ How to Diagnose and Correct Search Path Problems

Steps

1. **Display the current search path to verify that the directory for the command is not in your path or that it isn't misspelled.**

   ```
   $ echo $PATH
   ```

2. **Check the following:**
   - Is the search path correct?
   - Is the search path listed before other search paths where another version of the command is found?
   - Is the command in one of the search paths?

   If the path needs correction, go to step 3. Otherwise, go to step 4.

3. **Add the path to the appropriate file, as shown in this table.**

Shell	File	Syntax	Notes
Bourne and Korn	`$HOME/.profile`	`$ PATH=`*$HOME/bin:/sbin:/usr/local /bin ...* `$ export PATH`	A colon separates path names.

C	`$HOME/.cshrc`
	or
	`$HOME/.login`
		hostname% `set path=(`*~bin /sbin /usr/local/bin ...*`)`	A blank space separates path names.

4. **Activate the new path as follows:**

Shell	File Where Path Is Located	Use this Command to Activate The Path
Bourne and Korn	`.profile`	`$` **`. ./.profile`**
C	`.cshrc`	*hostname%* **`source .cshrc`**
	`.login`	*hostname%* **`source .login`**

5. **Verify the new path.**

   `$` **`which`** *command*

Example 26–1

Diagnosing and Correcting Search Path Problems

This example shows that the `mytool` executable is not in any of the directories in the search path using the `which` command.

```
venus% mytool
mytool: Command not found
venus% which mytool
no mytool in /sbin /usr/sbin /usr/bin /etc /home/ignatz/bin  .
venus% echo $PATH
/sbin /usr/sbin /usr/bin /etc /home/ignatz/bin
venus% vi ~.cshrc
(Add appropriate command directory to the search path)
venus% source .cshrc
venus% mytool
```

If you cannot find a command, look at the man page for its directory path. For example, if you cannot find the `lpsched` command (the `lp` printer daemon), the `lpsched`(1M) man page tells you the path is `/usr/lib/lp/lpsched`.

Solving File Access Problems

When users cannot access files or directories that they previously could access, the permissions or ownership of the files or directories probably has changed.

Changing File and Group Ownerships

Frequently, file and directory ownerships change because someone edited the files as superuser. When you create home directories for new users, be sure to make the user the owner of the dot (`.`) file in the home directory. When users do not own "`.`" they cannot create files in their own home directory.

Access problems can also arise when the group ownership changes or when a group of which a user is a member is deleted from the `/etc/group` database.

For information about how to change the permissions or ownership of a file that you are having problems accessing, see Chapter 6, "Controlling Access to Files (Tasks)," in *System Administration Guide: Security Services*.

Recognizing Problems With Network Access

If users have problems using the `rcp` remote copy command to copy files over the network, the directories and files on the remote system may have restricted access by setting permissions. Another possible source of trouble is that the remote system and the local system are not configured to allow access.

See "Strategies for NFS Troubleshooting" in *System Administration Guide: Network Services* for information about problems with network access and problems with accessing systems through AutoFS.

CHAPTER 27

Troubleshooting Printing Problems (Tasks)

This chapter explains how to troubleshoot printing problems that might occur when you set up or maintain printing services.

This is a list of the step-by-step instructions in this chapter.

For information about printing and the LP print service, see Chapter 1.

Tips on Troubleshooting Printing Problems

Sometimes after setting up a printer, you find that nothing prints. Or, you might get a little farther in the process: something prints, but it is not what you expect, the output is incorrect or illegible. Then, when you get past these problems, other problems might occur, such as:

- `lp` commands hanging
- Printers becoming idle
- Users getting conflicting messages

Note – Although many of the suggestions in this chapter are relevant to parallel printers, they are geared toward the more common serial printers.

Troubleshooting No Output (Nothing Prints)

When nothing prints, there are three general areas to check:

- The printer hardware
- The network
- The LP print service

If you get a banner page, but nothing else, this is a special case of incorrect output. See "Troubleshooting Incorrect Output" on page 433.

Check the Hardware

The hardware is the first area to check. As obvious as it sounds, you should make sure that the printer is plugged in and turned on. In addition, you should refer to the manufacturer's documentation for information about hardware settings. Some computers use hardware switches that change the characteristics of a printer port.

The printer hardware includes the printer, the cable that connects it to the computer, and the ports into which the cable plugs at each end. As a general approach, you should work your way from the printer to the computer. Check the printer. Check where the cable connects to the printer. Check the cable. Check where the cable connects to the computer.

Check the Network

Problems are more common with remote print requests that are going from a print client to a print server. You should make sure that network access between the print server and print clients is enabled.

If the network is running the Network Information Service Plus (NIS+), see *System Administration Guide: Naming and Directory Services (NIS+)* for instructions to enable access between systems. If the network is not running the Network Information Service (NIS) or NIS+, before you set up print servers and print clients, include the Internet address and system name for each client system in the `/etc/hosts` file on the print server. Also, the IP address and system name for the print server must be included in the `/etc/hosts` file of each print client system.

Check the LP Print Service

For printing to work, the LP scheduler must be running on both the print server and print client. If it is not running, you need to start it using the `/usr/lib/lp/lpsched` command. If you have trouble starting the scheduler, see "How to Restart the Print Scheduler" on page 90.

In addition to the scheduler running, a printer must be enabled and accepting requests before it will produce any output. If the LP print service is not accepting requests for a printer, the submitted print requests are rejected. Usually, in that instance, the user receives a warning message after submitting a print request. If the LP print service is not enabled for a printer, print requests remain queued on the system until the printer is enabled.

In general, you should analyze a printing problem as follows:

- Follow the path of the print request step-by-step.
- Examine the status of the LP print service at each step.
  - Is the configuration correct?
  - Is the printer accepting requests?
  - Is the printer enabled to process requests?
- If the request is hanging on transmission, set up `lpr.debug` in `syslog.conf` to display the flow.
- If the request is hanging locally, examine the `lpsched` log (`/var/lp/logs/lpsched`).
- If the request is hanging locally, have notification of the printer device errors (faults) mailed to you, and re-enable the printer.

The procedures found in "Troubleshooting Printing Problems" on page 436 use this strategy to help you troubleshoot various problems with the LP print service.

Troubleshooting Incorrect Output

If the printer and the print service software are not configured correctly, the printer might print, but it might provide output that is not what you expect.

Check the Printer Type and File Content Type

If you used the wrong printer type when you set up the printer with the LP print service, inappropriate printer control characters can be sent to the printer. The results are unpredictable: nothing might print, the output might be illegible, or the output might be printed in the wrong character set or font.

If you specified an incorrect file content type, the banner page might print, but that is all. The file content types specified for a printer indicate the types of files the printer can print directly, without filtering. When a user sends a file to the printer, the file is sent directly to the printer without any attempt to filter it. The problem occurs if the printer cannot handle the file content type.

When setting up print clients, you increase the chance for a mistake because the file content types must be correct on both the print server and the print client. If you set up the print client as recommended with `any` as the file content type, files are sent directly to the print server and the print server determines the need for filtering. Therefore, the file content types have to be specified correctly only on the server.

You can specify a file content on the print client to off-load filtering from the server to the client, but the content type must be supported on the print server.

Check the `stty` Settings

Many formatting problems can result when the default `stty` (standard terminal) settings do not match the settings required by the printer. The following sections describe what happens when some of the settings are incorrect.

Wrong Baud Settings

When the baud setting of the computer does not match the baud setting of the printer, usually you get some output, but it does not look like the file you submitted for printing. Random characters are displayed, with an unusual mixture of special characters and undesirable spacing. The default for the LP print service is 9600 baud.

Note – If a printer is connected by a parallel port, the baud setting is irrelevant.

Wrong Parity Setting

Some printers use a parity bit to ensure that data received for printing has not been garbled during transmission. The parity bit setting for the computer and the printer must match. If they do not match, some characters either will not be printed at all, or will be replaced by other characters. In this case, the output looks approximately correct. The word spacing is all right and many letters are in their correct place. The LP print service does not set the parity bit by default.

Wrong Tab Settings

If the file contains tabs, but the printer expects no tabs, the printed output might contain the complete contents of the file, but the text might be jammed against the right margin. Also, if the tab settings for the printer are incorrect, the text might not have a left margin, it might run together, it might be concentrated to a portion of the page, or it might be incorrectly double-spaced. The default is for tabs to be set every eight spaces.

Wrong Return Setting

If the output is double-spaced, but it should be single-spaced, either the tab settings for the printer are incorrect or the printer is adding a line feed after each return. The LP print service adds a return before each line feed, so the combination causes two line feeds.

If the print zigzags down the page, the stty option onlcr that sends a return before every line feed is not set. The stty=onlcr option is set by default, but you might have cleared it while trying to solve other printing problems.

Troubleshooting Hung lp Commands

If you type any of the lp commands (such as lpsystem, lpadmin, or lpstat) and nothing happens (no error message, status information, or prompt is displayed), chances are something is wrong with the LP scheduler. Such a problem can usually be resolved by stopping and restarting the LP scheduler. See "How to Stop the Print Scheduler" on page 89 for instructions.

Troubleshooting Idle (Hung) Printers

You might find a printer that is idle, even though it has print requests queued to it. A printer might seem idle when it should not be for one of the following reasons:

- The current print request is being filtered.
- The printer has a fault.
- Networking problems might be interrupting the printing process.

