



System Administration Guide: Advanced Administration

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

Part No: 806-4074-10
May 2002

Copyright 2002 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

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

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

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, SunOS, Solstice, Solstice AdminSuite, Solstice DiskSuite, Solaris Solve, Java, JavaStation, DeskSet, OpenWindows, NFS and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. DecWriter, LaserWriter, Epson, NEC, Adobe

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

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

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

Copyright 2002 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés

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

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

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, SunOS, Solstice, Solstice AdminSuite, Solstice DiskSuite, Solaris Solve, Java, JavaStation, DeskSet, OpenWindows, NFS et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc. DecWriter, LaserWriter, Epson, NEC, Adobe

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

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



020115@3062



Contents

Preface	19
1 Managing Printing Services Topics	23
2 Managing Printing Services (Overview)	25
What's New in Printing?	25
Changes to LP Scheduler (lpsched)	25
USB Printer Support	26
Managing Printer Information With LDAP	26
Where to Find Printer Tasks	27
Printing in the Solaris Operating Environment	27
Solaris Print Manager	28
Printing Support in the Name Service Switch	28
Choosing a Method to Manage Printers	29
The LP Print Service	30
Managing Network Printers	30
Administering Printers	31
Setting Definitions for Printers	31
Administering Character Sets, Filters, Forms, and Fonts	31
Customizing the LP Print Service	32
The Solaris Print Client-Server Process	32
The Print Client Process	32
Using Print Clients	33
Printer Configuration Resources	34
Using Print Servers	36

3	Planning Printers on Your Network (Overview)	37
	Where to Find Printer Tasks	37
	Distributing Printers on the Network	38
	Assigning Print Servers and Print Clients	38
	Print Server Requirements and Recommendations	39
	Spooling Space	39
	Disk Space	40
	Memory	40
	Swap Space	40
	Hard Disk	41
	Planning for Printer Setup	41
	Setting Definitions for Printers	41
	Selecting a Printer Type	45
	Selecting a File Content Type	45
4	Setting Up Printers (Tasks)	49
	Setting Up Printing (Task Map)	49
	Setting Up Printing With Solaris Print Manager	50
	Setting Up a Printer With Solaris Print Manager	51
	Solaris Print Manager and Special Characters	52
	Starting Solaris Print Manager	52
	▼ How to Start Solaris Print Manager	52
	Setting Up a Print Server	54
	▼ How to Add a New Attached Printer With Solaris Print Manager	55
	Setting Up a Print Client	57
	▼ How to Add Printer Access With Solaris Print Manager	57
	Setting Up a .printers File	58
	▼ How to Set Up a .printers File	58
	Adding a Network Printer	59
	Printer Vendor Supplied Software for Network Printers	60
	Sun Support for Network Printers	60
	Invoking the Network Printer Support	61
	Selecting the Protocol	61
	Selecting the Printer Node Name	61
	Selecting the Destination (or Network Printer Access) Name	62
	Setting the Timeout Value	63
	Managing Network Printer Access	63

▼ How to Add a Network Printer With Printer Vendor Supplied Tools	63
▼ How to Add A Network Printer With LP Commands	64
Converting Printer Configuration Information	67
Converting Printer Configuration Information (Task Map)	67
Converting Existing Printer Configuration Information	68
▼ How to Convert Printer Information For a System Running the Solaris 2.5.1 Release	69
▼ How to Convert Printer Information For a System Running an lpd-based Print Server	69
5 Administering Printers (Tasks)	71
Managing Printers and the Print Scheduler	72
Deleting Printers and Printer Access	72
▼ How to Delete a Printer and Remote Printer Access	72
Checking Printer Status	75
▼ How to Check the Status of Printers	75
Restarting the Print Scheduler	76
▼ How to Stop the Print Scheduler	76
▼ How to Restart the Print Scheduler	77
Setting or Resetting Miscellaneous Printer Definitions	77
▼ How to Add a Printer Description	77
Setting Up a Default Printer Destination	78
▼ How to Set a System's Default Printer	78
Printing Banner Pages	79
▼ How to Make Banner Pages Optional	80
▼ How to Turn Off Banner Pages	81
Setting Up Printer Classes	81
▼ How to Define a Class of Printers	82
Setting Up Printer Fault Alerts	83
▼ How to Set Fault Alerts for a Printer	84
Setting Up Printer Fault Recovery	85
▼ How to Set Printer Fault Recovery	86
Limiting User Access to a Printer	87
▼ How to Limit User Access to a Printer	88
Managing Print Requests	89
▼ How to Check the Status of Print Requests	90
Enabling or Disabling a Printer From Printing	91
▼ How to Enable or Disable a Printer	92

Accepting or Rejecting Print Requests	93
▼ How to Accept or Reject Print Requests for a Printer	94
Canceling a Print Request	95
▼ How to Cancel a Print Request	95
▼ How to Cancel a Print Request From a Specific User	96
Moving a Print Request	97
▼ How to Move Print Requests to Another Printer	97
Changing the Priority of Print Requests	99
▼ How to Change the Priority of a Print Request	99
6 Managing Character Sets, Filters, Forms, and Fonts (Tasks)	101
Managing Character Sets	101
Selectable Character Sets	102
Hardware-Mounted Character Sets	103
Tracking Print Wheels	104
Alerts for Mounting Print Wheels or Cartridges	104
▼ How to Define a Print Wheel or Font Cartridge	104
▼ How to Unmount and Mount a Print Wheel or Font Cartridge	105
▼ How to Set an Alert to Mount a Print Wheel or Font Cartridge	106
▼ How to Set Up an Alias for a Selectable Character Set	108
Managing Print Filters	109
Creating Print Filters	110
Adding, Changing, Removing, and Restoring Print Filters	110
▼ How to Add a Print Filter	111
▼ How to Delete a Print Filter	112
▼ How to View Information About a Print Filter	112
Managing Forms	114
Adding, Changing, or Deleting Forms	114
Mounting Forms	114
Tracking Forms	115
Defining Alerts for Mounting Forms	115
Checking Forms	115
Limiting Access to Forms	116
▼ How to Add a Form	116
▼ How to Delete a Form	117
▼ How to Unmount and Mount a Form	117
▼ How to Set an Alert to Mount a Form	119

▼ How to View Information About a Form	120
▼ How to View the Current Status of a Form	120
▼ How to Limit User Access to a Form	121
▼ How to Limit Printer Access to a Form	122
Managing Fonts	123
Managing Printer-Resident Fonts	124
Downloading Host-Resident Fonts	124
Installing and Maintaining Host-Resident Fonts	125
▼ How to Install Downloaded PostScript Fonts	125
▼ How to Install Host-Resident PostScript Fonts	125
7 Customizing the LP Print Service (Tasks)	127
Adjusting Printer Port Characteristics	127
▼ How to Adjust the Printer Port Characteristics	128
Adding a <code>terminfo</code> Entry for an Unsupported Printer	129
▼ How to Add a <code>terminfo</code> Entry for an Unsupported Printer	132
Customizing the Printer Interface Program	133
The Standard Printer Interface Program	133
Customizing <code>stty</code> Modes	134
Exit Codes	134
Fault Messages	135
Using a Customized Printer Interface Program	136
▼ How to Set Up a Custom Printer Interface Program	136
Creating a New Print Filter	137
Writing a Print Filter Program	138
Creating a Print Filter Definition	140
▼ How to Create a New Print Filter	146
Creating a New Printer Form	148
▼ How to Create a New Form Definition	150
8 LP Print Service (Reference)	151
Where to Find Printer Tasks	151
The LP Print Service	152
The Structure of the LP Print Service	152
LP Print Service Commands	160
Functions of the LP Print Service	161
How LP Administers Files and Schedules Local Print Requests	161

	Scheduling Network Print Requests	162
	Filtering Print Files	163
	What the Printer Interface Program Does	163
	How the lpsched Daemon Tracks the Status of Print Requests	163
	Cleaning Out Log Files	164
	▼ How to Change Frequency of Printer Request Log Rotation	164
	How Local Printing Works	164
	How Remote Printing Works	166
9	Managing Terminals and Modems Topics	169
10	Managing Terminals and Modems (Overview)	171
	What's New in Managing Terminals and Modems?	171
	Terminals, Modems, Ports, and Services	172
	Terminals	172
	Modems	172
	Ports	172
	Services	173
	Port Monitors	173
	Tools for Managing Terminals and Modems	174
	Serial Ports Tool	174
	Service Access Facility	175
11	Setting Up Terminals and Modems (Tasks)	177
	Setting Up Terminals and Modems with Serial Ports Tool	177
	Setting Up Terminals	178
	Setting Up Modems	178
	▼ How to Set Up a Terminal	180
	▼ How to Set Up a Modem	181
	▼ How to Initialize a Port	181
	Troubleshooting Terminal and Modem Problems	182
12	Managing Serial Ports With the Service Access Facility (Tasks)	185
	Using the Service Access Facility	186
	Overall Administration: sacadm Command	187
	Service Access Controller: SAC Program	187

SAC Initialization Process	187
Port Monitor Service Administrator: <code>pmadm</code> Command	188
A Port Monitor at Work: <code>ttymon</code>	188
Port Initialization Process	190
Bidirectional Service	190
Port Monitors: TTY Monitor and Network Listener	190
TTY Port Monitor: <code>ttymon</code>	191
<code>ttymon</code> and the Console Port	191
Special <code>ttymon</code> -Specific Administrative Command: <code>ttyadm</code>	192
Network Listener Service: <code>listen</code>	192
Special <code>listen</code> -Specific Administrative Command: <code>nlsadmin</code>	192
Administering <code>ttymon</code> Port Monitors	193
▼ How to Add a <code>ttymon</code> Port Monitor	193
▼ How to View <code>ttymon</code> Port Monitor Status	193
▼ How to Stop a <code>ttymon</code> Port Monitor	194
▼ How to Start a <code>ttymon</code> Port Monitor	194
▼ How to Disable a <code>ttymon</code> Port Monitor	195
▼ How to Enable a <code>ttymon</code> Port Monitor	195
▼ How to Remove a <code>ttymon</code> Port Monitor	195
Administering <code>ttymon</code> Services	196
▼ How to Add a Service	196
▼ How to View the Status of a TTY Port Service	197
▼ How to Enable a Port Monitor Service	199
▼ How to Disable a Port Monitor Service	200
Service Access Facility Administration (Reference)	200
Files Associated With SAF	200
The <code>/etc/saf/_sactab</code> File	201
The <code>/etc/saf/pmtab/_pmtab</code> File	201
Service States	203
Port Monitor States	203
Port States	204

13 Managing System Resources Topics 205

14 Managing System Resources (Overview) 207

What's New in Managing System Resources?	207
New <code>df</code> , <code>du</code> , and <code>ls</code> Options	207

	Changing System Parameters	208
	Solaris Resource Management	208
	High-Level View of Managing System Resources (Task Map)	209
15	Displaying and Changing System Information (Tasks)	211
	Displaying System Information	211
	▼ How to Determine If a System Can Run the 64-bit Solaris Operating Environment	212
	▼ How to Determine If a System Has 64-bit Solaris Capabilities Enabled	212
	▼ How to Display System and Software Release Information	213
	▼ How to Display General System Information	214
	▼ How to Display a System's Host ID Number	214
	▼ How to Display a System's Installed Memory	215
	▼ How to Display the Date and Time	215
	Changing System Information	216
	▼ How to Set a System's Date and Time Manually	216
	▼ How to Set Up a Message-of-the-Day	216
	▼ How to Change a System's Host Name	217
16	Managing Disk Use (Tasks)	219
	Displaying Information About Files and Disk Space	219
	▼ How to Display Information About Files and Disk Space	220
	Checking the Size of Files	222
	▼ How to Display the Size of Files	222
	▼ How to Find Large Files	223
	▼ How to Find Files That Exceed a Specified Size Limit	223
	Checking the Size of Directories	224
	▼ How to Display the Size of Directories, Subdirectories, and Files	224
	▼ How to Display the User Ownership of Local UFS File Systems	226
	Finding and Removing Old or Inactive Files	227
	▼ How to List the Newest Files	227
	▼ How to Find and Remove Old or Inactive Files	228
	▼ How to Clear Out Temporary Directories	229
	▼ How to Find and Delete core Files	230
	▼ How to Delete Crash Dump Files	230

17	Managing Quotas (Tasks)	233
	What Are Quotas?	233
	Using Quotas	234
	Setting Soft Limits and Hard Limits for Quotas	234
	The Difference Between Disk Block and File Limits	235
	Setting Up Quotas	235
	Guidelines for Setting Up Quotas	236
	Setting Up Quotas (Task Map)	237
	▼ How to Configure File Systems for Quotas	237
	▼ How to Set Up Quotas for a User	238
	▼ How to Set Up Quotas for Multiple Users	239
	▼ How to Check Quota Consistency	240
	▼ How to Turn On Quotas	240
	Checking Quotas	241
	▼ How to Check for Exceeded Quotas	241
	▼ How to Check Quotas on a File System	242
	Changing and Removing Quotas	243
	▼ How to Change the Soft Time Limit Default	244
	▼ How to Change Quotas for a User	245
	▼ How to Disable Quotas for a User	246
	▼ How to Turn Off Quotas	247
18	Scheduling System Tasks (Tasks)	249
	Ways to Automatically Execute System Tasks	249
	For Scheduling Repetitive Jobs: <code>crontab</code>	250
	For Scheduling a Single Job: <code>at</code>	251
	Scheduling a Repetitive System Task (<code>cron</code>)	251
	Inside a <code>crontab</code> File	251
	How the <code>cron</code> Daemon Handles Scheduling	252
	Syntax of <code>crontab</code> File Entries	253
	Creating and Editing <code>crontab</code> Files	254
	▼ How to Create or Edit a <code>crontab</code> File	254
	▼ How to Verify a <code>crontab</code> File	255
	Displaying <code>crontab</code> Files	255
	▼ How to Display a <code>crontab</code> File	255
	Removing <code>crontab</code> Files	256
	▼ How to Remove a <code>crontab</code> File	257