Check the Print Filters

Slow print filters run in the background to avoid tying up the printer. A print request that requires filtering will not print until it has been filtered.

Check Printer Faults

When the LP print service detects a fault, printing resumes automatically, but not immediately. The LP print service waits about five minutes before trying again, and continues trying until a request is printed successfully. You can force a retry immediately by enabling the printer.

Check Network Problems

When printing files over a network, you might encounter the following types of problems:

- Requests sent to print servers might back up in the client system (local) queue.
- Requests sent to print servers might back up in the print server (remote) queue.

Print Requests Backed Up in the Local Queue

Print requests submitted to a print server might back up in the client system queue for the following reasons:

- The print server is down.
- The printer is disabled on the print server.
- The network between the print client and print server is down.
- Underlying network software was not set up properly.

While you are tracking the source of the problem, you should stop new requests from being added to the queue. See "How to Accept or Reject Print Requests for a Printer" on page 110 for more information.

Print Requests Backed Up in the Remote Queue

If print requests back up in the print server queue, the printer has probably been disabled. When a printer is accepting requests, but not processing them, the requests are queued to print. Unless there is a further problem, once the printer is enabled, the print requests in the queue should print.

Troubleshooting Conflicting Status Messages

A user might enter a print request and be notified that the client system has accepted it, then receive mail from the print server that the print request has been rejected. These conflicting messages might occur for the following reasons:

- The print client might be accepting requests, while the print server is rejecting requests.
- The definition of the printer on the print client might not match the definition of that printer on the print server. More specifically, the definitions of the print job components, like filters, character sets, print wheels, or forms are not the same on the client and server systems.

You should check that identical definitions of these job components are registered on both the print clients and print servers so that local users can access printers on the print servers.

Troubleshooting Printing Problems

This section contains step-by-step instructions that explain:

- How to troubleshoot no output
- How to troubleshoot incorrect output
- How to unhang the `lp` commands
- How to troubleshoot an idle (hung) printer
- How to resolve conflicting status messages

How to Troubleshoot No Printer Output

This task includes the following troubleshooting procedures to try when you submit a print request to a printer and nothing prints:

- Check the hardware (“How to Check the Printer Hardware” on page 437).
- Check the network (“How To Check the Printer Network Connections” on page 438).
- Check the LP print service basic functions (“How to Check the Basic Functions of the LP Print Service” on page 438).
- Check printing from a Solaris print client to a Solaris print server (“How to Check Printing From a Solaris Print Client to a Solaris Print Server” on page 442).

Try the first three procedures in the order in which they are listed, before going to the print client/server section. However, if the banner page prints, but nothing else does, turn to the instructions under “How to Troubleshoot Incorrect Output” on page 442.

▼ How to Check the Printer Hardware

Steps

1. **Check that the printer is plugged in and turned on.**
2. **Check that the cable is connected to the port on the printer and to the port on the system or server.**
3. **Make sure that the cable is the correct cable and that it is not defective.**

   Refer to the manufacturer's documentation. If the printer is connected to a serial port, verify that the cable supports hardware flow control. A NULL modem adapter supports this. The following table shows the pin configuration for NULL modem cables.

	Host	Printer
Mini-Din-8	**25-Pin D-sub**	**25-Pin D-sub**
-	1 (FG)	1(FG)
3(TD)	2(TD)	3(RD)
5(RD)	3(RD)	2(TD)
6(RTS)	4(RTS)	5(CTS)
2(CTS)	5(CTS)	4(RTS)
4(SG)	7(SG)	7(SG)
7(DCD)	6(DSR), 8(DCD)	20(DTR)
1(DTR)	20(DTR)	6(DSR), 8(DCD)

4. **Check that any hardware switches for the ports are set properly.**

   See the printer documentation for the correct settings.

5. **Check that the printer is operational.**

   Use the printer's self-test feature, if the printer has one. Check the printer documentation for information about printer self-testing.

6. **Check that the baud settings for the computer and the printer are correct.**

   If the baud settings are not the same for both the computer and the printer, sometimes nothing will print, but more often you get incorrect output. For instructions, see "How to Troubleshoot Incorrect Output" on page 442.

▼ How To Check the Printer Network Connections

Steps

1. **Check that the network link between the print server and the print client is set up correctly.**

   ```
   print_client# ping print_server
   print_server is alive
   print_server# ping print_client
   print_client not available
   ```

   If the message says the system is alive, you know you can reach the system, so the network is all right. The message also tells you that either a name service or the local `/etc/hosts` file has translated the host (system) name you entered into an IP address; otherwise, you would need to enter the IP address.

   If you get a `not available` message, try to answer the following questions: How is NIS or NIS+ set up at your site? Do you need to take additional steps so that print servers and print clients can communicate with one another? If your site is not running NIS or NIS+, have you entered the IP address for the print server in each print client's `/etc/hosts` file, and entered all print client IP addresses in the `/etc/hosts` file of the print server?

2. **(On a SunOS 5.0–5.1 print server only) Check that the `listen` port monitor is configured correctly.**

3. **(On a SunOS 5.0–5.1 print server only) Check that the network listen services are registered with the port monitor on the print server.**

▼ How to Check the Basic Functions of the LP Print Service

This procedure uses the printer `luna` as an example of checking basic LP print service functions.

Steps 1. **On the print server, make sure that the LP print service is running.**

   a. **Check whether the LP scheduler is running.**

   ```
   # svcs application/print/server
   scheduler is running
   ```

   b. **If the scheduler is not running, become superuser, `lp`, or assume an equivalent role, and start the scheduler.**

   ```
   # /usr/lib/lp/lpsched
   ```

   If you have trouble starting the scheduler, see "How to Unhang the LP Print Service" on page 447.

2. **On both the print server and print client, make sure that the printer is accepting requests.**

   a. **Check that the printer is accepting requests.**

   ```
   # lpstat -a
   mars accepting requests since Jan 12 14:23 2004
   luna not accepting requests since Jan 12 14:23 2004
   unknown reason
   ```

   This command verifies that the LP system is accepting requests for each printer configured for the system.

   b. **If the printer is not accepting requests, become superuser, `lp`, or assume an equivalent role, and allow the printer to accept print requests.**

   ```
   # accept luna
   ```

   The specified printer now accepts requests.

3. **On both the print server and print client, make sure that the printer is enabled to print submitted print requests.**

   a. **Check that the printer is enabled.**

   ```
   # lpstat -p luna
   printer luna disabled since Jan 12 14:25 2004.
   available.
   unknown reason
   ```

   This command displays information about printer status. You can omit the printer name to obtain information about all printers set up for the system. The following example shows a printer that is disabled.

   b. **If the printer is disabled, become superuser, `lp`, or assume an equivalent role, and enable the printer.**

   ```
   # enable luna
   printer "luna" now enabled.
   ```

   The specified printer is enabled to process print requests.

4. **On the print server, make sure that the printer is connected to the correct serial port.**

   a. **Check that the printer is connected to the correct serial port.**

   ```
   # lpstat -t
   scheduler is running
   system default destination: luna
   device for luna: /dev/term/a
   ```

   The message `device for` *printer-name* shows the port address. Is the cable connected to the port to which the LP print service says is connected? If the port is correct, skip to Step 5.

   b. **Become superuser, `lp`, or assume an equivalent role.**

   c. **Change the file ownership of the device file that represents the port.**

   ```
   # chown lp device-filename
   ```

   This command assigns the special user `lp` as the owner of the device file. In this command, *device-filename* is the name of the device file.

   d. **Change the permissions on the printer port device file.**

   ```
   # chmod 600 device-filename
   ```

   This command allows only superuser, `lp`, or an equivalent role to access the printer port device file.

5. **On both the print server and print client, make sure that the printer is configured properly.**

   a. **Check that the printer is configured properly.**

   ```
   # lpstat -p luna -l
   printer luna is idle. enabled since Jan 12 14:24 2004. available
           Content types: postscript
           Printer types: PS
   ```

   The above example shows a PostScript printer that is configured properly, and that is available to process print requests. If the printer type and file content type are correct, skip to Step 6.

   b. **If the printer type or file content type is incorrect, try setting the print type to `unknown` and the content type to `any` on the print client.**

   ```
   # lpadmin -p printer-name -T printer-type -I file-content-type
   ```

6. **On the print server, make sure that the printer is not faulted.**

   a. **Check that the printer is not waiting because of a printer fault.**

   ```
   # lpadmin -p printer-name -F continue
   ```

   This command instructs the LP print service to continue if it is waiting because of a fault.

b. Force an immediate retry by re-enabling the printer.

```
# enable printer-name
```

c. (Optional) Instruct the LP print service to enable quick notification of printer faults.

```
# lpadmin -p printer-name -A 'write root'
```

This command instructs the LP print service to set a default policy of writing root, sending the printer fault message to the terminal on which root is logged in, if the printer fails. This might help you get quick notification of faults as you try to fix the problem.

7. Make sure that the printer is not set up incorrectly as a login terminal.

Note – It is easy to mistakenly set up a printer as a login terminal, so be sure to check this possibility even if you think it does not apply.

---

a. Look for the printer port entry in the `ps -ef` command output.

```
# ps -ef
    root   169   167  0   Apr 04 ?        0:08 /usr/lib/saf/listen tcp
    root   939     1  0 19:30:47 ?        0:02 /usr/lib/lpsched
    root   859   858  0 19:18:54 term/a   0:01 /bin/sh -c \ /etc/lp
/interfaces/luna
luna-294 rocket!smith "passwd\n##
#
```

In the output from this command, look for the printer port entry. In the above example, port `/dev/term/a` is set up incorrectly as a login terminal. You can tell by the `"passwd\n##` information at the end of the line. If the port is set correctly, skip the last steps in this procedure.

b. Cancel the print request(s).

```
# cancel request-id
```

In this command, *request-id* is the request ID number for a print request to be canceled.

c. Set the printer port to be a nonlogin device.

```
# lpadmin -p printer-name -h
```

d. Check the `ps -ef` command output to verify that the printer port is no longer a login device.

If you do not find the source of the printing problem in the basic LP print service functions, continue to one of the following procedures for the specific client/server case that applies.

▼ How to Check Printing From a Solaris Print Client to a Solaris Print Server

Steps

1. **Check the basic functions of the LP print service on the print server, if you have not done so already.**

   For instructions on checking basic functions, see "How to Check the Basic Functions of the LP Print Service" on page 438. Make sure that the printer works locally before trying to figure out why nothing prints when a request is made from a print client.

2. **Check the basic functions of the LP print service on the print client, if you have not done so already.**

   On the print client, the printer has to be enabled and accepting requests before any request from the client will print.

Note – For most of the following steps, you must be logged in as root or `lp`.

---

3. **Make sure that the print server is accessible.**

   a. **On the print client, send an "are you there?" request to the print server.**

   print_client# **`ping`** *print_server*

   If you receive the message *print_server* `not available`, you might have a network problem.

4. **Verify that the print server is operating properly.**

```
# lpstat -t luna
scheduler is running
system default destination: luna
device for luna: /dev/term/a
luna accepting requests since Jan 12 14:26 2004
printer luna now printing luna-314. enabled since Jan 12 14:26 2004.
available.
luna-129            root                488   Jan 12 14:32
#
```

   The above example shows a print server up and running.

5. **If the print server is not operating properly, go back to step 1.**

▼ How to Troubleshoot Incorrect Output

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role.**

2. **Make sure that the printer type is correct.**

An incorrect printer type might cause incorrect output. For example, if you specify printer type PS and the pages print in reverse order, try printer type PSR. (These type names must be in uppercase.) Also, an incorrect printer type might cause missing text, illegible text, or text with the wrong font. To determine the printer type, examine the entries in the terminfo database. For information on the structure of the terminfo database, see "Printer Type" on page 51.

a. **On the print server, display the printer's characteristics.**

```
$ lpstat -p luna -l
printer luna is idle. enabled since Thu Jul 12 15:02:32 ...
        Form mounted:
        Content types: postscript
        Printer types: PS
        Description:
        Connection: direct
        Interface: /usr/lib/lp/model/standard
        After fault: continue
        Users allowed:
                (all)
        Forms allowed:
                (none)
        Banner not required
        Character sets:

        Default pitch:
        Default page size: 80 wide 66 long
        Default port settings:
$
```

b. **Consult the printer manufacturer's documentation to determine the printer model.**

c. **If the printer type is not correct, change it with Solaris Print Manager's Modify Printer Properties window, or use the following lpadmin command.**

```
# lpstat -p printer-name -T printer-type
```

On the print client, the printer type should be unknown. On the print server, the printer type must match a terminfo entry that is defined to support the model of printer you have. If there is no terminfo entry for the type of printer you have, see "How to Add a terminfo Entry for an Unsupported Printer" on page 152.

3. **If the banner page prints, but there is no output for the body of the document, check the file content types.**

   File content types specified for a printer indicate the types of files the printer can print directly without filtering. An incorrect file content type causes filtering to be bypassed when it might be needed.

   a. **Note the information on file content type that was supplied in the previous step by the lpstat command.**

   On the print client, the file content type should be any, unless you have good reason to specify one or more explicit content types. If a content is specified on the client, filtering is done on the print client, rather than the print server. In addition, content types on the client must match the content types specified on the print server, which in turn must reflect the capabilities of the printer.

b. **Consult your printer manufacturer's documentation to determine which types of files the printer can print directly.**

   The names you use to refer to these types of files do not have to match the names used by the manufacturer. However, the names you use must agree with the names used by the filters known to the LP print service.

c. **If the file content type is not correct, change it with Solaris Print Manager's Modify Printer Properties window, or the following `lpadmin` command.**

   ```
   # lpadmin -p printer-name -I file-content-type(s)
   ```

   Run this command on either the print client, or print server, or both, as needed. Try `-I any` on the print client, and `-I ""` on the print server. The latter specifies a null file content type list, which means an attempt should be made to filter all files, because the printer can directly print only files that exactly match its printer type.

   This combination is a good first choice when files are not printing. If it works, you might want to try specifying explicit content types on the print server to reduce unnecessary filtering. For a local PostScript printer, you should use `postscript`, or `postscript,simple`, if the printer supports these types. Be aware that `PS` and `PSR` are not file content types; they are printer types.

   If you omit `-I`, the file content list defaults to `simple`. If you use the `-I` option and want to specify file content types in addition to `simple`, `simple` must be included in the list.

   When specifying multiple file content types, separate the names with commas. Or you can separate names with spaces and enclose the list in quotation marks. If you specify `any` as the file content type, no filtering will be done and only file types that can be printed directly by the printer should be sent to it.

4. **Check that the print request does not bypass filtering needed to download fonts.**

   If a user submits a print request to a PostScript printer with the `lp -T PS` command, no filtering is done. Try submitting the request with the `lp -T postscript` command to force filtering, which might result in the downloading of non-resident fonts needed by the document.

5. **Make sure that the `stty` settings for the printer port are correct.**

   a. **Read the printer documentation to determine the correct `stty` settings for the printer port.**

   Note – If a printer is connected to a parallel or USB port, the baud setting is irrelevant.

   ---

   b. **Examine the current settings by using the `stty` command.**

```
# stty -a < /dev/term/a
speed 9600 baud;
rows = 0; columns = 0; ypixels = 0; xpixels = 0;
eucw 1:0:0:0, scrw 1:0:0:0
intr = ^c; quit = ^|; erase = ^?; kill = ^u;
eof = ^d; eol = <undef>; eol2 = <undef>; swtch = <undef>;
start = ^q; stop = ^s; susp = ^z; dsusp = ^y;
rprnt = ^r; flush = ^o; werase = ^w; lnext = ^v;
parenb -parodd cs7 -cstopb -hupcl cread -clocal -loblk -parext
-ignbrk brkint -ignpar -parmrk -inpck istrip -inlcr -igncr icrnl -iuclc
ixon -ixany -ixoff imaxbel
isig icanon -xcase echo echoe echok -echonl -noflsh
-tostop echoctl -echoprt echoke -defecho -flusho -pendin iexten
opost -olcuc onlcr -ocrnl -onocr -onlret -ofill -ofdel tab3
#
```

This command shows the current `stty` settings for the printer port.

The following table shows the default `stty` options used by the LP print service's standard printer interface program.

Option	Meaning
`-9600`	Set baud rate to 9600
`-cs8`	Set 8-bit bytes
`-cstopb`	Send one stop bit per byte
`-parity`	Do not generate parity
`-ixon`	Enable XON/XOFF (also known as START/STOP or DC1/DC3)
`-opost`	Do "output post-processing" using all the settings that follow in this table
`-olcuc`	Do not map lowercase to uppercase
`-onlcr`	Change line feed to carriage return/line feed
`-ocrnl`	Do not change carriage returns into line feeds
`-onocr`	Output carriage returns even at column 0
`-n10`	No delay after line feeds
`-cr0`	No delay after carriage returns
`-tab0`	No delay after tabs
`-bs0`	No delay after backspaces
`-vt0`	No delay after vertical tabs
`-ff0`	No delay after form feeds

c. **Change the `stty` settings.**

```
# lpadmin -p printer-name -o "stty= options"
```

Use the following table to choose `stty` options to correct various problems affecting print output.

`stty` Values	Result	Possible Problem From Incorrect Setting
`110`, `300`, `600`, `1200`, `1800`, `2400`, `4800`, `9600`, `19200`, `38400`	Sets baud rate to the specified value (enter only one baud rate)	Random characters and special characters might be printed and spacing might be inconsistent
`oddp`	Sets odd parity	Missing or incorrect characters appear randomly
`evenp`	Sets even parity	
`-parity`	Sets no parity	
`-tabs`	Sets no tabs	Text is jammed against right margin
`tabs`	Sets tabs every eight spaces	Text has no left margin, is run together, or is jammed together
`-onlcr`	Sets no carriage return at the beginning of line(s)	Incorrect double spacing
`onlcr`	Sets carriage return at beginning of line(s)	The print zigzags down the page

You can change more than one option setting by enclosing the list of options in single quotation marks and separating each option with spaces. For example, suppose the printer requires you to enable odd parity and set a 7-bit character size. You would type a command similar to that shown in the following example:

```
# lpadmin -p neptune -o "stty='parenb parodd cs7'"
```

The `stty` option `parenb` enables parity checking/generation, `parodd` sets odd parity generation, and `cs7` sets the character size to 7 bits.