Controlling Access to the crontab Command	257
▼ How to Deny crontab Access	258
▼ How to Limit crontab Access to Specified Users	259
▼ How to Verify Limited crontab Access	260
Scheduling a Single System Task (at)	260
Description of the at Command	260
▼ How to Create an at Job	261
▼ How to Display the at Queue	262
▼ How to Verify an at Job	262
▼ How to Display at Jobs	263
▼ How to Remove at Jobs	263
▼ How to Deny Access to the at Command	264
▼ How to Verify that at Access Is Denied	265
19 Managing System Accounting (Tasks)	267
What is System Accounting?	267
How System Accounting Works	268
System Accounting Components	268
Setting Up System Accounting	272
▼ How to Set Up System Accounting	273
Billing Users	274
▼ How to Bill Users	275
Maintaining Accounting Information	275
Fixing Corrupted Files and wtmpx Errors	275
▼ How to Fix a Corrupted wtmpx File	276
Fixing tacct Errors	276
▼ How to Fix tacct Errors	276
Restarting the runacct Script	277
▼ How to Restart the runacct Script	277
Stopping and Disabling System Accounting	278
▼ How to Temporarily Stop System Accounting	278
▼ How to Permanently Disable System Accounting	278
20 System Accounting (Reference)	281
The runacct Script	281
Daily Accounting Reports	283
Daily Report	284

Daily Usage Report	285
Daily Command Summary	286
Monthly Command Summary	288
Last Login Report	288
Looking at the <code>pacct</code> File With <code>acctcom</code>	289
System Accounting Files	291
Files Produced by the <code>runacct</code> Script	293

21 Managing System Performance Topics 295

22 Managing System Performance (Overview) 297

What's New in Managing System Performance?	297
DNLC Improvements	297
The <code>pargs</code> and <code>preap</code> Commands	298
Performance Tool	299
New Fair Share (FSS) and Fixed (FX) Schedulers	299
Where to Find System Performance Tasks	300
System Performance and System Resources	300
Processes and System Performance	301
About Monitoring Performance	302
Monitoring Tools	303

23 Managing System Processes (Tasks) 305

Commands for Managing System Processes (Overview)	305
The <code>ps</code> Command	306
▼ How to List Processes	307
The <code>/proc</code> File System and Commands	308
Managing Processes With <code>/proc</code> Tools	309
▼ How to Display Information About Processes	310
▼ How to Control Processes	311
Killing a Process (<code>kill</code>)	312
▼ How to Kill a Process	312
Managing Process Class Information	313
Changing the Scheduling Priority of Processes With <code>prionctl</code>	313
▼ How to Display Basic Information About Process Classes	314
▼ How to Display the Global Priority of a Process	314
▼ How to Designate a Process Priority	315

▼ How to Change Scheduling Parameters of a Timesharing Process	315
▼ How to Change the Class of a Process	316
Changing the Priority of a Timesharing Process With the <code>nice</code> Command	317
▼ How to Change the Priority of a Process	318
Troubleshooting Problems With System Processes	318
24 Monitoring System Performance (Tasks)	321
What's New in Monitoring System Performance?	322
Displaying Virtual Memory Statistics (<code>vmstat</code>)	322
▼ How to Display Virtual Memory Statistics (<code>vmstat</code>)	323
▼ How to Display System Event Information (<code>vmstat -s</code>)	324
▼ How to Display Swapping Statistics (<code>vmstat -S</code>)	325
▼ How to Display Cache Flushing Statistics (<code>vmstat -c</code>)	325
▼ How to Display Interrupts Per Device (<code>vmstat -i</code>)	326
Displaying Disk Utilization Information (<code>iostat n</code>)	326
▼ How to Display Disk Utilization Information (<code>iostat</code>)	327
▼ How to Display Extended Disk Statistics (<code>iostat -xtc</code>)	328
Displaying Disk Space Statistics (<code>df</code>)	329
▼ How to Display Disk Space Information (<code>df</code>)	329
Monitoring System Activities (<code>sar</code>)	330
▼ How to Check File Access (<code>sar -a</code>)	330
▼ How to Check Buffer Activity (<code>sar -b</code>)	331
▼ How to Check System Call Statistics (<code>sar -c</code>)	332
▼ How to Check Disk Activity (<code>sar -d</code>)	333
▼ How to Check Page-Out and Memory (<code>sar -g</code>)	335
▼ How to Check Kernel Memory Allocation (<code>sar -k</code>)	336
▼ How to Check Interprocess Communication (<code>sar -m</code>)	337
▼ How to Check Page-In Activity (<code>sar -p</code>)	338
▼ How to Check Queue Activity (<code>sar -q</code>)	340
▼ How to Check Unused Memory (<code>sar -r</code>)	341
▼ How to Check CPU Utilization (<code>sar -u</code>)	341
▼ How to Check System Table Status (<code>sar -v</code>)	342
▼ How to Check Swapping Activity (<code>sar -w</code>)	343
▼ How to Check Terminal Activity (<code>sar -y</code>)	345
▼ How to Check Overall System Performance (<code>sar -A</code>)	346
Collecting System Activity Data Automatically (<code>sar</code>)	346
Running the <code>sadc</code> Command When Booting	347

	Running the <code>sadc</code> Command Periodically With the <code>sa1</code> Script	347
	Producing Reports With the <code>sa2</code> Shell Script	347
	Setting Up Automatic Data Collection (<code>sar</code>)	347
	How to Set Up Automatic Data Collection	349
25	Troubleshooting Solaris Software Topics	351
26	Troubleshooting Software Problems (Overview)	353
	What's New in Troubleshooting Software Problems?	353
	New System Log Rotation	353
	New Fall Back Shell for root Account	355
	Where to Find Software Troubleshooting Tasks	355
	Troubleshooting a System Crash	355
	What to Do if the System Crashes	356
	Gathering Troubleshooting Data	356
	Troubleshooting a System Crash Checklist	357
	Viewing System Messages	358
	▼ How to View System Messages	359
	Customizing System Message Logging	360
	▼ How to Customize System Message Logging	361
	Enabling Remote Console Messaging	362
	Using Auxiliary Console Messaging During Run Level Transitions	362
	Using the <code>consadm</code> Command During an Interactive Login Session	363
	▼ How to Enable an Auxiliary (Remote) Console	364
	▼ How to Display a List of Auxiliary Consoles	365
	▼ How to Enable an Auxiliary (Remote) Console Across System Reboots	365
	▼ How to Disable an Auxiliary (Remote) Console	365
27	Managing Core Files (Tasks)	367
	Managing Core Files (Task Map)	367
	Managing Core Files Overview	368
	Configurable Core File Paths	368
	Expanded Core File Names	369
	Setting the Core File Name Pattern	369
	Enabling <code>setuid</code> Programs to Produce Core Files	370
	How to Display the Current Core Dump Configuration	370
	How to Set a Core File Name Pattern	370

	▼ How to Enable a Per-Process Core File Path	371
	▼ How to Enable a Global Core File Path	371
	Troubleshooting Core File Problems	372
	Examining Core Files	372
	Example—Examining Core Files With proc Tools	372
28	Managing System Crash Information (Tasks)	373
	Managing System Crash Information (Task Map)	373
	System Crashes (Overview)	374
	System Crash Dump Files	374
	Saving Crash Dumps	375
	The dumpadm Command	375
	Dump Devices and Volume Managers	376
	Managing System Crash Dump Information	377
	▼ How to Display the Current Crash Dump Configuration	377
	▼ How to Modify a Crash Dump Configuration	377
	▼ How to Examine a Crash Dump	379
	▼ How to Recover From a Full Crash Dump Directory (Optional)	380
	▼ How to Disable or Enable Saving Crash Dumps	380
29	Troubleshooting Miscellaneous Software Problems (Tasks)	381
	What to Do If Rebooting Fails	381
	SPARC: Troubleshooting 64-bit Solaris Boot Problems	382
	What to Do if You Forgot Root Password	383
	SPARC: Example—What to Do if You Forgot Root Password	383
	IA: Example—What to Do if You Forgot Root Password	384
	What to Do if a System Hangs	385
	What to Do if a File System Fills Up	386
	File System Fills Up Because a Large File or Directory Was Created	386
	A TMPFS File System is Full Because the System Ran Out of Memory	387
	What to Do if File ACLs Are Lost After Copy or Restore	387
	Troubleshooting Backup Problems	387
	The root (/) File System Fills Up After You Back Up a File System	387
	Make Sure the Backup and Restore Commands Match	388
	Check to Make Sure You Have the Right Current Directory	388

30	Troubleshooting File Access Problems (Tasks)	391
	Solving Problems With Search Paths (Command not found)	391
	▼ How to Diagnose and Correct Search Path Problems	392
	Solving File Access Problems	394
	Changing File and Group Ownerships	394
	Recognizing Problems With Network Access	394
31	Troubleshooting Printing Problems (Tasks)	395
	Tips on Troubleshooting Printing Problems	395
	Troubleshooting No Output (Nothing Prints)	396
	Troubleshooting Incorrect Output	397
	Troubleshooting Hung LP Commands	399
	Troubleshooting Idle (Hung) Printers	399
	Troubleshooting Conflicting Status Messages	400
	Troubleshooting Printing Problems	401
	▼ How to Troubleshoot No Printer Output	401
	▼ How to Check the Printer Hardware	401
	▼ How To Check the Printer Network Connections	402
	▼ How to Check the Basic Functions of the LP Print Service	403
	▼ How to Check Printing From a Solaris Print Client to a Solaris Print Server	406
	▼ How to Troubleshoot Incorrect Output	407
	▼ How to Unhang the LP Print Service	411
	▼ How to Troubleshoot an Idle (Hung) Printer	411
	▼ How to Check That the Printer is Ready to Print	411
	▼ How to Check for Print Filtering	412
	▼ How to Resume Printing After a Printer Fault	412
	▼ How to Send Print Requests to a Remote Printer When They Back Up in the Local Queue	412
	▼ How to Free Print Requests From a Print Client That Back Up in the Print Server Queue	413
	▼ How to Resolve Conflicting Printer Status Messages	413
32	Resolving UFS File System Inconsistencies (Tasks)	415
	<code>fsck</code> Error Messages	415
	General <code>fsck</code> Error Messages	417
	Initialization Phase <code>fsck</code> Messages	418
	Phase 1: Check Blocks and Sizes Messages	421

Phase 1B: Rescan for More DUPS Messages	424
Phase 2: Check Path Names Messages	424
Phase 3: Check Connectivity Messages	430
Phase 4: Check Reference Counts Messages	432
Phase 5: Check Cylinder Groups Messages	435
Cleanup Phase Messages	436
33 Troubleshooting Software Package Problems (Tasks)	437
Troubleshooting Software Package Symbolic Link Problems	437
Specific Software Package Installation Errors	438
General Software Package Installation Problems	439
Index	441

Preface

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

This book assumes that you have already installed the SunOS™ 5.9 operating system, and you have set up any networking software that you plan to use. The SunOS 5.9 operating system is part of the Solaris 9 product family, which also includes many features, including the Solaris Common Desktop Environment (CDE). The SunOS 5.9 operating system is compliant with AT&T's System V, Release 4 operating system.

For the Solaris 9 release, new features interesting to system administrators are covered in sections called *What's New in ... ?* in the appropriate chapters.

Note – The Solaris operating environment runs on two types of hardware, or platforms—SPARC and IA. The Solaris operating environment runs on both 64-bit and 32-bit address spaces. The information in this document pertains to both platforms and address spaces unless called out in a special chapter, section, note, bullet, figure, table, example, or code example.

Who Should Use This Book

This book is intended for anyone responsible for administering one or more systems running the Solaris 9 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
<i>System Administration Guide: Basic Administration</i>	User accounts and groups, server and client support, shutting down and booting a system, removable media, managing software (packages and patches), disks and devices, file systems, and backing up and restoring data
<i>System Administration Guide: Advanced Administration</i>	Printing services, terminals and modems, system resources (disk quotas, accounting, and crontabs), system processes, and troubleshooting Solaris software problems
<i>System Administration Guide: IP Services</i>	TCP/IP networks, IPv4 and IPv6, DHCP, IP Security, Mobile IP, and IP Network Multipathing
<i>System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)</i>	DNS, NIS, and LDAP naming and directory services
<i>System Administration Guide: Naming and Directory Services (FNS and NIS+)</i>	FNS and NIS+ naming and directory services
<i>System Administration Guide: Resource Management and Network Services</i>	Resource management, remote file systems, mail, SLP, and PPP
<i>System Administration Guide: Security Services</i>	Auditing, PAM, RBAC, and SEAM

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

Accessing Sun Documentation Online

The docs.sun.comSM 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 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 <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with on-screen computer output	<code>machine_name%</code> su Password:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	To delete a file, type rm <i>filename</i> .
<i>AaBbCc123</i>	Book titles, new words or terms, or words to be emphasized.	Read Chapter 6 in <i>User's Guide</i> . These are called <i>class</i> options. Do <i>not</i> 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	<code>machine_name%</code>
C shell superuser prompt	<code>machine_name#</code>
Bourne shell and Korn shell prompt	<code>\$</code>
Bourne shell and Korn shell superuser prompt	<code>#</code>

General Conventions

Be aware of the following conventions 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 5.9 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 options from the SunOS 5.9 commands.

Managing Printing Services Topics

This topic map lists the chapters that provide information for managing printing services.

Chapter 2	Provides overview information for managing printing services on a network, which includes information on print servers, print clients, and the LP print service.
Chapter 3	Provides overview information for planning printing services on a network, which includes information on allocating system resources and defining printers on a network.
Chapter 4	Provides step-by-step instructions for setting up a printer on a system and making it available to other systems on the network.
Chapter 5	Provides step-by-step instructions for administering printers, such as deleting printers, setting print policies, and managing print requests.
Chapter 6	Provides step-by-step instructions for setting up and maintaining character sets, print filters, forms, and fonts.
Chapter 7	Provides step-by-step instructions for customizing the LP print service, such as adjusting printer port characteristics or adding a <code>terminfo</code> entry for a unsupported printer.
Chapter 8	Provides background information on the LP print service.

Managing Printing Services (Overview)

This chapter provides information about managing printers in the Solaris environment.