6. **Verify that the document prints correctly.**

```
# lp -d printer-name filename
```

▼ How to Unhang the LP Print Service

Steps

1. **Log in as superuser, `lp`, or assume an equivalent role.**
2. **Stop the LP print service.**

   ```
   # svcadm disable /application/print/server
   ```

3. **Restart the LP print service.**

   ```
   # svcadm enable /application/print/server
   ```

   The LP print service should restart. If you are having trouble restarting the scheduler, see "How to Restart the Print Scheduler" on page 90.

How to Troubleshoot an Idle (Hung) Printer

This task includes a number of procedures to use when a printer appears idle but it should not be. It makes sense to try the procedures in order, but the order is not mandatory.

▼ How to Check That the Printer is Ready to Print

Steps

1. **Display printer status information.**

   ```
   # lpstat -p printer-name
   ```

   The information displayed shows you whether the printer is idle or active, enabled or disabled, or available or not accepting print requests. If everything looks all right, continue with other procedures in this section. If you cannot run the `lpstat` command, see "How to Unhang the LP Print Service" on page 447.

2. **If the printer is not available (not accepting requests), allow the printer to accept requests.**

   ```
   # accept printer-name
   ```

   The printer begins to accept requests into its print queue.

3. **If the printer is disabled, re-enable it.**

   ```
   # enable printer-name
   ```

   This command re-enables the printer so that it will act on the requests in its queue.

How to Check for Print Filtering

Check for print filtering by using the `lpstat -o` command.

```
$ lpstat -o luna
luna-10          fred         1261   Mar 12 17:34 being filtered
luna-11          iggy         1261   Mar 12 17:36 on terra
luna-12          jack         1261   Mar 12 17:39 on terra
$
```

See if the first waiting request is being filtered. If the output looks like the above example, the file is being filtered; the printer is not hung, it just is taking a while to process the request.

▼ How to Resume Printing After a Printer Fault

Steps

1. **Look for a message about a printer fault and try to correct the fault if there is one.**

   Depending on how printer fault alerts have been specified, messages might be sent to root by email or written to a terminal on which root is logged in.

2. **Re-enable the printer.**

   ```
   # enable printer-name
   ```

   If a request was blocked by a printer fault, this command will force a retry. If this command does not work, continue with other procedures in this section.

▼ How to Send Print Requests to a Remote Printer When They Back Up in the Local Queue

Steps

1. **On the print client, stop further queuing of print requests to the print server.**

   ```
   # reject printer-name
   ```

2. **On the print client, send an "are you there?" request to the print server.**

   ```
   print_client# ping print_server
   print_server is alive
   ```

   If you receive the message *print_server* `not available`, you might have a network problem.

3. **After you fix the above problem, allow new print requests to be queued.**

   ```
   # accept printer-name
   ```

4. **If necessary, re-enable the printer.**

   ```
   # enable printer-name
   ```

▼ How to Free Print Requests From a Print Client That Back Up in the Print Server Queue

Steps

1. **On the print server, stop further queuing of print requests from any print client to the print server.**

   ```
   # reject printer-name
   ```

2. **Display the `lpsched` log file.**

   ```
   # more /var/lp/logs/lpsched
   ```

   The information displayed might help you pinpoint what is preventing the print requests from the print client to the print server from being printed.

3. **After you fix the problem, allow new print requests to be queued.**

   ```
   # accept printer-name
   ```

4. **If necessary, re-enable the printer on the print server.**

   ```
   # enable printer-name
   ```

▼ How to Resolve Conflicting Printer Status Messages

Steps

1. **On the print server, verify the printer is enabled and is accepting requests.**

   ```
   # lpstat -p printer-name
   ```

   Users will see conflicting status messages when the print client is accepting requests, but the print server is rejecting requests.

2. **On the print server, check that the definition of the printer on the print client matches the definition of the printer on the print server.**

   ```
   # lpstat -p -l printer-name
   ```

   Look at the definitions of the print job components, like print filters, character sets, print wheels, and forms, to be sure they are the same on both the client and server systems so that local users can access printers on print server systems.

CHAPTER 28

Resolving UFS File System Inconsistencies (Tasks)

This chapter describes the `fsck` error messages and the possible responses you can make to resolve the error messages.

This is a list of the information in this chapter.

For information about the `fsck` command and how to use it to check file system integrity, see Chapter 21, "Checking UFS File System Consistency (Tasks)," in *System Administration Guide: Devices and File Systems*.

`fsck` Error Messages

Normally, the `fsck` command is run non-interactively to *preen* the file systems after an abrupt system halt in which the latest file system changes were not written to disk. Preening automatically fixes any basic file system inconsistencies and does not try to repair more serious errors. While preening a file system, the `fsck` command fixes the inconsistencies it expects from such an abrupt halt. For more serious conditions, the command reports the error and terminates.

When you run the fsck command interactively, it reports each inconsistency found and fixes innocuous errors. However, for more serious errors, the command reports the inconsistency and prompts you to choose a response. When you run the fsck command with the -y or -n options, your response is predefined as yes or no to the default response suggested by the fsck command for each error condition.

Some corrective actions will result in some loss of data. The amount and severity of data loss might be determined from the fsck diagnostic output.

The fsck command is a multipass file system check program. Each pass invokes a different phase of the fsck command with different sets of messages. After initialization, the fsck command performs successive passes over each file system, checking blocks and sizes, path names, connectivity, reference counts, and the map of free blocks (possibly rebuilding it). It also performs some cleanup.

The phases (passes) performed by the UFS version of the fsck command are:

- Initialization
- Phase 1 – Check blocks and sizes
- Phase 2 – Check path names
- Phase 3 – Check connectivity
- Phase 4 – Check reference counts
- Phase 5 – Check cylinder groups

The next sections describe the error conditions that might be detected in each phase, the messages and prompts that result, and possible responses you can make.

Messages that might appear in more than one phase are described in "General fsck Error Messages" on page 453. Otherwise, messages are organized alphabetically by the phases in which they occur.

The following table lists many of the abbreviations included in the fsck error messages.

TABLE 28–1 Error Message Abbreviations

Abbreviation	Meaning
BLK	Block number
DUP	Duplicate block number
DIR	Directory name
CG	Cylinder group
MTIME	Time file was last modified
UNREF	Unreferenced

Many of the messages also include variable fields, such as inode numbers, which are represented in this book by an italicized term, such as *inode-number*. For example, this screen message:

INCORRECT BLOCK COUNT I=2529

is shown as follows:

INCORRECT BLOCK COUNT I=*inode-number*

General fsck Error Messages

The error messages in this section might be displayed in any phase after initialization. Although they offer the option to continue, it is generally best to regard them as fatal. They reflect a serious system failure and should be handled immediately. When confronted with such a message, terminate the program by entering n(o). If you cannot determine what caused the problem, contact your local service provider or another qualified person.

CANNOT SEEK: BLK *block-number* (CONTINUE)

Cause
: A request to move to a specified block number, *block-number*, in the file system failed. This message indicates a serious problem, probably a hardware failure.

  If you want to continue the file system check, fsck will retry the move and display a list of sector numbers that could not be moved. If the block was part of the virtual memory buffer cache, fsck will terminate with a fatal I/O error message.

Action
: If the disk is experiencing hardware problems, the problem will persist. Run fsck again to recheck the file system.

  If the recheck fails, contact your local service provider or another qualified person.

CANNOT READ: BLK *block-number* (CONTINUE)

Cause
: A request to read a specified block number, *block-number*, in the file system failed. The message indicates a serious problem, probably a hardware failure.

  If you want to continue the file system check, fsck will retry the read and display a list of sector numbers that could not be read. If the block was part of the virtual memory buffer cache, fsck will terminate with a fatal I/O error message. If fsck tries to write back one of the blocks on which the read failed, it will display the following message:

  WRITING ZERO'ED BLOCK *sector-numbers*TO DISK

Action
: If the disk is experiencing hardware problems, the problem will persist. Run fsck again to recheck the file system. If the recheck fails, contact your local service provider or another qualified person.

CANNOT WRITE: BLK *block-number* (CONTINUE)

Cause
A request to write a specified block number, *block-number*, in the file system failed.

If you continue the file system check, `fsck` will retry the write and display a list of sector numbers that could not be written. If the block was part of the virtual memory buffer cache, `fsck` will terminate with a fatal I/O error message.

Action
The disk might be write-protected. Check the write-protect lock on the drive. If the disk has hardware problems, the problem will persist. Run `fsck` again to recheck the file system. If the write-protect is not the problem or the recheck fails, contact your local service provider or another qualified person.

Initialization Phase `fsck` Messages

In the initialization phase, command-line syntax is checked. Before the file system check can be performed, `fsck` sets up tables and opens files.

The messages in this section relate to error conditions resulting from command-line options, memory requests, the opening of files, the status of files, file system size checks, and the creation of the scratch file. All such initialization errors terminate `fsck` when it is preening the file system.

`Can't roll the log for` *device-name*`.`

```
DISCARDING THE LOG MAY DISCARD PENDING TRANSACTIONS.
DISCARD THE LOG AND CONTINUE?
```

Cause
`fsck` was unable to flush the transaction log of a logging UFS file system prior to checking the file system for errors.