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

- “What’s New in Printing?” on page 25
- “Printing in the Solaris Operating Environment” on page 27
- “The LP Print Service” on page 30
- “The Solaris Print Client-Server Process” on page 32

What’s New in Printing?

This section describes new printing features in the Solaris 9 release.

Changes to LP Scheduler (`lpsched`)

In previous Solaris releases, the `lpsched` process was automatically started at system boot time whether or not there were local printers configured on the system.

In this Solaris release, the `lpadmin` command automatically starts the `lpsched` process when local printers are added to the system and stops it when the last local printer is removed.

This change does not alter either the use of or the administration of the print subsystem from the perspective of print clients. The `lpsched` process is not running on systems that do not have local printers configured. This does not effect the ability to print nor the ability to configure new printers.

USB Printer Support

You can use Solaris Print Manager to set up a USB printer that is attached to a SPARC or IA system with USB ports.

The new logical device names for USB printers are:

```
/dev/printers/[0..N]*
```

When you add a USB printer to a printer server, select one of these devices for a USB printer under Printer Port on the Add New Attached Printer screen. For more information on using Solaris Print Manager to set up printers, see Chapter 4.

Although the new Solaris USB printer driver supports all USB printer-class compliant printers, a list of recommended PostScript™ printers is in the `usbprn(7D)` man page.

The `usbprn` driver is compliant with non-PostScript printers that utilize third-party PostScript conversion packages like GhostScript. You can obtain conversion packages from the Solaris Software Companion CD, available at <http://www.sun.com/software/solaris/binaries/package.html>.

For information and cautions about hot-plugging USB printers, see the Notes and Diagnostics sections of the `usbprn(7D)` man page.

Managing Printer Information With LDAP

You can manage printer information in the following name services in this Solaris release:

- LDAP
- NIS
- NIS+
- NIS+ with Federated Naming Service (xfn)
- files

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

- If the LDAP server is the Netscape™ Directory Server (NSDS), the default distinguished name is `cn=Directory Manager`. If the LDAP server is Sun™ Directory Server, the distinguished name might be something like `cn=admin, dc=xyz, dc=com`. Solaris Print Manager uses the `ldapclient` command to determine the default LDAP server name. If there is more than one server specified, the first one is automatically selected. For more information, see `ldapclient(1M)`.
- Solaris Print Manager always displays printer entries from the current LDAP server. If this is not the domain's LDAP master server, the list of printers displayed may not be the current list of printers. This is because the LDAP replica server might not have been updated by the master server, and so, is out of sync with the

master. Replica servers can have various update replication agreements. For example, when there is a change on the master, the replica servers can be updated immediately or updated once a day.

- If the selected LDAP server is a LDAP replica server, any updates are referred to the master server and done there. This again means the printer list could be out of sync with the master. For example, a deleted printer might still appear in the displayed printer list until the replica is updated from the master.
- Users can use the LDAP command line utilities, `ldapadd` and `ldapmodify`, to update printer entries in the directory, but this 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 it is not unique, the result of modifications done by Solaris Print Manager, or the `lpset` command, might not be predictable.

Where to Find Printer Tasks

Printer Task	For More Information
Set up printers with Solaris Print Manager and the <code>lp</code> commands	Chapter 4
Administer printers with the <code>lp</code> commands after they are setup	Chapter 5
Manage character sets, filters, forms, and fonts	Chapter 6
Customize the LP print service	Chapter 7

Printing in the Solaris Operating Environment

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 components:

- Solaris Print Manager, a graphic user interface, 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 used to set up and manage printers on a local system or in a name service. They also provide functionality above and 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 commands to completely manage printing in the Solaris environment. For more information, see Chapter 5.

The Solaris print software limitations 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-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: NIS, NIS+, NIS+ with Federated Naming Service (xfn), and files. You must be logged in as superuser to use this tool.

Using Solaris Printer Manager is the preferred method for managing printer configuration information. It is preferred over `Admintool: Printers` because it centralizes printer information when used in conjunction with a name service. Using a name service for storing printer configuration information is desirable because it makes printer information available to all systems on the network, making printing administration easier.

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

The Solaris Print Manager package is `SUNWppm`.

For step-by-step instructions on setting up new printers with Solaris Print Manager, see Chapter 4.

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. The `xfn` keyword represents the FNS printer contexts.

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 xfn

For example, if your name service is NIS, printer configuration information on print clients is looked 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 `nsswitch.conf(4)` and *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*.

Choosing a Method to Manage Printers

Adding printer information to a name service makes access to printers available to all systems on the network and generally 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, NIS+, or NIS+ (xfn) database makes the printer available to all systems on the network.
Don't use a name service	<p>Adding the printer adds the printer information to the printer server's configuration files only. Print clients will not know about the printer automatically.</p> <p>You will have to add the printer information to every print client that needs to use the printer.</p>

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

TABLE 2-1 Solaris Printing Component Features

Component	Availability	Graphical User Interface?	Configures Network Printers?	Manages Print Clients and Servers?	Uses LDAP, NIS, NIS+, or NIS+ (xfn)?
Solaris Print Manager	Solaris 9 and compatible versions and Solaris Easy Access Server 3.0	Yes	Yes	Yes	Yes
Admintool	Solaris 9 and compatible versions	Yes	No	Yes	No
LP commands	Solaris 9 and compatible versions	No	Yes	Yes	Yes

After using the table above to determine which printing tool is best for your network environment, see Chapter 4 for printer setup information.

Most printing configuration tasks can be accomplished with Solaris Print Manager. However, if you have special needs, such as writing interface scripts or adding your own filters, you can use the LP print service commands, which underlie Solaris Print Manager. Performing printing administration tasks with LP commands are described in Chapter 5.

The LP Print Service

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

For background information about the LP print service, see Chapter 8.

Other LP print service topics covered and their chapter references are described in the following sections.

Managing Network Printers

A *network printer* is a hardware device that is connected directly to the network. It 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, but 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 4.

Administering Printers

After you set up print servers and print clients, there are a number of administration tasks you might need to perform frequently:

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

For step-by-step instructions on how to perform the printer administration tasks, see Chapter 5.

Setting Definitions for Printers

Establishing 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 3.

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 6.

Customizing the LP Print Service

Although the LP print service is designed to be flexible enough to handle most printers and printing needs, it 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 can 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 to customize the LP print service, see Chapter 7.

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 it is printed.

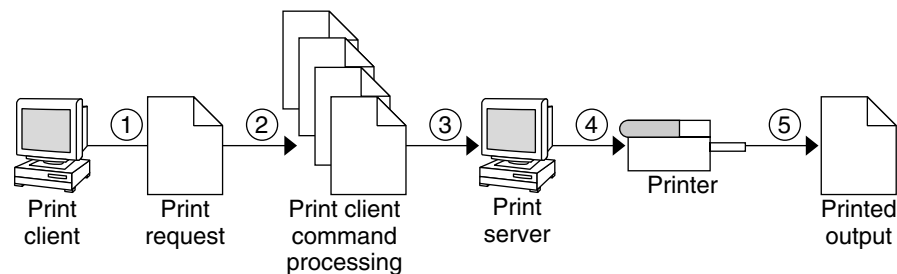


FIGURE 2-1 Overview of the Print Client Process

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 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, and 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.

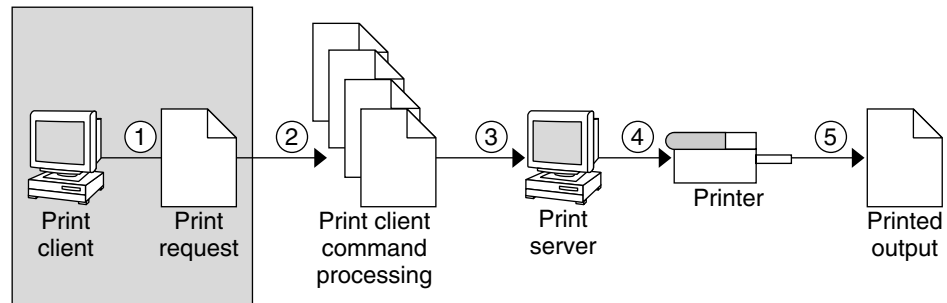


FIGURE 2-2 The User Submits a Print Request from a Print Client

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 for the NIS name service
- The LDAP or NIS name service database

The print client sends its requests to the print server's queue. The 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 printer configuration information for all printers on the network. The name service (LDAP, NIS, NIS+, or NIS+ (xfn)) simplifies printer configuration maintenance: When you add a printer in the name service, all print clients on the network can access it.

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.

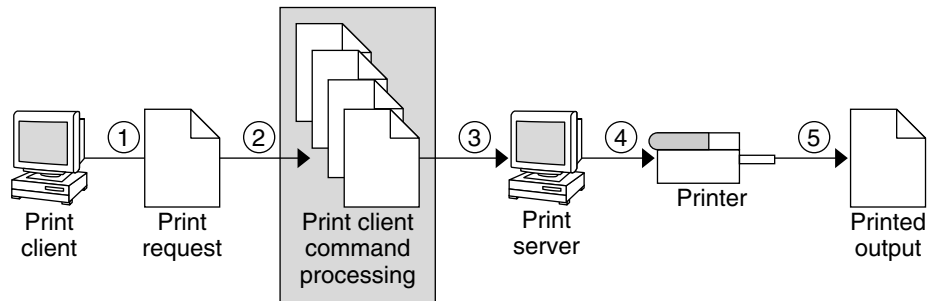



FIGURE 2-3 The Print Client Checks Resources to Locate Printers

How the Print Software Locates Printers

As shown in the following figure, the print software use more options to locate printers and printer configuration information.



- | | |
|--|---|
| A. Atomic, POSIX, or Context-Based Printer Name or Class | E. Local <code>/etc/printers.conf</code> File |
| B. User's <code>PRINTER</code> or <code>LPDEST</code> Environment Variable for Default Printer | F. LDAP directory information tree for printers in the domain's 'ou=printers' container |
| C. <code>_default</code> Variable in Printers Database in <code>/etc/nsswitch.conf</code> | G. NIS <code>printers.conf.byname</code> Map |
| D. User's <code>\$HOME/.printers</code> File | H. NIS+ <code>printers.org_dir</code> Table |
| | I. NIS+ (xfn) FNS Printing Contexts |

FIGURE 2-4 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 the following example.

```
% lpr -P galaxy:neptune filename
```

- Context-based style 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:

- It 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.

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 2-5 The Print Server Sends a Print Request to 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 5.9 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. It is widely used and 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.

Where to Go From Here

Go to Chapter 4 for step-by-step instructions on setting up new printers with Solaris Print Manager. If you need printer planning information, see Chapter 3.

Planning Printers on Your Network (Overview)

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” on page 38
- “Assigning Print Servers and Print Clients” on page 38
- “Print Server Requirements and Recommendations” on page 39

Where to Find Printer Tasks

Printer Task	For More Information
Set up printers with Solaris Print Manager and the lp commands	Chapter 4
Administer printers with the lp commands after they are setup	Chapter 5
Manage character sets, filters, forms, and fonts	Chapter 6
Customize the LP print service	Chapter 7

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, it usually works best to 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 users' systems for convenience or for improved response. A printer that is connected to a user's 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 3-1 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, and 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 environment. Besides physically connecting a printer to a system, you must define the printer characteristics to the LP print service and make the system a print server. 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 version of the UNIX operating system. Print clients running the Solaris 9 release and compatible versions can print to print servers 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 166.




FIGURE 3-2 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 they have been

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, which 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. Look at 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. If users in a small group typically print only short email messages, (simple ASCII files without sophisticated formatting requirements), 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 or bit-mapped or raster images, they will likely fill up 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 resides in a small partition, and if a large amount of disk space is available elsewhere, you can use that space as spooling space by mounting it on the `/var` directory on the print server. For information about mounting file systems and editing the `vfstab` file, see “Mounting and Unmounting File Systems (Tasks)” in *System Administration Guide: Basic Administration*.

Memory

The Solaris environment 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 filtering 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 “Configuring Additional Swap Space (Tasks)” in *System Administration Guide: Basic Administration*.

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 in the Solaris environment that includes:

- Setting definitions for printers such a printer name, printer description, printer port
- Selecting a printer type and file content type
- Setting up fault notification and a default printer destination
- Determining whether you want to print banner pages or limit user access to a printer
- Setting up printer classes and fault recovery

Setting Definitions for Printers

Establishing 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 print definitions, while Solaris Print Manager lets you set only some of them when you install or modify a printer. The following table lists the print definitions and shows whether you can assign the definition with Solaris Print Manager.

TABLE 3-1 Print Definitions Set With Solaris Print Manager

Print Definition	Can You Set It With Solaris Print Manager?
Printer name	Yes
Printer description	Yes
Printer port	Yes

TABLE 3-1 Print Definitions Set With Solaris Print Manager (Continued)

Print Definition	Can You Set It With Solaris Print Manager?
Printer type	Yes
File contents	Yes, but with less functionality than the <code>lpadmin</code> command
Fault notification	Yes, but with less functionality than the <code>lpadmin</code> command
Default printer destination	Yes
Printing banner pages	Yes, but with less functionality than the <code>lpadmin</code> command
Limiting user access to a printer	Yes, but with less functionality than the <code>lpadmin</code> 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 be:

- Unique among all printers within the bounds of an administrative domain
- A maximum of 14 alphanumeric characters, which can include dashes and underscores
- Easy to remember and can 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. If, however, 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 there are printing problems.

Users can look at 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, by using Solaris Print Manager or the `lpadmin -p printer-name -v device-name` command.

Most systems have two serial ports and a parallel port or USB ports. Unless you add ports, you cannot directly connect more than two serial printers and a parallel 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	<code>/dev/term/a</code>
Serial	<code>/dev/term/b</code>
Parallel	<code>/dev/printers/0 -> /dev/ecpp0</code>
USB	<code>/dev/printers/[1-9]</code>
Specify any port name that the print server recognizes	Other

The LP print service initializes the printer port using the settings from the standard printer interface program. For more information about printer interface programs, see “Managing Print Filters” on page 109. 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 127 for information about customizing the port settings.

IA only – If you use multiple ports on an IA based system, only the first port is enabled by default. The second 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 `lp` (parallel) port. The pathnames for the IA port configuration files are:

```
/platform/i86pc/kernel/drv/asy.conf
```

```
/platform/i86pc/kernel/drv/lp.conf
```

For information about configuring serial and parallel ports on IA based systems, see the *Solaris 9 (Intel Platform Edition) Hardware Compatibility List*.

Printer Type

The printer type is a generic name for a type of printer. It 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 Apple LaserWriter® I and Apple LaserWriterII printers.

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

Solaris Print Manager lets you select a printer type from a menu or choose `Other` and specify any printer type in the `terminfo` database. This provides you as much capability as the `lpadmin` command.

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  epon2500      esprit
envision230    ep40        epon2500-80   ethernet
ep2500+basic   ep4000      epon2500-hi   ex3000
ep2500+color   ep4080      epon2500-hi80 exidy
$
```

The entries for Epson printers are included in the preceding example.

If you have a 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           nuclterm
ncr7900i        nec          netty-Tabs    newhp
ncr7900i-na     net         netty-vi      newhpkeyboard
$

```

The entry in this directory for NEC is included in the preceding example.

Selecting a Printer Type

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—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 them 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® Document Structuring Conventions in Appendix C of the *PostScript Language Reference Manual* (written by Adobe Systems Incorporated, and published by Addison-Wesley, 1990).

If a printer can emulate more than one kind of printer, you can assign it several types by using the `lpadmin -T` command. If you specify more than one printer type, the LP print service uses the 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 manufacturer or product names.

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 model of printer. If possible, find an entry in the `terminfo` database that works for your printer. It will be much easier than trying to create an entry. If you have to create your own entry, “Adding a `terminfo` Entry for an Unsupported Printer” on page 129 contains some useful tips.

Selecting a File Content Type

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.)

You can specify the file content type for a printer by using the `lpadmin -I` command or Solaris Print Manager. With Solaris Print Manager, you can select a file content type from a menu. Not all available file content types are listed on the menu. You must use the `lpadmin` command to specify file content types that are not included on the Solaris Print Manager menu.

Most printers can print the following types of files directly:

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

When submitting a file for printing, the user can indicate the content type of the file (`lp -T content-type`). If no file content type is supplied when the request is submitted, the LP server looks at 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.

Solaris Print Manager provides a list of file content types from which you can choose when you install or modify a local printer. The choices are translated to the names that the LP print service uses. The following table describes the file content types you can choose with Solaris Print Manager.

TABLE 3-2 Choosing File Content Type With Solaris Print Manager

File Contents Choice	LP Print Service Name	Description
PostScript	<code>postscript</code>	PostScript files do not require filtering.
ASCII	<code>simple</code>	ASCII files do not require filtering.
Both PostScript and ASCII	<code>simple, postscript</code>	PostScript files and ASCII files do not require filtering.
None	<code>" "</code>	All files require filtering, except those matching the printer's type.
Any	<code>any</code>	No filtering 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 (which means filtering is not needed for PostScript files) is the default choice in Solaris Print Manager and is probably correct most of the time.

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` 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 3-3 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 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 132.

Setting Up Printers (Tasks)

This chapter describes how to set up a printer and make it accessible to systems on the network with Solaris Print Manager.

For information on the procedures associated with setting up printers with Solaris Print Manager, see “Setting Up Printing (Task Map)” on page 49.

For overview information about printers, see Chapter 2.

Setting Up Printing (Task Map)

Task	Description	For Instructions
1. Add New Attached Printer	Using Solaris Print Manager – After you physically attach the printer to a system, make the printer available for printing.	“How to Add a New Attached Printer With Solaris Print Manager” on page 55
2. Add Access to a Printer	Using Solaris Print Manager – Add printer access on the print client.	“How to Add Printer Access With Solaris Print Manager” on page 57
3. (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 58

Task	Description	For Instructions
4. Add a New Network Printer	<p>Using Printer Vendor Supplied Tools – After you physically connect the printer to the network, use vendor-supplied software to configure the network printer.</p> <p>Using LP Commands – After you physically connect the printer to the network, configure the network printer.</p>	<p>“How to Add a Network Printer With Printer Vendor Supplied Tools” on page 63</p> <p>“How to Add A Network Printer With LP Commands” on page 64</p>
5. (Optional) Turn Off Banner Pages	You can turn off banner pages so they are never printed.	“How to Turn Off Banner Pages” on page 81
6. (Optional) Set Up Fault Alerts	You can set up more specific fault alerts for the printer than Solaris Print Manager provides.	“How to Set Fault Alerts for a Printer” on page 84
7. (Optional) Set Up Fault Recovery	You can set up how a printer should recover after it faults.	“How to Set Printer Fault Recovery” on page 86
8. (Optional) Limit Access to the Printer	Solaris Print Manager enables you to set up an allow list, but if you want to limit a few users’ access to the printer, you might want to set up a deny list.	“How to Limit User Access to a Printer” on page 88

Setting Up Printing With 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+, NIS+ with Federated Naming Service (xfn), and files. You must be logged in as superuser to use this tool.

You can use Solaris Print Manager to set up print servers (New Attached Printer) and print clients (Add Access to Printer). A local or attached printer is one which is physically cabled to the print server. A network printer is physically attached to the network. 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.

Setting Up a Printer With Solaris Print Manager

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

Printer Attribute	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 printer server	venus	The local system	Required to install an attached or network printer and to add access to a printer
Description	User defined string	laser printer near breakroom	N/A	Optional
Printer Port	Device 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 or network printer
File Contents	Content to be printed	any	PostScript	Required to install an attached or network printer
Destination	Destination name for network printers	See "Selecting the Destination (or Network Printer Access) Name" on page 62 for examples	N/A	Required to install a network printer
Protocol	Protocol used to communicate with printer	TCP	BSD	Required to install a network printer
Fault Notification	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	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 meta characters, 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 use Solaris Print Manager to set up your printers, start Solaris Print Manager either by selecting Printer Administrator from the CDE Workspace menu or by starting it from the command line. See the following section for details.

▼ How to Start Solaris Print Manager

1. Verify that the following prerequisites are met. To use Solaris Print Manager, you must:

- Have a bit-mapped display monitor. Solaris Print Manager can be used only on a system with a console that is a bit-mapped screen, such as a standard display monitor that comes with a Sun workstation.
- 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 on the printer 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, NIS+, or NIS+ (xfn) 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 is usually displayed and selected automatically.

For more information about using LDAP to manage printer information, see “Managing Printer Information With LDAP” on page 26.

- If your name service is NIS, you must have the root password for the NIS master.
- If your name service is NIS+, you might need to do the following:
 1. Log in to the NIS+ master as superuser.
 2. Identify the group that owns the printers table:

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

3. 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 host_name
```

4. Log in to the system that runs Solaris Print Manager as superuser. Depending on your NIS+ configuration, you might also need to run the /usr/bin/keylogin command. For more information, see keylogin(1).

- If your name service is NIS+ (xfn), you might need to do the following:
 1. Log in to the NIS+ master as superuser.
 2. Identify the group that owns the federated naming table:

```
# niscat -o fns.ctx_dir.domain_name.com
.
.
.
Group          : "admin.domain_name.com"
```

3. If necessary, add the system that runs Solaris Print Manager to the NIS+ admin group authorized to update the fns.ctx_dir.<domain> file.

```
# nisgrpadm -a admin.domain_name.com host_name
```

4. Log in to the system that runs Solaris Print Manager as superuser. Depending on your NIS+ configuration, you might also need to run the /usr/bin/keylogin command. See keylogin(1) for more information.

- Have the SUNWppm package installed.

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

2. Start Solaris Print Manager by 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, and click the Printer Administrator icon in the Application Manager’s System_Admin window
- From the command line:

```
# /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, like this:

```
# 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:

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

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

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

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

```
$ pkginfo | grep SUNWppm
```

3. **Select the name service used in your network from the Select Naming Service window. Choices are: NIS+ (xfn), 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 install an attached printer and/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 sections describe how to use Solaris Print Manager to add an attached printer or a network printer on a printer server. The example that follows each Solaris Print Manager procedure describes how to add a printer with LP commands.

▼ How to Add a New Attached Printer With Solaris Print Manager

1. **Select the system to be the print server.**
2. **Connect the printer to the print server and 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 52.
4. **Select New Attached Printer from the Printer menu.**
The New Attached Printer window is displayed.
5. **Fill in the window.**
If you need information to complete a field, click the Help button.
6. **Click OK.**
7. **Verify that the printer has been installed by checking 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.

Example—Adding a New Attached Printer With LP Commands

This example shows how to make a local PostScript printer available for printing on a print server. 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 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
# cd /etc/lp/fd
# for filter in *.fd;do
  > name=`basename $filter .fd`
  > lpfilter -f $name -F $filter
  > done 5
# accept luna
  destination "luna" now accepting requests
# enable luna 6
printer "luna" now enabled
# lpadmin -p luna -D "Room 1954 ps" 7
# lpstat -p luna 8
printer luna is idle. enabled since Jul 12 11:17 2001. available.
```

1. Gives lp ownership and sole access to a port device.
2. Defines the printer name and the port device 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 print filters to the print server.
6. Accepts print requests for the printer and enables the printer.
7. Adds a description for the printer.
8. Verifies that the printer is ready.

Where to Go From Here

Use the following table to determine which tasks to complete next.

Task	For More Information
Add access to the newly installed printer on the print clients because you did not add the printer information to the name service database.	"How to Add Printer Access With Solaris Print Manager" on page 57
Set up a .printers file.	"How to Set Up a .printers File" on page 58

Setting Up a Print Client

A print client is a system that is not the server for the printer, yet 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 depending on whether you add the printer information to the name service database.

The following sections describe how to use the new Solaris Print Manager to add access to a printer on a print client. The example that follows this procedure describes how to add printer access with LP commands.

▼ How to Add Printer Access With Solaris Print Manager

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 52.

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

The Add Access to Printer window is displayed.

3. **Fill in the window.**

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—Adding Printer Access With LP 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.

Setting Up a `.printers` File

There is no need to set up a `.printers` file in your users' home directories if they 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. You can also set up a special `_all` alias to define a list of printers affected when you cancel a print request or check the status of printers.

Keep in mind that the use of the `.printers` file by the LP print service is controlled by the name service switch (`/etc/nsswitch.conf`). The default configuration is that the print service checks a user's home directory to locate printer configuration information before it checks the other name services. This means 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 `printers(4)`. For more information about the name service switch, see `nsswitch.conf(4)`.

▼ How to Set Up a `.printers` File

1. **Log in to the user's system as superuser.**
2. **Start the file editor 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 your default printer, using an entry similar to the one shown in the following example.**

```
_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, using an entry similar to the one shown in the next example.

```
_all printer1,printer2,printer3
```

5. Save the file as `.printers`.

Adding a Network Printer

A *network printer* is a hardware device that is directly connected to the network. This means it can be accessed from a print server without actually connecting it to the print server with a cable. It has its own system name and IP address. Even though a network printer is not connected to a print server, it is necessary to set up a print server for it. The print server provides queuing capabilities, filtering, and printing administration for the network printer.

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

The vendor might supply an SVR4 printer interface script to replace the standard printer interface script. If so, their SVR4 interface script will call the vendor-supplied printing program to send the job to the printer. If not, you will need to modify the standard interface script to call the vendor-supplied printing program. You can do this by editing the per-printer copy of the standard interface script to call the vendor-supplied printing program.

The terms used in network printer configuration are:

- **Print server** – The machine that spools and schedules the jobs for a printer. This is the machine on which the printer is configured.
- **Printer-host device** – The printer-host device is 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 it creates a *network printer*.
- **Printer node** – This is either the physical printer or the printer-host device. It is the physical printer when the network support resides in the physical printer. It is the printer-host device when an external box is used to provide the network interface. The printer node name is the machine name given with the IP address. This name is selected by the system administrator and has no default or vendor requirement. The printer node name, as with all nodes, must be unique.

- **Printer name** – The name entered on the command line when using any of the printer commands. It is selected by the system administrator at the time of printer configuration. Any one physical printer can have several printer or queue names; each 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. It 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. It is printer specific. In the case where 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_name:port_designation

- **Protocol** – The over-the-wire protocol used to communicate with the printer. The printer vendor documentation supplies the information regarding the protocol to select. The network printer support supplies both BSD Printer Protocol and raw TCP. Due to implementation variations, you might want to try both.
- **Timeout, or retry interval Timeout** – Is a seed number representing 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, and increases with 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 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 under an LP print system.

Sun Support for Network Printers

If the network printer vendor does not provide software support, the 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 4. The following is a discussion of printer management using the Sun supplied software.

Invoking the Network Printer Support

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

```
lpadmin -p printer_name -m netstandard
```

Selecting the Protocol

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

The command to select the protocol is:

```
lpadmin -p printer_name -o protocol=bsd
```

or

```
lpadmin -p printer_name -o protocol=tcp
```

If the protocol 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; others the reverse. For this information, see the printer vendor documentation. The default is to send the control file first.

The command to select the ordering is:

```
lpadmin -p printer_name -o bsdctrl=first
```

or

```
lpadmin -p printer_name -o bsdctrl=last
```

Selecting the Printer Node Name

The system administrator selects 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 (or Network Printer Access) Name

The print subsystem requires access information for the printer. This is the name that the subsystem uses when making the network connection to the printer. This name is supplied by the system administrator to the print subsystem by using the `lpadmin` command. It 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 will find information about port designation in the documentation that is provided with the printer by the printer vendor. The format of printer access name is:

```
printer_node-name [:port_designation]
```

Example 1—Destination (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:

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

Example 2—Destination (or Network Printer Access Name) With Port Designation (Name)

When using 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:

```
lpadmin -p printer_name -o dest=cardboard:xxx_parallel_1
```

Example 3—Destination (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:

```
lpadmin -p printer_name -o dest=newspaper
```

Setting the Timeout Value

The timeout option is provided to allow for individual selection of 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. It is a seed value which 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 alerts the user that intervention might be required. This could be anything from the printer being turned off, to out of paper. Should these messages be produced too often, for example when the printer is warming up, increasing the timeout value will eliminate spurious messages.

The system administrator can experiment to find the optimal timeout value. The command to set the timeout is:

```
lpadmin -p printer_name -o timeout=n
```

Managing Network Printer Access

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

The default device for the network printer is `/dev/null`. This is sufficient when there is only one queue for the printer. Should more queues be required, set the device to a file. This 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 With Printer Vendor Supplied Tools

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 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 will be 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 clients. For more information, see “Setting Up a Print Client” on page 57.

4. Optional tasks to complete.

There are several optional tasks you might want to complete when setting up a network printer. For pointers to the remaining tasks, see “Setting Up Printing (Task Map)” on page 49.

▼ How to Add A Network Printer With LP Commands

This procedure describes the steps necessary to setup a network printer using the network printer support software. The use of this software is intended for those printers that do not come with vendor supplied software.

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 is equivalent to adding any node to the network.

2. Collect the information required to configure a network printer.

- Printer name
- Printer server
- Network printer access name
- Protocol
- Timeout

For more information, see the terms described in “Adding a Network Printer” on page 59.

3. Define the printer name, the device, the printer type and content type.

a. Define the printer name and the port device the printer will use.

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

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

b. Identify the interface script the printer will use.

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

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

c. Set the printer destination, protocol, and timeout values.

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

<code>-p printer-name</code>	Specifies the network printer name.
<code>-o dest=access-name:port</code>	Sets the printer destination to the network printer access name and a designated printer vendor port, if it is defined in the printer vendor documentation.
<code>-o protocol=protocol</code>	Sets the over-the-wire protocol used to communicate with the printer. Both BSD and raw TCP are supported.
<code>-o timeout=value</code>	Sets a retry timeout value that represents a number of seconds to wait between attempting connections to the printer.

d. Specify the file content types of the printer and the printer type.