Action
Answering yes means the file system operations that were in the log, but had not been applied to the file system, are lost. In this case, `fsck` runs the same checks it always runs and asks the following question in phase 5:

```
FREE BLK COUNT(S) WRONG IN SUPERBLK (SALVAGE)
```

Answering yes at this point reclaims the blocks that were used for the log. The next time the file system is mounted with logging enabled, the log will be recreated.

Answering no preserves the log and exits, but the file system isn't mountable.

`bad inode number` *inode-number* `to ginode`

Cause
An internal error occurred because of a nonexistent inode *inode-number*. `fsck` exits.

Action
Contact your local service provider or another qualified person.

`cannot alloc` *size-of-block map* `bytes for blockmap`
`cannot alloc` *size-of-free map* `bytes for freemap`
`cannot alloc` *size-of-state map* `bytes for statemap`
`cannot alloc` *size-of-lncntp* `bytes for lncntp`

Cause
: Request for memory for its internal tables failed. `fsck` terminates. This message indicates a serious system failure that should be handled immediately. This condition might occur if other processes are using a very large amount of system resources.

Action
: Killing other processes might solve the problem. If not, contact your local service provider or another qualified person.

`Can't open checklist file:` *filename*

Cause
: The file system checklist file *filename* (usually `/etc/vfstab`) cannot be opened for reading. `fsck` terminates.

Action
: Check if the file exists and if its access modes permit read access.

`Can't open` *filename*

Cause
: `fsck` cannot open file system *filename*. When running interactively, `fsck` ignores this file system and continues checking the next file system given.

Action
: Check to see if read and write access to the raw device file for the file system is permitted.

`Can't stat root`

Cause
: `fsck` request for statistics about the root directory failed. `fsck` terminates.

Action
: This message indicates a serious system failure. Contact your local service provider or another qualified person.

`Can't stat` *filename*
`Can't make sense out of name` *filename*

Cause
: `fsck` request for statistics about the file system *filename* failed. When running interactively, `fsck` ignores this file system and continues checking the next file system given.

Action
: Check if the file system exists and check its access modes.

filename`: (NO WRITE)`

Cause

Either the `-n` option was specified or `fsck` could not open the file system *filename* for writing. When `fsck` is running in no-write mode, all diagnostic messages are displayed, but `fsck` does not attempt to fix anything.

Action

If `-n` was not specified, check the type of the file specified. It might be the name of a regular file.

`IMPOSSIBLE MINFREE=`*percent* `IN SUPERBLOCK (SET TO DEFAULT)`

Cause

The superblock minimum space percentage is greater than 99 percent or less than 0 percent.

Action

To set the `minfree` parameter to the default 10 percent, type `y` at the default prompt. To ignore the error condition, type `n` at the default prompt.

filename`: BAD SUPER BLOCK:` *message*
```
USE AN ALTERNATE SUPER-BLOCK TO SUPPLY NEEDED INFORMATION;
e.g., fsck[-f ufs] -o b=# [special ...]
where # is the alternate superblock.  See fsck_ufs(1M)
```

Cause

The superblock has been corrupted.

Action

One of the following messages might be displayed:

```
CPG OUT OF RANGE
FRAGS PER BLOCK OR FRAGSIZE WRONG
INODES PER GROUP OUT OF RANGE
INOPB NONSENSICAL RELATIVE TO BSIZE
MAGIC NUMBER WRONG
NCG OUT OF RANGE
NCYL IS INCONSISTENT WITH NCG*CPG
NUMBER OF DATA BLOCKS OUT OF RANGE
NUMBER OF DIRECTORIES OUT OF RANGE
ROTATIONAL POSITION TABLE SIZE OUT OF RANGE
SIZE OF CYLINDER GROUP SUMMARY AREA WRONG
SIZE TOO LARGE
BAD VALUES IN SUPERBLOCK
```

Try to rerun `fsck` with an alternative superblock. Specifying block 32 is a good first choice. You can locate an alternative copy of the superblock by running the `newfs -N` command on the slice. Be sure to specify the `-N` option; otherwise, `newfs` overwrites the existing file system.

`UNDEFINED OPTIMIZATION IN SUPERBLOCK (SET TO DEFAULT)`

Cause

The superblock optimization parameter is neither `OPT_TIME` nor `OPT_SPACE`.

Action
: To minimize the time to perform operations on the file system, type `y` at the `SET TO DEFAULT` prompt. To ignore this error condition, type `n`.

Phase 1: Check Blocks and Sizes Messages

This phase checks the inode list. It reports error conditions encountered while:

- Checking inode types
- Setting up the zero-link-count table
- Examining inode block numbers for bad or duplicate blocks
- Checking inode size
- Checking inode format

All errors in this phase except `INCORRECT BLOCK COUNT`, `PARTIALLY TRUNCATED INODE`, `PARTIALLY ALLOCATED INODE`, and `UNKNOWN FILE TYPE` terminate `fsck` when it is preening a file system.

These messages (in alphabetical order) might occur in phase 1:

block-number `BAD I=`*inode-number*

Cause
: Inode *inode-number* contains a block number *block-number* with a number lower than the number of the first data block in the file system or greater than the number of the last block in the file system. This error condition might generate the `EXCESSIVE BAD BLKS` error message in phase 1 if inode *inode-number* has too many block numbers outside the file system range. This error condition generates the `BAD/DUP` error message in phases 2 and 4.

Action
: N/A

`BAD MODE: MAKE IT A FILE?`

Cause
: The status of a given inode is set to all 1s, indicating file system damage. This message does not indicate physical disk damage, unless it is displayed repeatedly after `fsck -y` has been run.

Action
: Type `y` to reinitialize the inode to a reasonable value.

`BAD STATE` *state-number* `TO BLKERR`

Cause
: An internal error has scrambled the `fsck` state map so that it shows the impossible value *state-number*. `fsck` exits immediately.

Action
: Contact your local service provider or another qualified person.

block-number `DUP I=`*inode-number*

Cause

Inode *inode-number* contains a block number *block-number*, which is already claimed by the same or another inode. This error condition might generate the `EXCESSIVE DUP BLKS` error message in phase 1 if inode *inode-number* has too many block numbers claimed by the same or another inode. This error condition invokes phase 1B and generates the `BAD/DUP` error messages in phases 2 and 4.

Action

N/A

`DUP TABLE OVERFLOW (CONTINUE)`

Cause

There is no more room in an internal table in `fsck` containing duplicate block numbers. If the `-o p` option is specified, the program terminates.

Action

To continue the program, type `y` at the `CONTINUE` prompt. When this error occurs, a complete check of the file system is not possible. If another duplicate block is found, this error condition repeats. Increase the amount of virtual memory available (by killing some processes, increasing swap space) and run `fsck` again to recheck the file system. To terminate the program, type `n`.

`EXCESSIVE BAD BLOCKS I=inode-number (CONTINUE)`

Cause

Too many (usually more than 10) blocks have a number lower than the number of the first data block in the file system or greater than the number of the last block in the file system associated with inode *inode-number*. If the `-o p` (preen) option is specified, the program terminates.

Action

To continue the program, type `y` at the `CONTINUE` prompt. When this error occurs, a complete check of the file system is not possible. You should run `fsck` again to recheck the file system. To terminate the program, type `n`.

`EXCESSIVE DUP BLKS I=inode-number (CONTINUE)`

Cause

Too many (usually more than 10) blocks are claimed by the same or another inode or by a free-list. If the `-o p` option is specified, the program terminates.

Action

To continue the program, type `y` at the `CONTINUE` prompt. When this error occurs, a complete check of the file system is not possible. You should run `fsck` again to recheck the file system. To terminate the program, type `n`.

`INCORRECT BLOCK COUNT I=inode-number (`*number-of-BAD-DUP-or-missing-blocks*
`should be` *number-of-blocks-in-filesystem*`) (CORRECT)`

Cause

The block count for inode *inode-number* is *number-of-BAD-DUP-or-missing-blocks*, but should be *number-of-blocks-in-filesystem*. When preening, `fsck` corrects the count.

Action
: To replace the block count of inode *inode-number* by *number-of-blocks-in-filesystem*, type `y` at the `CORRECT` prompt. To terminate the program, type `n`.

`LINK COUNT TABLE OVERFLOW (CONTINUE)`

Cause
: There is no more room in an internal table for `fsck` containing allocated inodes with a link count of zero. If the `-o p` (preen) option is specified, the program exits and `fsck` has to be completed manually.

Action
: To continue the program, type `y` at the `CONTINUE` prompt. If another allocated inode with a zero-link count is found, this error condition repeats. When this error occurs, a complete check of the file system is not possible. You should run `fsck` again to recheck the file system. Increase the virtual memory available by killing some processes or increasing swap space, then run `fsck` again. To terminate the program, type `n`.

`PARTIALLY ALLOCATED INODE I=`*inode-number* `(CLEAR)`

Cause
: Inode *inode-number* is neither allocated nor unallocated. If the `-o p` (preen) option is specified, the inode is cleared.

Action
: To deallocate the inode *inode-number* by zeroing out its contents, type `y`. This might generate the `UNALLOCATED` error condition in phase 2 for each directory entry pointing to this inode. To ignore the error condition, type `n`. A no response is appropriate only if you intend to take other measures to fix the problem.

`PARTIALLY TRUNCATED INODE I=`*inode-number* `(SALVAGE)`

Cause
: `fsck` has found inode *inode-number* whose size is shorter than the number of blocks allocated to it. This condition occurs only if the system crashes while truncating a file. When preening the file system, `fsck` completes the truncation to the specified size.