```
# lpadmin -p printer-name -I content-type -T printer-type
```

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

4. 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 printer requests and to print the requests.

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

6. Verify that the printer is correctly configured.

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

7. 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 57.

8. Optional tasks to complete.

There are several optional tasks you might want to complete when setting up a printer. For pointers to the remaining tasks, see “Setting Up Printing (Task Map)” on page 49.

Example—Adding A Network Printer With LP Commands

The commands in this example must be executed on the print server. The following information is used as an example. The information you provide will vary.

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

```
# 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
# cd /etc/lp/fd
# for filter in *.fd;do
    > name=`basename $filter .fd`
    > lpfilter -f $name -F $filter
    > done 5
# accept luna1
    destination "luna1" now accepting requests
# enable luna1 6
printer "luna1" now enabled
# lpadmin -p luna1 -D "Room 1954 ps" 7
# lpstat -p luna1 8
    printer luna1 is idle. enabled since Jul 12 11:17 2001. 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 directly, and the printer type.
5. Adds print filters to the print server.
6. Accepts print requests for the printer and enables the printer.
7. Adds a description for the printer.
8. Verifies that the printer is ready.

Converting Printer Configuration Information

This section explains how to convert the printer configuration information from systems running the SunOS 5.5.1 release and copy this information to print clients so they can access existing printers.

Note – If you have only a few existing printers, it might be easier to add access to the printers by using Solaris Print Manager rather than convert the printer configuration information and distribute it to print clients. For information on adding access to printers, see “Setting Up Printing (Task Map)” on page 49.

Converting Printer Configuration Information (Task Map)

The following task map identifies the tasks associated with converting printer configuration information.

Task	Description	For Instructions
Convert Existing Printer Configuration Information	Convert Printer Configuration Information for Systems Running the SunOS 5.5.1 Release – If your site uses the SunOS 5.5.1 release, you can convert the printer configuration information in the <code>/etc/lp/printers</code> directory to the <code>/etc/printers.conf</code> configuration file. This is usually a one-time task.	“How to Convert Printer Information For a System Running the Solaris 2.5.1 Release” on page 69

Task	Description	For Instructions
	Convert Printer Configuration Information for a System Running an lpd-based Print Server – If your site uses lpd-based software, you convert the printer configuration information in the <code>/etc/printcap</code> file to the <code>/etc/printers.conf</code> configuration file. This is usually a one-time task.	“How to Convert Printer Information For a System Running an lpd-based Print Server” on page 69

Converting Existing Printer Configuration Information

Existing printer configuration information is automatically converted when installing or upgrading to the Solaris 9 release. This section explains how to convert the printer configuration information for a system running SunOS 5.5.1 release or a system running a lpd-based printer server to the `/etc/printers.conf` printer configuration file. You'll use one of two print administration commands to automate the conversion task:

- The `conv_lp` command enables you to convert information in the `/etc/lp/printers` directory on a SunOS 5.9 system to entries in the system's `/etc/printers.conf` file. For instructions, see “How to Convert Printer Information For a System Running the Solaris 2.5.1 Release” on page 69.
- The `conv_lpd` command enables you to convert information in a `/etc/printcap` configuration file from a system running a lpd-based printer server to entries in a `/etc/printers.conf` file. For instructions, see “How to Convert Printer Information For a System Running an lpd-based Print Server” on page 69.

If you are not using a name service, you should create a master `/etc/printers.conf` file that includes the existing printers at your site. You can then copy the master file to all the print clients or by loading it into the NIS or NIS+ name service. This is a good way to initially enable all the new print clients access to the existing printers at your site.



Caution – If you are using the NIS or NIS+ name service to configure printer information, do not use a `/etc/printers.conf` file on your print clients. A print client uses the `/etc/printers.conf` file first to locate a printer; however, the `/etc/printers.conf` file might conflict with the printer information in the NIS or NIS+ maps and cause unexpected results. To avoid this problem, remove the `/etc/printers.conf` file on print clients when you want them to use NIS or NIS+ for printer information.

▼ How to Convert Printer Information For a System Running the Solaris 2.5.1 Release

1. Log in as superuser on the system running the Solaris 9 release.
2. Convert the printer configuration information in the system's `/etc/lp/printers` directory to the `/etc/printers.conf` file.

```
# /usr/lib/print/conv_lp
```

▼ How to Convert Printer Information For a System Running an lpd-based Print Server

1. Copy the `/etc/printcap` file from a system running an lpd-based print server to a system running the Solaris 9 release.
2. Log in as superuser on the system running the Solaris 9 release where you copied the `/etc/printcap` file.
3. Convert the printer configuration information in the `/etc/printcap` file to the `/etc/printers.conf` file.

```
# /usr/lib/print/conv_lpd
```


Administering Printers (Tasks)

This chapter describes the procedures for administer printers with the LP commands.

This is a list of the step-by-step instructions in this chapter.

- “How to Delete a Printer and Remote Printer Access” on page 72
- “How to Check the Status of Printers” on page 75
- “How to Stop the Print Scheduler” on page 76
- “How to Restart the Print Scheduler” on page 77
- “How to Add a Printer Description” on page 77
- “How to Set a System’s Default Printer” on page 78
- “How to Make Banner Pages Optional” on page 80
- “How to Turn Off Banner Pages” on page 81
- “How to Define a Class of Printers” on page 82
- “How to Set Fault Alerts for a Printer” on page 84
- “How to Set Printer Fault Recovery” on page 86
- “How to Limit User Access to a Printer” on page 88
- “How to Check the Status of Print Requests” on page 90
- “How to Accept or Reject Print Requests for a Printer” on page 94
- “How to Enable or Disable a Printer” on page 92
- “How to Cancel a Print Request” on page 95
- “How to Cancel a Print Request From a Specific User” on page 96
- “How to Move Print Requests to Another Printer” on page 97
- “How to Change the Priority of a Print Request” on page 99

For overview information about printing and the LP print service, see Chapter 2.

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.

Not only does the printer information need to be deleted from the print server, but it also needs to be deleted 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 print entry from the print clients or network name service and add access to the remote printer in its new location.

For detailed information on how to delete a local and remote printer, see “How to Delete a Printer and Remote Printer Access” on page 72. You can use Solaris Print Manager to delete a local or remote printer. However, Solaris Print Manager does not enable you to move queued print requests to another printer.

▼ How to Delete a Printer and Remote Printer Access

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

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 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 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 or 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 94.

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 92.

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 97.

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 ...]
```

<code>-r</code>	Removes the specified print client.
<code>print-client</code>	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 quotes.

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
```

You should receive an error indicating that the printer does not exist in the output of the above command.

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

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

You should receive an error indicating that the printer does not exist in the output of the above command.

Example—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`, and also 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 and examine the characteristics of those printers. You can use the `lpstat` command to find out status information about the LP print service or a specific printer.

▼ How to Check the Status of Printers

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 `lpstat(1)`.

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

<code>-d</code>	Shows the system's default printer.
<code>-p <i>printer-name</i></code>	Shows if a printer is active or idle, when it was enabled or disabled, and whether it 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 quotes. If you don't specify <i>printer-name</i> , the status of all printers is displayed.
<code>-D</code>	Shows the description of the specified <i>printer-name</i> .
<code>-l</code>	Shows the characteristics of the specified <i>printer-name</i> .
<code>-t</code>	Shows status information about the LP print service, including the status of all printers: whether they are active and whether they are accepting print requests.

Examples—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 Jul 12 11:35 2001. available.
unable to print: paper misfeed jam

Description: Printer by break room
printer luna is idle. enabled since Jul 12 11:36 2001. 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 ...
  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:
```

Restarting the Print Scheduler

The print scheduler, `lpsched`, handles print requests on print servers. However, there might be times when the print scheduler stops running on a system, so print requests stop being accepted or printed.

The following section describes how to 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

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

```
# lpstat -r
```

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

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

```
# /etc/init.d/lp stop
```

▼ How to Restart the Print Scheduler

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

```
# lpstat -r
```

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.

```
# /etc/init.d/lp start
```

Setting or Resetting Miscellaneous Printer Definitions

This section provides step-by-step instructions on setting or resetting printer definitions. Some of the following printer definitions can be set using Solaris Print Manager. The following procedures use the LP commands to quickly set or reset printer definitions.

▼ How to Add a Printer Description

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

Name of the printer for which you are adding a description.

-D "comment"

Specifies the characteristics of the printer, such as location or administrative contact. Enclose characters that the shell might interpret (like *, ?, \, !, ^) in single quotation marks.

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

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—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 the user doesn't 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 5-1 Search Order for Default Printer Destinations

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

▼ How to Set a System's Default Printer

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` 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 entered in the system's `/etc/lp/default` file.

3. Check the system's default printer.

```
$ lpstat -d
```

Example—Setting a System's Default Printer

The following example shows how to set the printer `luna` as the system's default printer. This means that `luna` is used as the system's default printer if the `LPDEST` or `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 will also have 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. Printing banner pages uses more paper, however, and might not be necessary 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, like 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, you can turn off banner pages for a printer so they are never printed. This is important if you have a situation where you don't need or want banner pages. You can turn off banner page printing by using the `lpadmin` command.

TABLE 5-2 Banner Page Printing

Command Used	Banner Page Printing	Override?
<code>lpadmin -p printer -o banner</code> or <code>lpadmin -p printer -o banner=always</code>	Required and printed	If you are a regular user and use <code>p -o nobanner</code> , the request is printed, but the <code>nobanner</code> argument is ignored. If you are root or another privileged user, the <code>nobanner</code> argument is honored.
<code>lpadmin -p printer -o nobanner</code> <code>lpadmin -p printer -o banner=optional</code>	On by default, but can be disabled on a per request basis with the <code>lp -o nobanner</code> command.	N/A
<code>lpadmin -p printer -o banner=never</code>	Disabled	No

For step-by-step command-line instructions, see “How to Turn Off Banner Pages” on page 81.

▼ How to Make Banner Pages Optional

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
```

<code>-p printer-name</code>	Name of the printer for which you are making banner pages optional.
<code>-o banner=optional</code>	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 entered in the print server’s `/etc/lp/printers/printer-name/configuration` file.

3. Verify that the output from the following command contains the line `Banner not required`.

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


Example—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

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` Name of the printer for which you are making banner pages optional.

`-o banner=never` Disables banner page printing under all circumstances.

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

3. Verify that the output from the following command contains the line `Banner not printed`.

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

4. Submit a print request to the printer to make sure a banner page does not print.

Example—Turning Off Printing Banner Pages

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

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

Setting Up Printer Classes

The 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.

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

There are no default printer classes 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, PostScript)
- By location (for example, 5th floor)
- By work group or department (for example, Accounting)

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, you would add it to the class before you add a low-speed printer. As a result, the high-speed printer would handle as many print requests as possible. The low-speed printer would be 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, like 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 them only if you determine that using printer classes would benefit users on the network.

▼ How to Define a Class of Printers

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* Name of the printer you are adding to a class of printers.

-c *printer-class* Name of a class of printers.

The specified printer is added to the end of the list in the class 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—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 print service can notify you when it detects a printer fault. You can select any of the following methods to receive printer fault notification with the `lpadmin -A` command or with Solaris Print Manager:

- Write 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. It 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 says 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 5-3 Values for Printing Problem Alerts

Value for <code>-A alert</code>	Description
<code>'mail [user-name]'</code>	Send the alert message by email to root or lp on the print server, or the specified <i>user-name</i> , which is a name of a user.
<code>'write [user-name]'</code>	Send the alert message to the root or lp console window on the print server, or to the console window of the specified <i>user-name</i> , which is a name of a user. The specified user must be logged in to the print server to get the alert message.
<code>'command'</code>	Run the <i>command</i> file for each alert. The environment variables and current directory are saved and restored when the file is executed.
<code>quiet</code>	Stop alerts until the fault is fixed. Use this when you (root or specified user) receive repeated alerts.

TABLE 5-3 Values for Printing Problem Alerts (Continued)

Value for <code>-A alert</code>	Description
none	Do not send any alerts. This is the default if you don't specify fault alerts for a printer.