Action
: To complete the truncation to the size specified in the inode, type `y` at the `SALVAGE` prompt. To ignore this error condition, type `n`.

`UNKNOWN FILE TYPE I=`*inode-number* `(CLEAR)`

Cause
: The mode word of the inode *inode-number* shows that the inode is not a pipe, special character inode, special block inode, regular inode, symbolic link, FIFO file, or directory inode. If the `-o p` option is specified, the inode is cleared.

Action
: To deallocate the inode *inode-number* by zeroing its contents, which results in the `UNALLOCATED` error condition in phase 2 for each directory entry pointing to this inode, type `y` at the `CLEAR` prompt. To ignore this error condition, type `n`.

Phase 1B: Rescan for More DUPS Messages

When a duplicate block is found in the file system, this message is displayed:

block-number `DUP I=`*inode-number*

Cause
: Inode *inode-number* contains a block number *block-number* that is already claimed by the same or another inode. This error condition generates the `BAD/DUP` error message in phase 2. Inodes that have overlapping blocks might be determined by examining this error condition and the `DUP` error condition in phase 1.

Action
: When a duplicate block is found, the file system is rescanned to find the inode that previously claimed that block.

Phase 2: Check Path Names Messages

This phase removes directory entries pointing to bad inodes found in phases 1 and 1B. It reports error conditions resulting from:

- Incorrect root inode mode and status
- Directory inode pointers out of range
- Directory entries pointing to bad inodes
- Directory integrity checks

When the file system is being preened (`-o p` option), all errors in this phase terminate `fsck`, except those related to directories not being a multiple of the block size, duplicate and bad blocks, inodes out of range, and extraneous hard links.

These messages (in alphabetical order) might occur in phase 2:

`BAD INODE` *state-number* `TO DESCEND`

Cause
: An `fsck` internal error has passed an invalid state *state-number* to the routine that descends the file system directory structure. `fsck` exits.

Action
: If this error message is displayed, contact your local service provider or another qualified person.

`BAD INODE NUMBER FOR '.' I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode*
`SIZE=`*file-size* `MTIME=`*modification-time* `DIR=`*filename* `(FIX)`

Cause
: A directory *inode-number* has been found whose inode number for "`.`" does not equal *inode-number*.

Action
: To change the inode number for "`.`" to be equal to *inode-number*, type `y` at the `FIX` prompt To leave the inode numbers for "`.`" unchanged, type `n`.

`BAD INODE NUMBER FOR '..' I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size* `MTIME=`*modification-time* `DIR=`*filename* `(FIX)`

Cause
: A directory *inode-number* has been found whose inode number for "`..`" does not equal the parent of *inode-number*.

Action
: To change the inode number for "`..`" to be equal to the parent of *inode-number*, type `y` at the `FIX` prompt. (Note that "`..`" in the root inode points to itself.) To leave the inode number for "`..`" unchanged, type `n`.

`BAD RETURN STATE` *state-number* `FROM DESCEND`

Cause
: An `fsck` internal error has returned an impossible state *state-number* from the routine that descends the file system directory structure. `fsck` exits.

Action
: If this message is displayed, contact your local service provider or another qualified person.

`BAD STATE` *state-number* `FOR ROOT INODE`

Cause
: An internal error has assigned an impossible state *state-number* to the root inode. `fsck` exits.

Action
: If this error message is displayed, contact your local service provider or another qualified person.

`BAD STATE` *state-number* `FOR INODE=`*inode-number*

Cause
: An internal error has assigned an impossible state *state-number* to inode *inode-number*. `fsck` exits.

Action
: If this error message is displayed, contact your local service provider or another qualified person.

`DIRECTORY TOO SHORT I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size* `MTIME=`*modification-time* `DIR=`*filename* `(FIX)`

Cause
: A directory *filename* has been found whose size *file-size* is less than the minimum directory size. The owner *UID*, mode *file-mode*, size *file-size*, modify time *modification-time*, and directory name *filename* are displayed.

Action
: To increase the size of the directory to the minimum directory size, type `y` at the `FIX` prompt. To ignore this directory, type `n`.

`DIRECTORY` *filename*`: LENGTH` *file-size* `NOT MULTIPLE OF` *block-number* `(ADJUST)`

Cause

A directory *filename* has been found with size *file-size* that is not a multiple of the directory block size *block-number*.

Action

To round up the length to the appropriate block size, type `y`. When preening the file system (`-o p` option), `fsck` only displays a warning and adjusts the directory. To ignore this condition, type `n`.

`DIRECTORY CORRUPTED I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode*
`SIZE=`*file-size* `MTIME=`*modification-time* `DIR=`*filename* `(SALVAGE)`

Cause

A directory with an inconsistent internal state has been found.

Action

To throw away all entries up to the next directory boundary (usually a 512-byte boundary), type `y` at the `SALVAGE` prompt. This drastic action can throw away up to 42 entries. Take this action only after other recovery efforts have failed. To skip to the next directory boundary and resume reading, but not modify the directory, type `n`.

`DUP/BAD I=`*inode-number* `OWNER=O MODE=M SIZE=`*file-size*
`MTIME=`*modification-time* `TYPE=`*filename* `(REMOVE)`

Cause

Phase 1 or phase 1B found duplicate blocks or bad blocks associated with directory or file entry *filename*, inode *inode-number*. The owner *UID*, mode *file-mode*, size *file-size*, modification time *modification-time*, and directory or file name *filename* are displayed. If the `-p` (preen) option is specified, the duplicate/bad blocks are removed.

Action

To remove the directory or file entry *filename*, type `y` at the `REMOVE` prompt. To ignore this error condition, type `n`.

`DUPS/BAD IN ROOT INODE (REALLOCATE)`

Cause

Phase 1 or phase 1B has found duplicate blocks or bad blocks in the root inode (usually inode number 2) of the file system.

Action

To clear the existing contents of the root inode and reallocate it, type `y` at the `REALLOCATE` prompt. The files and directories usually found in the root inode will be recovered in phase 3 and put into the `lost+found` directory. If the attempt to allocate the root fails, `fsck` will exit with: `CANNOT ALLOCATE ROOT INODE`. Type `n` to get the `CONTINUE` prompt. Type: `y` to respond to the `CONTINUE` prompt, and ignore the `DUPS/BAD` error condition in the root inode and continue running the file system check. If the root inode is not correct, this might generate many other error messages. Type `n` to terminate the program.

`EXTRA '.' ENTRY I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode*
`SIZE=`*file-size* `MTIME=`*modification-time* `DIR=`*filename* `(FIX)`

Cause
: A directory *inode-number* has been found that has more than one entry for "`.`".

Action
: To remove the extra entry for "`.`" type `y` at the `FIX` prompt. To leave the directory unchanged, type `n`.

`EXTRA '..' ENTRY I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size* `MTIME=`*modification-time* `DIR=`*filename*`(FIX)`

Cause
: A directory *inode-number* has been found that has more than one entry for "`..`" (the parent directory).

Action
: To remove the extra entry for '`..`' (the parent directory), type `y` at the `FIX` prompt. To leave the directory unchanged, type `n`.

hard-link-number `IS AN EXTRANEOUS HARD LINK TO A DIRECTORY` *filename* `(REMOVE)`

Cause
: `fsck` has found an extraneous hard link *hard-link-number* to a directory *filename*. When preening (`-o p` option), `fsck` ignores the extraneous hard links.

Action
: To delete the extraneous entry *hard-link-number* type `y` at the `REMOVE` prompt. To ignore the error condition, type `n`.

inode-number `OUT OF RANGE I=`*inode-number* `NAME=`*filename* `(REMOVE)`

Cause
: A directory entry *filename* has an inode number *inode-number* that is greater than the end of the inode list. If the `-p` (preen) option is specified, the inode will be removed automatically.

Action
: To delete the directory entry *filename* type `y` at the `REMOVE` prompt. To ignore the error condition, type `n`.

`MISSING '.' I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size* `MTIME=`*modification-time* `DIR=`*filename* `(FIX)`

Cause
: A directory *inode-number* has been found whose first entry (the entry for "`.`") is unallocated.

Action
: To build an entry for "`.`" with inode number equal to *inode-number*, type `y` at the `FIX` prompt. To leave the directory unchanged, type `n`.

`MISSING '.' I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size* `MTIME=`*modification-time* `DIR=`*filename* `CANNOT FIX, FIRST ENTRY IN DIRECTORY CONTAINS` *filename*

Cause
: A directory *inode-number* has been found whose first entry is *filename*. `fsck` cannot resolve this problem.

Action
: If this error message is displayed, contact your local service provider or another qualified person.

`MISSING '.' I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size* `MTIME=`*modification-time* `DIR=`*filename* `CANNOT FIX, INSUFFICIENT SPACE TO ADD '.'`

Cause
: A directory *inode-number* has been found whose first entry is not "`.`". `fsck` cannot resolve the problem.

Action
: If this error message is displayed, contact your local service provider or another qualified person.

`MISSING '..' I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size* `MTIME=`*modification-time* `DIR=`*filename* `(FIX)`

Cause
: A directory *inode-number* has been found whose second entry is unallocated.