▼ How to Set Fault Alerts for a Printer

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]
```

<code>-p printer-name</code>	Name of the printer for which you are specifying an alert for printer faults.
<code>-A alert</code>	Specifies what kind of alert will occur when the printer faults. For detailed information about the valid values for <i>alert</i> , see Table 5-3. Some valid values are <i>mail</i> , <i>write</i> , and <i>quiet</i> .
<code>-W minutes</code>	Specifies how often (in minutes) the fault alert will occur. If you don't specify this option, the alert is sent once.

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

3. Check the information following the `On fault` heading from the output of the following command.

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

Examples—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`, with reminders every 5 minutes.

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

The following example shows how to sets up the printer `venus` to send fault alerts to the console window, with reminders 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 find out about printing faults so 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 tell the system administrator 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

A print filter is required to continue printing from the top of a page where the printing stopped. A print filter records the control sequences used by the printer to track page boundaries, which the default filters used by the print service cannot do. You will be notified by the 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 146.

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 you can set for a printer with the `lpadmin -F` command.

TABLE 5-4 Values for Printer Fault Recovery

Value for <code>-F recover-options</code>	Description
<code>beginning</code>	After a fault recovery, printing restarts from the beginning of the file.
<code>continue</code>	After a fault recovery, printing starts at the top of the page where the printing stopped. This recovery option requires a print filter.

TABLE 5-4 Values for Printer Fault Recovery (Continued)

Value for <code>-F recovery-options</code>	Description
<code>wait</code>	After a fault recovery, printing stops until you enable the printer. After you enable the printer with the <code>enable</code> 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

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
```

<code>-p printer-name</code>	Name of the printer for which you are specifying fault recovery.
<code>-F recovery-options</code>	One of the three valid recovery options: beginning, continue, or wait. For detailed information about the valid values for <i>recovery-options</i> , see Table 5-4.

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

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

3. Check the information following the `After fault` heading in the output of the following command.

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

Example—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 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 can create allow and deny lists using the `lpadmin -u` command on the print server. (Solaris Print Manager enables you to create only allow lists.) If you create neither, a printer is available to all users who can access the printer.

An allow list contains the names of users 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:

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 <code>all</code> in the allow list	All users can access the printer.
Specify <code>all</code> in the deny list	All users, except root and lp (on the server), are denied access to the printer.
Make any entry in the allow list	The deny list is ignored. Only those 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 actually controlling access to the printer, allow and deny lists can only be created on the print server itself. If you create allow and deny lists, the print server will exclusively control user access to printers.

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

TABLE 5-5 Values for Allow and Deny Lists

Value for <i>user-list</i>	Description
<i>user</i>	User on any system
<code>all</code>	All users on all systems
<code>none</code>	No user on any system
<i>system !user</i>	User on <i>system</i> only
<code>!user</code>	User on local system only

TABLE 5-5 Values for Allow and Deny Lists (Continued)

Value for <i>user-list</i>	Description
<code>all!user</code>	User on any system
<code>all!all</code>	All users on all systems
<code>system!all</code>	All users on <i>system</i>
<code>!all</code>	All users on local system

▼ How to Limit User Access to a Printer

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]
```

<code>-p printer-name</code>	Name of the printer to which the allow or deny user access list applies.
<code>-u allow:user-list</code>	User names to be added to the allow 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 quotes. Table 5-5 provides the valid values for <i>user-list</i> .
<code>-u deny:user-list</code>	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 quotes. Table 5-5 provides the valid values for <i>user-list</i> .

The specified users are added to the allow or deny user access 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 user access list, the following files are not created for 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. Check the information following the `Users allowed` or `Users denied` heading in the output of the following command.

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

Examples—Limiting User Access to a Printer

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

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

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

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

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 login to the print server to move, hold, resume, or change the priorities of print requests with LP 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 with the `lp H` command.

TABLE 5-6 Values for Changing the Priority of a Print Request

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

▼ How to Check the Status of Print Requests

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 `lpstat(1)`.

```
$ 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.

You can 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 quotes.

If you don't specify *list*, the status of print requests 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.

You can 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 quotes.

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 request ID, the user, the output size in bytes, the date and time of the request, and information about the request, such as “being filtered.”

Examples—Checking the Status of Print Requests

The following example shows that user `fred` has one print request queued to 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 there are two print requests 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 queue. However, requests are still added to the queue. (You must set the printer to reject print requests so requests are not added to the queue. For information about rejecting print requests, see “Accepting or Rejecting Print Requests” on page 93.)

A printer is enabled to print and accepts print requests when it 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 processing of print requests is interrupted when a printer is disabled.




FIGURE 5-1 What Happens When a Printer Is Enabled or Disabled

▼ How to Enable or Disable a Printer

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
```

<code>disable</code>	Cancels the current job, then disables the printer. The current job is saved to reprint when the printer is enabled.
<code>-c</code>	Cancels the current job, then disables the printer. The current job is not printed later.
<code>-W</code>	Waits until the current job is finished before disabling the printer.
<code>-r "reason"</code>	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 with the <code>lpstat -p</code> command.
<code>printer-name</code>	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
```

Examples—Enabling or Disabling a Printer

The following example shows how to stop the current job on the printer `luna`, save it 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. 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 5-7 Functions of `accept/reject` and `enable/disable` Commands

Command	Function
<code>accept</code>	Accept print requests that are sent to the print queue.
<code>enable</code>	Print the requests that are in the print queue.
<code>reject</code>	Reject print requests that are sent to the print queue.
<code>disable</code>	Stop printing requests that are currently in the print queue.

If a print request is rejected, the print service writes or mails a message to the user who submitted the request, saying 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 users' systems when one tries to check the printer's queue. The following figure shows the point at which processing of print requests is interrupted when a print queue rejects print requests.




FIGURE 5-2 What Happens When a Print Queue Accepts or Rejects Requests

▼ How to Accept or Reject Print Requests for a Printer

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 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 with the `lpstat -p` command.

printer-name

Name of the printer that will stop accepting print requests.

The queued requests will continue printing as long as the printer is enabled. For instructions on disabling a printer so it stops printing, see “How to Enable or Disable a Printer” on page 92.

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
```

Examples—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 from printer queues or to cancel jobs that are printing. There are three ways to use the `cancel` command:

- Cancel requests by request identification number (request ID)
- Cancel requests from a specific user on all, or specified, printers
- Cancel the job currently printing

When you use the `cancel` command, a message tells you the request(s) are canceled, and the next request in queue is printed. You can cancel a print request only if you are:

- The user who submitted the request and you are logged in on the system from which you submitted the request
- 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.
- Logged in as superuser or `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 is comprised of 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

1. If you are going to cancel print requests of other users, 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 90.

3. Cancel a print request.

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

request-id 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 quotes.

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 quotes.

4. Verify that the print requests are canceled.

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

Examples—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

1. (Optional) Become superuser, lp, or assume an equivalent role if you are going to cancel print requests of other users.
2. Cancel a print request from a specific user.

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


<i>-u user-list</i>	<p>Cancels the print request for a specified user.</p> <p><i>user-list</i> 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 quotes.</p>
<i>printer-name</i>	<p>Specifies the printer for which you want to cancel the specified user's print requests.</p> <p><i>printer-name</i> 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 quotes.</p> <p>If you don't specify <i>printer-name</i>, the user's print requests will be canceled on all printers.</p>

Examples—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, and 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 file content type or forms) that cannot be met by the newly specified printer cannot be moved. They must be canceled.

▼ How to Move Print Requests to Another Printer

To move all print requests from one printer to another, you do not need to know the request IDs. However, it is a good idea to see how many print requests are affected before you move them.

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

2. **(Optional) Check 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` 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` Name of the printer from which all print requests will be moved.

`printer-name2` Name of the printer to which all print requests will be moved.

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

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

5. **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
```

Make sure all specified print requests were moved to the destination printer's queue by using the following command.

```
$ lpstat -o printer-name2
```

Example—Moving Print Requests to Another Printer

The following example shows how to move print requests from the printer `luna` to the printer `terra`, and then tells the original printer `luna` 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:

- Putting any print request on hold if it has not finished printing. Putting a request on hold stops it, 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 be 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, moving it in the queue so it is ahead of lower priority requests and behind requests at the same level or at a higher priority.

▼ How to Change the Priority of a Print Request

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 90.

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

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

-i request-id

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 quotes.

-H change-priority

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 5-6.

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—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
```

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

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

This is a list of the step-by-step instructions in this chapter.

- “How to Define a Print Wheel or Font Cartridge” on page 104
- “How to Unmount and Mount a Print Wheel or Font Cartridge” on page 105
- “How to Set an Alert to Mount a Print Wheel or Font Cartridge” on page 106
- “How to Set Up an Alias for a Selectable Character Set” on page 108
- “How to Add a Print Filter” on page 111
- “How to Delete a Print Filter” on page 112
- “How to View Information About a Print Filter” on page 112
- “How to Add a Form” on page 116
- “How to Delete a Form” on page 117
- “How to Unmount and Mount a Form” on page 117
- “How to Set an Alert to Mount a Form” on page 119
- “How to View Information About a Form” on page 120
- “How to View the Current Status of a Form” on page 120
- “How to Limit User Access to a Form” on page 121
- “How to Limit Printer Access to a Form” on page 122
- “How to Install Downloaded PostScript Fonts” on page 125
- “How to Install Host-Resident PostScript Fonts” on page 125

For overview information about printing, see Chapter 2.

Managing Character Sets

Printers differ in the method 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, and 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, using either print wheels, font cartridges, or preprogrammed selectable character sets. Usually, only one of these printing methods applies to a given printer type.

Print wheels and font cartridges, from the perspective of the LP print service, are similar, because someone must intervene and mount the hardware on the printer, when needed. Character sets that require you to physically mount a wheel or cartridge are referred to as *hardware character sets*. Character sets that do not require hardware mounting, that come preprogrammed with the printer, and can be selected by a print request, 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:

```
tput -T printer-type csn
```

The `csn` 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. You can repeat the command, using `-1`, `-2`, `-3`, and so on in place of the `-0`, to display the names of the other character sets. 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 used in the manufacturer's documentation for the printer. Because manufacturers do not all use the same character set names, the `terminfo` character set names can differ from one printer type to the next.

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.

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.

Note – Character sets for PostScript printers are not listed when you use the `lpstat -p -l` command because 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 123.

Hardware-Mounted 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 names you want to use for the available print wheels, and how you want to be alerted when a printer needs a different print wheel. Then, when a user requests a particular character set with the `lp -S` command, the scheduler sends an alert to mount the print wheel, and 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 105.

If you do not specify multiple print wheels or cartridges for a printer, the LP print service assumes that the printer has a single, fixed print wheel or cartridge, and users cannot specify a special print wheel or cartridge when using the printer.

Unlike selectable character sets, the names you use for print wheels or cartridges are not tied to entries in the `terminfo` database. Print wheel or cartridge names are used only for the purpose of communicating with the LP print service and its users.

The names you choose for print wheels or cartridges, however, 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 or cartridges, or selectable character sets. That way, users can ask for a font style (character set) without regard to which printer—or even whether a print wheel or cartridges—or selectable character set will be used.

Of course, you and the printer users should agree on the meanings of print wheel or 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, such as print wheels or print cartridges, that provide a particular font or character set. 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 Cartridges

You request alerts for mounting print wheels or cartridges in the same way you request other alerts from the LP print service. For general information about alerts, see “Setting Up Printer Fault Alerts” on page 83.

▼ How to Define a Print Wheel or Font Cartridge

1. **Log in as superuser, lp, or assume an equivalent role on the print server.**
2. **Define a print wheel or font cartridge that can be used with the printer.**

```
print-server# lpadmin -p printer-name -s hard-charset1 [,hard-charset2...]
```

<code>-p printer-name</code>	Name of the printer for which you are defining a print wheel or font cartridge.
<code>-s hard-charset</code>	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 quotes. Define names that are meaningful to users, and inform the users of the names.

The print wheel or font cartridge definition is added 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 those in Step 2.

The print wheel or font cartridge definition is added in 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 information following the `Character sets` heading in the following output is correct on both the print server and the print client.

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

Example—Defining a Print Wheel

The following example shows how to define the `pica` print wheel on the printer `luna` for a print client named `asteroid`.

```
asteroid# lpadmin -p luna -S pica
```

▼ How to Unmount and Mount a Print Wheel or Font Cartridge

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 in the printer.

```
# lpadmin -p printer-name -M -S none
```

`-p printer-name` Printer on which you are unmounting a print wheel or font cartridge.

`-M -S none` Specifies unmounting the current print wheel or font cartridge.

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

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
```

<code>-p printer-name</code>	Printer on which you are mounting a print wheel or font cartridge.
<code>-M -S hard-charset</code>	Hardware character set name of the print wheel or font cartridge you want to mount.

The print wheel or font cartridge is added in 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 that 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—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

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]
```

<code>-S hard-charset</code>	Hardware character set name of the print wheel or font cartridge for which you want to set an alert.
------------------------------	--

<code>-A alert</code>	Specifies what kind of alert will occur when a print wheel or font cartridge is requested. For detailed information about the valid values for <i>alert</i> , see Table 5-3. Some valid values are <i>mail</i> , <i>write</i> , and <i>quiet</i> . If you specify <i>mail</i> or <i>write</i> , a predefined alert message says to mount the specified print wheel or font cartridge and includes the names of one or more printers that have been set up to use such a print wheel or cartridge.
<code>-Q requests</code>	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.
<code>-W minutes</code>	Specifies how often (in minutes) the alert will occur. If you don't specify this option, the alert is sent only once.

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

The alert is added in 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
```

Otherwise, if you have set a low number of print requests to trigger the alert, submit enough print requests to meet the minimum requirement and make sure you receive an alert to mount the print wheel or font cartridge.

Examples—Setting an Alert to Mount a Print Wheel or Font Cartridge

The following example shows how to set email alerts to occur every five minutes for the `elite` print wheel when there are ten print requests for `elite` in the print queue.

```
# lpadmin -S elite -A mail -Q 10 -W 5
```

The following example shows how to set email alerts to occur every minute for the `finnish` font cartridge when there are five print requests for `finnish` in the print queue.

```
# lpadmin -S finnish -A mail -Q 5 -W 1
```

The following example show how to set console-window alerts to occur every 10 minutes for the `elite` print wheel when there are five print requests for `elite` in the print queue.

```
# lpadmin -S elite -A write -Q 5 -W 10
```

The following example shows how to set no alerts to occur 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 129.

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
```

<code>-T printer-type</code>	Printer type found in the <code>terminfo</code> database. For information on entries in the <code>terminfo</code> database, see “Printer Type” on page 44.
<code>n</code>	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, <code>cs1</code> could cause the system to display <code>english#</code> .

For more information, see `tput(1)`.

3. Set up an alias for a selectable character set.

```
# lpadmin -p printer-name -S select-charset1=alias1[,select-charset2=alias2...]
```

<code>-p printer-name</code>	Printer on which you are setting up aliases for selectable character sets.
<code>-S select-charset</code>	Selectable character set name for which to set an alias. The name can be found in Step 2.
<code>alias</code>	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 quotes.

The alias is added in 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 those in Step 3.

The alias is added in 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. Verify that the selectable character set alias is listed in the output of the following command on the print server and print clients.

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

Otherwise, submit a print request that uses the alias for the selectable character set and check for output.

Example—Setting Up an Alias for a Selectable Character Set

The following example shows how to display the names of selectable character sets and specify `text` as an alias for the `usascii` selectable character set on the printer `luna`, which is an `ln03` printer type.

```
# 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

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:

- 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, like two-sided printing, landscape printing, or draft- and letter-quality printing
- Detect printer faults and notify the LP print service of them so 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 6-1. The filter programs are located in the `/usr/lib/lp/postscript` directory. For PostScript printing, you usually do not need to do anything beyond installing 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 background information and step-by-step instructions, see “Creating a New Print Filter” on page 137.

Adding, Changing, Removing, and Restoring Print Filters

Print filters are added, changed, or removed 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 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 one 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, and remove any filters you have added.

SunOS software provides a default set of PostScript filters, which Solaris Print Manager automatically adds to a print server. Some of the TranScript filters used with `lpd`-based print services have SunOS equivalents, but others do not. The following table lists the default PostScript filters and identifies the TranScript filters, where applicable.

TABLE 6-1 Default PostScript Filters

Filter	Action	TranScript Equivalent
download	Download fonts	
dpost	ditroff to PostScript	psdit
postdaisy	daisy to PostScript	
postdmd	dmd to PostScript	
postio	Serial interface for PostScript printer	pscomm
postior	Communicate with printer	
postmd	Matrix gray scales to PostScript	
postplot	plot to PostScript	psplot
postprint	simple to PostScript	enscript
postreverse	Reverse or select pages	psrev
posttek	TEK4014 to PostScript	ps4014

The SunOS software does *not* provide the following filters:

- TEX
- oscar (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 146 for information about writing a custom print filter.

▼ How to Add a Print Filter

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-def
```

-f *filter-name* Name you choose for the print filter.

-F *filter-def* Name of the print filter definition.

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

The print filter is added in the print server's `/etc/lp/filter.table` file.

3. **Verify that the print filter was added by checking for information about the print filter in the output of the following command.**

```
# lpfilter -f filter-name -l
```

Example—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

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` 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 filter was deleted.**

```
# lpfilter -f filter-name -l
```

You should receive an error indicating that no filter by the specified name exists.

Example—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

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
```


<code>-f <i>filter-name</i></code>	Print filter for which you want to view information. Specify all for <i>filter-name</i> to view information about all the available print filters.
<code>-l</code>	Displays information about the specified filter.

Information about the specified print filter(s) is displayed.

Examples—Viewing Information About a Print Filter

The following example shows how to request information for the `postdaisy` print filter, and the information that is displayed in response.

```
# 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
```

The following example shows how to redirect information about the `daisytroff` filter to a file (create the filter definition for that filter). This is useful if a filter definition is removed unintentionally.

```
# lpfilter -f daisytroff -l > daisytroff.fd
```

The following example shows how to display all the print filters that have been added to the system, and the information that is displayed in response.

```
# 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
```

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, removing, mounting, and controlling 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 enter 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.

Mounting Forms

To print a form, you must load the paper in the printer and use a command to *mount* the form, which 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:

- 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, make sure 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 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.

Tracking Forms

The LP print service helps you track which forms are mounted on each printer and 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, 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 root, 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 you request other alerts from the LP print service. For general information about alerts, see “Setting Up Printer Fault Alerts” on page 83.

Checking Forms

When you have defined a form for the LP print service, you can check it with either of two commands, depending on the type of information you want to check.

- Show the attributes of the form by using the `lpforms` command. You can also redirect the output of the command into a file to save it for future reference.
- Display the current status of the form by using 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 list the contents of the `/etc/lp/forms` directory to see the names of the forms there.

Limiting Access to Forms

You can control which printers and 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.

To limit user access to forms, see “How to Limit User Access to a Form” on page 121. To limit printer access to a form, see “How to Limit Printer Access to a Form” on page 122.

▼ How to Add a Form

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 *form-name* Name you choose for the form.

-F */etc/lp/forms/form* Name of the form definition.

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

The form is added in the print server's `/etc/lp/forms/form-name/describe` file.

3. Verify that the form was added by checking for a listing of information about the form in the output of the following command.

```
# lpforms -f form-name -l
```

Example—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
```

Note – 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 122.

▼ How to Delete a Form

1. Log in as superuser, lp, or assume an equivalent role on the print server.
2. Delete the form.

```
# lpforms -f form-name -x
```

<code>-f form-name</code>	Form to be deleted.
<code>-x</code>	Deletes the specified form.

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

The form is deleted from `/etc/lp/forms/form-name` file.

3. Verify that 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—Deleting a Form

The following example shows how to delete the `medical` form.

```
# lpforms -f medical -x
```

▼ How to Unmount and Mount a Form

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
```

<code>printer-name</code>	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 as 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 form paper for the next print request.**
6. **Mount the form.**

```
# lpadmin -p printer-name -M -f form-name [-a -o filebreak]
```

-p <i>printer-name</i>	Printer on which you are mounting a form.
-M -f <i>form-name</i>	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 it has one.

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 by checking for the form name under the Form mounted heading in the output of the following command.**

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

Otherwise, submit a print request that requires the new form and check the printer for output.

Examples—Unmounting and Mounting 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
```

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
```

▼ How to Set an Alert to Mount a Form

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]
```

<code>-f form-name</code>	Form for which you want to set a request alert.
<code>-A alert</code>	Specifies what kind of alert will occur when a form is requested. See Table 5-3 for detailed information about the valid values for <i>alert</i> . Some valid values are <code>mail</code> , <code>write</code> , and <code>quiet</code> . If you choose <code>mail</code> or <code>write</code> , a predefined alert message says to mount the specified form and includes the names of one or more printers that have been set up to use the form.
<code>-Q requests</code>	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.
<code>-W minutes</code>	Specifies how often (in minutes) the alert will occur. If you don't specify this option, the alert is sent once.

The request alert is added in the print server's `/etc/lp/forms/form-name/alert.sh` file.

3. Verify that the alert has been added for the form by checking the output of the following command.

```
# lpforms -f form-name -A list
```

Otherwise, if you have set a low number of print requests to trigger the alert, submit print requests to meet the minimum requirement and make sure you receive an alert to mount the form.

Examples—Setting an Alert to Mount a Form

The following example shows how to set email alerts to occur every five minutes for the `letterhead` form when there are 10 print requests for `letterhead` in the print queue.

```
# lpforms -f letterhead -A mail -Q 10 -W 5
```

The following example shows how to set console window alerts to occur every 10 minutes for the `letterhead` form when there are five requests for `letterhead` in the print queue.

```
# lpforms -f letterhead -A write -Q 5 -W 10
```

The following example shows how to set no request alerts for the `invoice` form.

```
# lpforms -f invoice -A none
```

▼ How to View Information About a Form

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 form-name` 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(s) is displayed.

Examples—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 is useful if a form definition gets removed unintentionally.

```
# lpforms -f medical -l > medical.fmd
```

▼ How to View the Current Status of a Form

1. Log in on the print server.
2. Request information about the current status of a form.


```
$ lpstat -f form-name
```

`-f form-name` 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(s) is displayed.

Example—Viewing the Current Status of a Form

the following example shows how to display the status of the `medical` form.

```
$ lpstat -f medical
form medical is available to you
```

▼ How to Limit User Access to a Form

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 form-name` Name of the form for which the allow or deny user access list is being created.

`-u allow:user-list` Represents users to be added to the allow access list. Use a comma or a space to separate users' login IDs. If you use spaces, enclose the list of IDs in quotes.

Table 5-5 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 users' login IDs. If you use spaces, enclose the list of IDs in quotes. Table 5-5 provides the valid values for *user-list*.

The specified user(s) are added to the allow or deny user access list for the specified form in one of the following files on the print server:

```
/etc/lp/forms/form-name/allow or
/etc/lp/forms/form-name/deny
```

3. Verify the allow and deny user access lists.

```
# lpforms -f form-name -l
```

Examples—Limiting User Access to a Form

The following example shows how to allow only the users `nathan` and `marcia` access to the `check` form.

```
# lpforms -f check -u allow:nathan,marcia
```

The following 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

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
```

<code>-p printer-name</code>	Name of the printer for which the allow or deny forms list is being created.
<code>-f allow:form-list deny:form-list</code>	Form names to be added to the allow or 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 quotes.

The specified form(s) are added to the allow or 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 the allow and deny forms lists.

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

Examples—Limiting Printer Access to a Form

The following 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
```

The following 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

If you have a laser printer, you might need to install and maintain PostScript fonts. You might also have to decide where to install PostScript fonts and how to manage them. 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, it 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 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 there are numerous available fonts, 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 are to be used by the users on one system and the 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. It also supplies `troff` width tables for the 35 standard PostScript fonts which reside on many PostScript printers, for use by the `troff(1)` program.

Managing Printer-Resident Fonts

Most PostScript printers come equipped with fonts 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 the file:

```
/etc/lp/printers/printer-name/residentfonts
```

When the printer is attached to a print server, make sure the list in the `residentfonts` file includes fonts that are on the print server and which are available for downloading to the printer.

You must edit the files containing 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, and it must have access to the printer while the filter is processing. A slow filter takes longer to convert a file, and it does not need to access the printer while the filter is processing. An example of a slow filter is ASCII to PostScript.

The download filter is a fast filter; it 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 this, 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. It 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. It 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 (by getting the appropriate file name from the map table) to determine if the requested font is available.
4. If the font is available, the filter takes the file for that font and appends it to the file to be printed.

5. It sends the font definition file and the source file (the file to be printed) to the PostScript printer.

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 that shows the correspondence between the names of fonts and the names of the files containing those fonts. Both the map 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 you create called `/usr/share/lib/hostfontdir/typeface/font`, where *typeface* is replaced by a name like *palatino* or *helvetica*, and *font* is replaced by a name like *bold* or *italic*.

▼ How to Install Downloaded PostScript Fonts

1. **Log in as superuser, lp, 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
```

printer-name Name of the printer on which you want to install downloaded PostScript fonts.

3. **Create the `residentfonts` file, if it does not already exist.**

```
# touch residentfonts
```

This file might not exist if this is the first time you are adding permanently downloaded fonts.
4. **Edit the `residentfonts` file and add all the printer-resident fonts and fonts to be permanently downloaded.**

▼ How to Install Host-Resident PostScript Fonts

1. **Log in as superuser, lp, or assume an equivalent role on the print server or print client.**
2. **Create the `hostfontdir` directory, if it does not already exist.**

```
# cd /usr/share/lib
# mkdir hostfontdir
# 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 the name of the font and the name of the file in which it resides to the map table.

a. Change to the /usr/share/lib/hostfontdir directory.

b. Edit the map file using a text editor such as vi.

Add a one-line entry for each font you want to add to the table, with 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 table on the appropriate system, users will be able to apply the font (for example, Palatino Bold) in their print jobs. When they submit a print request containing this font, the LP print service appends a copy of the file /usr/share/lib/hostfontdir/palatino/bold to that file before sending it to the printer.

6. If you are using troff, you must create new width tables for this font in the standard troff font directory.

Customizing the LP Print Service (Tasks)

This chapter provides background information and procedures for customizing the LP print service.

This is a list of the step-by-step instructions in this chapter.

- “How to Adjust the Printer Port Characteristics” on page 128
- “How to Add a `terminfo` Entry for an Unsupported Printer” on page 132
- “How to Set Up a Custom Printer Interface Program” on page 136
- “How to Create a New Print Filter” on page 146
- “How to Create a New Form Definition” on page 150

For overview information about printers, see Chapter 2.

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 find out 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 used by the LP print service.

TABLE 7-1 `stty` Default Settings Used by the LP Print Service

Option	Meaning
<code>-9600</code>	Set baud to 9600

TABLE 7-1 stty Default Settings Used by the LP Print Service (Continued)

Option	Meaning
-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

▼ How to Adjust the Printer Port Characteristics

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 printer-name	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 stty(1). Table 7-1 shows the default stty settings used by the LP print service.

3. Verify that the printer port characteristics have been changed.


```
# stty -a
```

Examples—Adjusting the Printer Port Characteristics

The following example shows how to set the port characteristics for the printer `luna`. The `parenb` option enables parity checking/generation, `parodd` sets odd parity generation, and `cs7` sets the character size to 7 bits.