Action
: To build an entry for "`..`" with inode number equal to the parent of *inode-number*, type `y` at the `FIX` prompt. (Note that "`..`" in the root inode points to itself.) To leave the directory unchanged, type `n`.

`MISSING '..' I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size* `MTIME=`*modification-time* `DIR=`*filename* `CANNOT FIX, SECOND ENTRY IN DIRECTORY CONTAINS` *filename*

Cause
: A directory *inode-number* has been found whose second entry is *filename*. `fsck` cannot resolve this problem.

Action
: If this error message is displayed, contact your local service provider or another qualified person.

`MISSING '..' I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size* `MTIME=`*modification-time* `DIR=`*filename* `CANNOT FIX, INSUFFICIENT SPACE TO ADD '..'`

Cause
: A directory *inode-number* has been found whose second entry is not "`..`" (the parent directory). `fsck` cannot resolve this problem.

Action
: If this error message is displayed, contact your local service provider or another qualified person.

`NAME TOO LONG` *filename*

Cause

An excessively long path name has been found, which usually indicates loops in the file system name space. This error can occur if a privileged user has made circular links to directories.

Action

Remove the circular links.

`ROOT INODE UNALLOCATED (ALLOCATE)`

Cause

The root inode (usually inode number 2) has no allocate-mode bits.

Action

To allocate inode 2 as the root inode, type `y` at the `ALLOCATE` prompt. The files and directories usually found in the root inode will be recovered in phase 3 and put into the `lost+found` directory. If the attempt to allocate the root inode fails, `fsck` displays this message and exits: `CANNOT ALLOCATE ROOT INODE`. To terminate the program, type `n`.

`ROOT INODE NOT DIRECTORY (REALLOCATE)`

Cause

The root inode (usually inode number 2) of the file system is not a directory inode.

Action

To clear the existing contents of the root inode and reallocate it, type `y` at the `REALLOCATE` prompt. The files and directories usually found in the root inode will be recovered in phase 3 and put into the `lost+found` directory. If the attempt to allocate the root inode fails, `fsck` displays this message and exits: `CANNOT ALLOCATE ROOT INODE`. To have `fsck` prompt with `FIX`, type `n`.

`UNALLOCATED I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size* `MTIME=`*modification-time* `type=`*filename*`(REMOVE)`

Cause

A directory or file entry *filename* points to an unallocated inode *inode-number*. The owner *UID*, mode *file-mode*, size *file-size*, modify time *modification-time*, and file name *filename* are displayed.

Action

To delete the directory entry *filename*, type `y` at the `REMOVE` prompt. To ignore the error condition, type `n`.

`ZERO LENGTH DIRECTORY I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size* `MTIME=`*modification-time* `DIR=`*filename* `(REMOVE)`

Cause

A directory entry *filename* has a size *file-size* that is zero. The owner UID, mode *file-mode*, size *file-size*, modify time *modification-time*, and directory name *filename* are displayed.

Action

To remove the directory entry *filename*, type `y` at the `REMOVE` prompt. This results in the `BAD/DUP` error message in phase 4. To ignore the error condition, type `n`.

Phase 3: Check Connectivity Messages

This phase checks the directories examined in phase 2 and reports error conditions resulting from:

- Unreferenced directories
- Missing or full `lost+found` directories

These messages (in alphabetical order) might occur in phase 3:

`BAD INODE` *state-number* `TO DESCEND`

Cause
: An internal error has caused an impossible state *state-number* to be passed to the routine that descends the file system directory structure. `fsck` exits.

Action
: If this occurs, contact your local service provider or another qualified person.

`DIR I=`*inode-number1* `CONNECTED. PARENT WAS I=`*inode-number2*

Cause
: This is an advisory message indicating a directory inode *inode-number1* was successfully connected to the `lost+found` directory. The parent inode *inode-number2* of the directory inode *inode-number1* is replaced by the inode number of the `lost+found` directory.

Action
: N/A

`DIRECTORY` *filename* `LENGTH` *file-size* `NOT MULTIPLE OF` *block-number* `(ADJUST)`

Cause
: A directory *filename* has been found with size *file-size* that is not a multiple of the directory block size B. (This condition can recur in phase 3 if it is not adjusted in phase 2.)

Action
: To round up the length to the appropriate block size, type `y` at the `ADJUST` prompt. When preening, `fsck` displays a warning and adjusts the directory. To ignore this error condition, type `n`.

`lost+found IS NOT A DIRECTORY (REALLOCATE)`

Cause
: The entry for `lost+found` is not a directory.

Action
: To allocate a directory inode and change the `lost+found` directory to reference it, type `y` at the `REALLOCATE` prompt. The previous inode reference by the `lost+found` directory is not cleared and it will either be reclaimed as an unreferenced inode or have its link count adjusted later in this phase. Inability to create a `lost+found` directory displays the message: `SORRY. CANNOT CREATE lost+found DIRECTORY` and aborts the attempt to link up the lost inode, which

generates the UNREF error message in phase 4. To abort the attempt to link up the lost inode, which generates the UNREF error message in phase 4, type n.

NO lost+found DIRECTORY (CREATE)

Cause
: There is no lost+found directory in the root directory of the file system. When preening, fsck tries to create a lost+found directory.

Action
: To create a lost+found directory in the root of the file system, type y at the CREATE prompt. This might lead to the message NO SPACE LEFT IN / (EXPAND). If the lost+found directory cannot be created, fsck displays the message: SORRY. CANNOT CREATE lost+found DIRECTORY and aborts the attempt to link up the lost inode. This in turn generates the UNREF error message later in phase 4. To abort the attempt to link up the lost inode, type n.

NO SPACE LEFT IN /lost+found (EXPAND)

Cause
: Another entry cannot be added to the lost+found directory in the root directory of the file system because no space is available. When preening, fsck expands the lost+found directory.

Action
: To expand the lost+found directory to make room for the new entry, type y at the EXPAND prompt. If the attempted expansion fails, fsck displays: SORRY. NO SPACE IN lost+found DIRECTORY and aborts the request to link a file to the lost+found directory. This error generates the UNREF error message later in phase 4. Delete any unnecessary entries in the lost+found directory. This error terminates fsck when preening is in effect. To abort the attempt to link up the lost inode, type n.

UNREF DIR I=*inode-number* OWNER=*UID* MODE=*file-mode* SIZE=*file-size* MTIME=*modification-time* (RECONNECT)

Cause
: The directory inode *inode-number* was not connected to a directory entry when the file system was traversed. The owner *UID*, mode *file-mode*, size *file-size*, and modification time *modification-time* of directory inode *inode-number* are displayed. When preening, fsck reconnects the non-empty directory inode if the directory size is non-zero. Otherwise, fsck clears the directory inode.

Action
: To reconnect the directory inode *inode-number* into the lost+found directory, type y at the RECONNECT prompt. If the directory is successfully reconnected, a CONNECTED message is displayed. Otherwise, one of the lost+found error messages is displayed. To ignore this error condition, type n. This error causes the UNREF error condition in phase 4.

Phase 4: Check Reference Counts Messages

This phase checks the link count information obtained in phases 2 and 3. It reports error conditions resulting from:

- Unreferenced files
- A missing or full `lost+found` directory
- Incorrect link counts for files, directories, symbolic links, or special files
- Unreferenced files, symbolic links, and directories
- Bad or duplicate blocks in files and directories
- Incorrect total free-inode counts

All errors in this phase (except running out of space in the `lost+found` directory) are correctable when the file system is being preened.

These messages (in alphabetical order) might occur in phase 4:

`BAD/DUP` *type* `I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size*
`MTIME=`*modification-time* `(CLEAR)`

Cause
: Phase 1 or phase 1B found duplicate blocks or bad blocks associated with file or directory inode *inode-number*. The owner *UID*, mode *file-mode*, size *file-size*, and modification time *modification-time* of inode *inode-number* are displayed.

Action
: To deallocate inode *inode-number* by zeroing its contents, type `y` at the `CLEAR` prompt. To ignore this error condition, type `n`.

`(CLEAR)`

Cause
: The inode mentioned in the `UNREF` error message immediately preceding cannot be reconnected. This message does not display if the file system is being preened because lack of space to reconnect files terminates `fsck`.

Action
: To deallocate the inode by zeroing out its contents, type `y` at the `CLEAR` prompt. To ignore the preceding error condition, type `n`.

`LINK COUNT` *type* `I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode*
`SIZE=`*file-size*
`MTIME=`*modification-time* `COUNT` *link-count* `SHOULD BE`
`corrected-link-count (ADJUST)`

Cause
: The link count for directory or file inode *inode-number* is *link-count* but should be *corrected-link-count*. The owner *UID*, mode *file-mode*, size *file-size*, and modification time *modification-time* of inode *inode-number* are displayed. If the `-o p` option is specified, the link count is adjusted unless the number of references is increasing. This condition does not occur unless there is a hardware failure. When the number of references is increasing during preening, `fsck` displays this message and exits: `LINK COUNT INCREASING`

Action

To replace the link count of directory or file inode *inode-number* with *corrected-link-count*, type `y` at the `ADJUST` prompt. To ignore this error condition, type `n`.

`lost+found IS NOT A DIRECTORY (REALLOCATE)`

Cause

The entry for `lost+found` is not a directory.

Action

To allocate a directory inode and change the `lost+found` directory to reference it, type `y` at the `REALLOCATE` prompt. The previous inode reference by the `lost+found` directory is not cleared. It will either be reclaimed as an unreferenced inode or have its link count adjusted later in this phase. Inability to create a `lost+found` directory displays this message: `SORRY. CANNOT CREATE lost+found DIRECTORY` and aborts the attempt to link up the lost inode. This error generates the `UNREF` error message later in phase 4. To abort the attempt to link up the lost inode, type `n`.

`NO lost+found DIRECTORY (CREATE)`

Cause

There is no `lost+found` directory in the root directory of the file system. When preening, `fsck` tries to create a `lost+found` directory.

Action

To create a `lost+found` directory in the root of the file system, type `y` at the `CREATE` prompt. If the `lost+found` directory cannot be created, `fsck` displays the message: `SORRY. CANNOT CREATE lost+found DIRECTORY` and aborts the attempt to link up the lost inode. This error in turn generates the `UNREF` error message later in phase 4. To abort the attempt to link up the lost inode, type `n`.

`NO SPACE LEFT IN / lost+found (EXPAND)`

Cause

There is no space to add another entry to the `lost+found` directory in the root directory of the file system. When preening, `fsck` expands the `lost+found` directory.

Action

To expand the `lost+found` directory to make room for the new entry, type `y` at the `EXPAND` prompt. If the attempted expansion fails, `fsck` displays the message: `SORRY. NO SPACE IN lost+found DIRECTORY` and aborts the request to link a file to the `lost+found` directory. This error generates the `UNREF` error message later in phase 4. Delete any unnecessary entries in the `lost+found` directory. This error terminates `fsck` when preening (`-o p` option) is in effect. To abort the attempt to link up the lost inode, type `n`.

`UNREF FILE I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size* `MTIME=`*modification-time* `(RECONNECT)`

Cause

File inode *inode-number* was not connected to a directory entry when the file system was traversed. The owner *UID*, mode *file-mode*, size *file-size*, and modification time *modification-time* of inode *inode-number* are displayed. When `fsck` is preening, the file is cleared if either its size or its link count is zero; otherwise, it is reconnected.

Action

To reconnect inode *inode-number* to the file system in the `lost+found` directory, type `y`. This error might generate the `lost+found` error message in phase 4 if there are problems connecting inode *inode-number* to the `lost+found` directory. To ignore this error condition, type `n`. This error always invokes the `CLEAR` error condition in phase 4.

`UNREF` *type* `I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode* `SIZE=`*file-size*
`MTIME=`*modification-time* `(CLEAR)`

Cause

Inode *inode-number* (whose *type* is directory or file) was not connected to a directory entry when the file system was traversed. The owner *UID*, mode *file-mode*, size *file-size*, and modification time *modification-time* of inode *inode-number* are displayed. When `fsck` is preening, the file is cleared if either its size or its link count is zero; otherwise, it is reconnected.

Action

To deallocate inode *inode-number* by zeroing its contents, type `y` at the `CLEAR` prompt. To ignore this error condition, type `n`.

`ZERO LENGTH DIRECTORY I=`*inode-number* `OWNER=`*UID* `MODE=`*file-mode*
`SIZE=`*file-size* `MTIME=`*modification-time*`(CLEAR)`

Cause

A directory entry *filename* has a size *file-size* that is zero. The owner *UID*, mode *file-mode*, size *file-size*, modification time *modification-time*, and directory name *filename* are displayed.

Action

To deallocate the directory inode *inode-number* by zeroing out its contents, type `y`. To ignore the error condition, type `n`.

Phase 5: Check Cylinder Groups Messages

This phase checks the free-block and used-inode maps. It reports error conditions resulting from:

- Allocated inodes missing from used-inode maps
- Free blocks missing from free-block maps
- Free inodes in the used-inode maps
- Incorrect total free-block count
- Incorrect total used inode count

These messages (in alphabetical order) might occur in phase 5:

`BLK(S) MISSING IN BIT MAPS (SALVAGE)`

Cause
: A cylinder group block map is missing some free blocks. During preening, `fsck` reconstructs the maps.

Action
: To reconstruct the free-block map, type `y` at the `SALVAGE` prompt. To ignore this error condition, type `n`.

`CG` *character-for-command-option*`: BAD MAGIC NUMBER`

Cause
: The magic number of cylinder group *character-for-command-option* is wrong. This error usually indicates that the cylinder group maps have been destroyed. When running interactively, the cylinder group is marked as needing reconstruction. `fsck` terminates if the file system is being preened.

Action
: If this occurs, contact your local service provider or another qualified person.

`FREE BLK COUNT(S) WRONG IN SUPERBLK (SALVAGE)`

Cause
: The actual count of free blocks does not match the count of free blocks in the superblock of the file system. If the `-o p` option was specified, the free-block count in the superblock is fixed automatically.

Action
: To reconstruct the superblock free-block information, type `y` at the `SALVAGE` prompt. To ignore this error condition, type `n`.

`SUMMARY INFORMATION BAD (SALVAGE)`

Cause
: The summary information is incorrect. When preening, `fsck` recomputes the summary information.

Action
: To reconstruct the summary information, type `y` at the `SALVAGE` prompt. To ignore this error condition, type `n`.

Cleanup Phase Messages

Once a file system has been checked, a few cleanup functions are performed. The cleanup phase displays the following status messages.

number-of `files,` *number-of-files*
`used,` *number-of-files* `free (`*number-of* `frags,` *number-of* `blocks,`
percent `fragmentation)`

This message indicates that the file system checked contains *number-of* files using *number-of* fragment-sized blocks, and that there are *number-of* fragment-sized blocks free in the file system. The numbers in parentheses break the free count down into *number-of* free fragments, *number-of* free full-sized blocks, and the *percent* fragmentation.

`***** FILE SYSTEM WAS MODIFIED *****`

This message indicates that the file system was modified by `fsck`. If this file system is mounted or is the current root (`/`) file system, reboot. If the file system is mounted, you might need to unmount it and run `fsck` again; otherwise, the work done by `fsck` might be undone by the in-core copies of tables.

filename `FILE SYSTEM STATE SET TO OKAY`

This message indicates that file system *filename* was marked as stable. Use the `fsck -m` command to determine if the file system needs checking.

filename `FILE SYSTEM STATE NOT SET TO OKAY`

This message indicates that file system *filename* was *not* marked as stable. Use the `fsck -m` command to determine if the file system needs checking.

CHAPTER 29

Troubleshooting Software Package Problems (Tasks)

This chapter describes problems you may encounter when installing or removing software packages. There are two sections: Specific Software Package Installation Errors, which describes package installation and administration errors you might encounter, and General Software Package Installation Problems, which describes behavioral problems that might not result in a particular error message.

This is a list of information in this chapter.

For information about managing software packages, see Chapter 14, "Managing Software (Overview)," in *System Administration Guide: Basic Administration*.

Troubleshooting Software Package Symbolic Link Problems

In previous Solaris releases, there was no way to specify a symbolic link target in the `pkgmap` file when creating a software package. This meant a package or patch-related symbolic link was always followed to the source of the symbolic link rather than to the target of the symbolic link when a package was added with the `pkgadd` command. This created problems when upgrading a package or a patch package that needed to change a symbolic link target destination to something else.

Now, the default behavior is that if a package needs to change the target of a symbolic link to something else, the target of the symbolic link and not the source of the symbolic link is inspected by the `pkgadd` command.

Unfortunately, this means that some packages may or may not conform to the new `pkgadd` behavior.

The PKG_NONABI_SYMLINKS environment variable might help you transition between the old and new pkgadd symbolic link behaviors. If this environment variable is set to true, pkgadd follows the source of the symbolic link.

Setting this variable enables a non-conforming package to revert to the old behavior if set by the administrator before adding a package with the pkgadd command.

The new pkgadd symbolic link behavior might cause an existing package to fail when added with the pkgadd command. You might see the following error message in this situation:

```
unable to create symbolic link to <path>
```

If a package doesn't install due to this problem, do the following:

1. If this is a Sun-supplied package, call the Resolution Center and report the non-conforming package name.
2. Set the PKG_NONABI_SYMLINKS environment variable and try adding the package with the pkgadd command again.

```
# PKG_NONABI_SYMLINKS=true
# export PKG_NONABI_SYMLINKS
# pkgadd pkg-name
```

Specific Software Package Installation Errors

```
WARNING: filename <not present on Read Only file system>
```

Reason Error Occurred	How to Fix the Problem
This error message indicates that not all of a package's files could be installed. This usually occurs when you are using pkgadd to install a package on a client. In this case, pkgadd attempts to install a package on a file system that is mounted from a server, but pkgadd doesn't have permission to do so.	If you see this warning message during a package installation, you must also install the package on the server. See Chapter 14, "Managing Software (Overview)," in *System Administration Guide: Basic Administration* for details.

General Software Package Installation Problems

Reason Error Occurred	How to Fix the Problem
There is a known problem with adding or removing some packages developed prior to the Solaris 2.5 release and compatible versions. Sometimes, when adding or removing these packages, the installation fails during user interaction or you are prompted for user interaction and your responses are ignored.	Set the following environment variable and try to add the package again. `NONABI_SCRIPTS=TRUE`

Index

B

C

D

E

F

G

H

I

K

L

M

N

O

P

Q

T

U

V

W