```
# lpadmin -p luna -o "stty='parenb parodd cs7'"
```

The following 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 a selected page size, character pitch, line pitch, and character set.

Each printer is identified in the `terminfo` database with a short name. The 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 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 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 there is no `terminfo` entry for your type of printer and you want to keep the printer set in the correct modes, you can either customize the interface program used with the printer or 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 7-2 Required terminfo Items for a Printer

Item	Meaning
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
cols	Number of columns in a line
cps	Average print rate in characters per second
it	Tabs initially every <i>n</i> 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	Change horizontal resolution
cpi	Change number of characters per inch
cr	Carriage return
csnm	List of character set names
cudl	Down one line
cud	Move carriage down <i>n</i> lines
cuf	Move carriage right <i>n</i> columns
cvr	Change vertical resolution
ff	Page eject
hpa	Horizontal position absolute
ht	Tab to next 8-space tab stop
if	Name of initialization file
iprogr	Path name of initialization program

TABLE 7-2 Required terminfo Items for a Printer (Continued)

Item	Meaning
is1	Printer initialization string
is2	Printer initialization string
is3	Printer initialization string
Strings:	
lpi	Change number of lines per inch
mgc	Clear all margins (top, bottom, and sides)
rep	Repeat a character <i>n</i> times
rwidm	Disable double-wide printing
scs	Select character set
scsd	Start definition of a character set
slines	Set page length to <i>n</i> lines per page
smgl	Set left margin at current column
smglp	Set left margin
smgr	Set right margin at current column
smgrp	Set right margin
smglr	Set both left and right margins
msgt	Set top margin at current line
smgtp	Set top margin
smgb	Set bottom margin at current line
smgbp	Set bottom margin
smgtb	Set both top and bottom margins
swidm	Enable double-wide printing
vpa	Vertical position absolute

▼ How to Add a `terminfo` Entry for an Unsupported Printer

Note – Before you create a `terminfo` entry for a printer, you should first make sure 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 there is one.

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 them as a guide for choosing a name for the printer.

3. **Create a `terminfo` entry file for the printer.**

Table 7-2 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 `terminfo(4)`.

To help you start writing a new `terminfo` entry, use the `infocmp` command to save an existing `terminfo` entry to a file. This is helpful if there is a `terminfo` entry that is similar to one 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
```

terminfo_entry 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. But first you need to understand what is in the standard program. The following section describes the standard program.

A printer interface program should:

- 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 gets the control sequences from the `terminfo` database and the `TERM` shell variable.
- Print a banner page, if necessary.
- Print the number of copies 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 5.9 or compatible LP print service. However, several `-o` options have been standardized in the SunOS 5.9 or compatible LP print service and 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, and the printer is identified as the “controlling terminal” for the printer interface program so that a “hang-up” of the port will cause a `SIGHUP` signal to be sent to the printer interface program.

The Standard Printer Interface Program

The standard (model) printer interface program, `/usr/lib/lp/model/standard`, is used by the LP print service to set the printing defaults shown in Table 7-3.

TABLE 7-3 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, like baud rate or output options, look for the section of the standard printer interface program that begins with the following comment:

```
## Initialize the printer port
```

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 7-4 Printer Interface Program Exit Codes

Code	Meaning to the LP Print Service
0	The print request has been successfully completed. If a printer fault occurred, it has been cleared.

TABLE 7-4 Printer Interface Program Exit Codes (Continued)

Code	Meaning to the LP Print Service
1 to 127	A problem was encountered when printing a request (for example, too many nonprintable characters or the request exceeds the printer capabilities). The LP print service notifies the person who submitted the request that there was an error when printing it. This error will not affect future print requests. If a printer fault has occurred, it has been cleared.
128	This code is reserved for internal 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 will 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 will not disable the printer, but it will try printing again in a few minutes.
>129	These codes are reserved for internal 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, but 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 zero exit code, just as if the fault had never occurred. An added advantage of this approach is that the interface program can detect when the fault is cleared automatically, so that 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, where it is put into the message that alerts the administrator to the printer fault. If its standard input is empty, `lp.tell` 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, but doing so will override 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 `kill(1)` and `signal(3C)`.) 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, and 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 so the signals are trapped at appropriate times. The standard interface interprets receipt of these signals as warnings that the printer has a problem; when it 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 in place 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

1. **Log in as superuser, lp, or assume an equivalent role on the print server.**
2. **Determine your next step based on whether you have a custom printer interface program.**

Custom Printer Interface Status	Action
Need to create a custom printer interface program	Go to Step 3.
Already have a custom printer interface program	Go to Step 5.

3. **Copy the standard printer interface program.**

```
# cp /var/spool/lp/model/standard custom-interface
```

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 133 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
```


<code>-p printer-name</code>	The printer that will use the custom printer interface program.
<code>-i custom-interface</code>	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.**

Examples—Setting Up a Custom Printer Interface Program

The following example shows how to set up a custom printer interface program named `custom` for the printer `luna`.

```
# lpadmin -p luna -i custom
```

The following example shows how to set up a custom printer interface program that the system `venus` is using on 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; 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 non-PostScript 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.

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. You can specify the accepted file content types for each printer. The user specifies the file content type when submitting a print request, and 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 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, and 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 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, and, 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 describing the reason for a fault. Other printers indicate a fault by using signals other than the signals indicating 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, and alerts are sent to the system administrator if alerts are enabled. For further information, see “Setting Up Printer Fault Alerts” on page 83.

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 like `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` allow

include files, they can fail when used as filters. If the program needs other files to complete its processing, it 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 non-zero 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. It 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 will stop the filter before alerting the administrator.
- The filter should send printer fault messages to its standard error as soon as the fault is recognized. It does not have to exit; it 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 non-zero exit code only if a printer fault has prevented it 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 it 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 or slow filter), and point to the filter programs (for example, `/usr/lib/lp/postscript/postdaisy`).

When defining a new print filter, in addition to writing a filter program, 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 it accepts

- Output types it produces
- Printer types to which it can send jobs
- Names of specific printers to which it 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. 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, but rather point to the filters.

Whether you store the information in a file, or enter it 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 will be 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, those shown in the following table are assigned by default. They are not very useful, which is why you should specify explicit values.

TABLE 7-5 Default Values for `lpfilter` Arguments

Item	Default
Input types	any
Output type	any
Printer types	any
Printers	any
Filter type	slow

Command

Use the full path of the filter program. If there are any fixed options that the program always needs, include them here.

Input Types

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 type. Several file types can be similar enough that the filter can deal with them. You can use whatever names you like, 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 they know how to identify the file type when submitting a file for printing.

Output Types

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. It is unlikely that you will 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.

Printer Types

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 it is identical to the list of output types. But it can be different. For example, you could have a printer with a single printer type for purposes of initialization, but which can recognize several different file content types. Essentially, this printer has an internal filter that converts the various file types into one 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 one 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 it as working only on type B printers.

Printers

A print filter is normally able to 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 or 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.

Filter Type

The LP print service recognizes fast and slow filters, as described in “Types of Filters” on page 138.

Slow filters that are invoked by printing modes (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.

Options

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 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 characteristics of the filter.

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 7-6 Print Filter Options Keywords

Characteristic	Keyword	Possible Patterns	Example
Content type (input)	INPUT	<i>content-type</i>	troff
Content type (output)	OUTPUT	<i>content-type</i>	postscript, impress
Printer type	TERM	<i>printer-type</i>	att495
Printer name	PRINTER	<i>printer-name</i>	lp1
Character pitch	CPI	<i>scaled-decimal</i>	10
Line pitch	LPI	<i>scaled-decimal</i>	6
Page length	LENGTH	<i>scaled-decimal</i>	66
Page width	WIDTH	<i>scaled-decimal</i>	80
Pages to print	PAGES	<i>page-list</i>	1-5, 13-20
Character set	CHARSET	<i>character-set</i>	finnish
Form name	FORM	<i>form-name</i>	invoice2
Number of copies	COPIES	<i>integer</i>	3
Special modes	MODES	<i>mode</i>	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 find out 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, it is 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 will reduce to one the number of copies of the filtered file it prints, since this "single copy" includes the multiple copies produced by the filter.
- The value of the `MODES` template comes from the `lp -y` command. Because a user can specify several `-y` options, there might be several values for the `MODES` template. The values will be applied in the left-to-right order given by the user.

The *replacement* part of a template shows how the value of a template should be given to the filter program. It is typically a literal option, sometimes with the placeholder asterisk (*) included to show where the value goes. The *pattern* and *replacement* 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, which can be used to extract portions of the *pattern* for copying into the *replacement*, and the `&`, which can be used to copy the entire *pattern* into the *replacement*. For more information, see `ed(1)`.

Note – If a comma or an equal sign (=) is included in a *pattern* or a *replacement*, precede it with a backslash (\). A backslash in front of any of these characters is removed when the *pattern* or *replacement* is used.

▼ How to Create a New Print Filter

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 138. 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 140. You should save the printer 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 111.

Examples—Creating a New Print Filter

The following example shows a print filter definition 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 will be handled as follows:

- When a user enters 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 enters 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
```

The following example shows a print filter definition to convert from `troff` to PostScript.

```
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\1
```

In this example, the filter program is named `dpost`. It 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 portrait mode or landscape mode. Because these options are not intrinsic to the LP print service, users must specify them 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 enters 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 luna ch1.doc
```

- 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 entering information about nine required characteristics (such as page length and page width) as input to the `lpforms` command. The LP print service uses this information to:

- 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 and underscores. The information must be in 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 phrase `[, mandatory]` means that the user cannot override the character set choice in the form. The *content-type* can be given, although this is optional, with an alignment pattern. If this attribute is given, the print service uses it to determine, as necessary, how to filter and print the file.

With two exceptions, the information can appear in any order. The exceptions are the `Alignment pattern` (which must always be last), and the *comment* (which must always follow the line with the `Comment :` prompt). If the comment contains a line beginning with a `>` character (like `Page length`, `Page width`, and so on), precede that line with a `>` character so 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 for the items listed in the following table the default values are assigned. Before running the `lpforms` command, gather the following information about the new form:

TABLE 7-7 Default Form Values

Item	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. This is also called leading. It 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.
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.

TABLE 7-7 Default Form Values (Continued)

Item	Default	Description
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 those logged in as root or lp can read it.

When you have gathered the information for the form, you enter it as input to the `lpforms` command. You should record this information first in a separate file so you can edit it before entering 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 Form Definition

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 print forms, see “Creating a New Printer Form” on page 148. 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 116.

LP Print Service (Reference)

This chapter provides background information on the LP print service.

- “The Structure of the LP Print Service” on page 152
- “LP Print Service Commands” on page 160
- “Functions of the LP Print Service” on page 161
- “How LP Administers Files and Schedules Local Print Requests” on page 161
- “Scheduling Network Print Requests” on page 162
- “Filtering Print Files” on page 163
- “What the Printer Interface Program Does” on page 163
- “How the `lpsched` Daemon Tracks the Status of Print Requests” on page 163
- “Cleaning Out Log Files” on page 164

Where to Find Printer Tasks

Printer Task	For More Information
Set up printers with Solaris Print Manager and the <code>lp</code> commands	Chapter 4
Administer printers with the <code>lp</code> commands after they are setup	Chapter 5
Manage character sets, filters, forms, and fonts	Chapter 6
Customize the LP print service	Chapter 7

The LP Print Service

The *LP print service* is a set of software utilities that allows users to print files while they continue to work. Originally, the print service was called the LP spooler. (LP stood for line printer, but its 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 (the printer, system, and network connections).

The Structure of the LP Print Service

This section describes the directory structure, files, logs, and commands of the LP print service.

LP Print Service Directories

The files of the LP print service are distributed among seven directories, as shown in the following table.

TABLE 8-1 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, <code>in.lpd</code> daemon and the <code>printd</code> daemon. The <code>printd</code> daemon transfers all pending jobs in the <code>/var/spool/print</code> directory once per minute. When there are no jobs remaining to transfer, it exits
/usr/sbin	The LP print service administrative commands
/usr/lib/lp	The LP daemons; directories for binary files and PostScript filters; and the <code>model</code> directory (which contains the standard printer interface program)
/var/lp/logs	The logs for LP activities: <code>lpsched.n</code> – Messages from <code>lpsched</code> and <code>requests.n</code> – Information about completed print requests
/var/spool/lp	The spooling directory where files are queued for printing

TABLE 8-1 Directories for the LP Print Service (Continued)

Directory	Contents
<code>/var/spool/print</code>	The LP print service client-side request staging area

LP Print Service Configuration Files

The scheduler stores configuration information in LP configuration files located in the `/etc/lp` directory, as described in the following table.



Caution – The configuration files listed in this table are private interfaces, and 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 8-2 Contents of the `/etc/lp` Directory

File	Type	Description
<code>classes</code>	Directory	Files identifying classes provided by the <code>lpadmin -c</code> command.
<code>fd</code>	Directory	Description of existing filters.
<code>filter.table</code>	File	Print filter lookup table.
<code>forms</code>	Directory	Location to put files for each form. Initially, this directory is empty.
<code>interfaces</code>	Directory	Printer interface program files.
<code>logs</code>	Link to <code>/var/lp/logs</code>	Log files of printing activities.
<code>model</code>	Link to <code>/usr/lib/lp/model</code>	The standard printer interface program.
<code>printers</code>	Directory	Directories for each local printer. Each directory contains configuration information and alert files for an individual printer.
<code>pwheels</code>	Directory	Print wheel or cartridge files.

These configuration files serve a similar function 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 them directly. Instead, use the `lpadmin` command to make configuration changes. Your changes will be 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 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 of the printer-specific directories.

File Name	Description
<code>alert.sh</code>	Shell to execute in response to alerts
<code>alert.vars</code>	Alert variables
<code>configuration</code>	Configuration file
<code>users.deny</code>	List of users to whom printer access is denied
<code>comment</code>	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, which contains definitions for many types of terminals and printers. The LP print service uses information in the `terminfo` database to initialize a printer, to establish a selected page size, character pitch, line pitch, and character set, as well as to communicate 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 44. If necessary, you can add entries to the `terminfo` database, but it is a tedious and time-consuming process. For more information, see “Adding a `terminfo` Entry for an Unsupported Printer” on page 129.

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 8-3 Contents of the `/usr/lib/lp` Directory

File	Type	Description
<code>bin</code>	Directory	Contains files for generating printing alerts, slow filters, and queue management programs.
<code>lpsched</code>	Daemon	Manages scheduling of LP print requests.
<code>model</code>	Directory	Contains the standard printer interface program.
<code>postscript</code>	Directory	Contains all PostScript filter programs provided by the LP print service. These filters come with descriptor files in the <code>/etc/lp/fd</code> directory that 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 described in the following table.

Log File Name	Description
<code>syslogd(1M)</code>	Set <code>lpr.debug</code> in <code>/etc/syslog.conf</code> to enable LP print service logging
<code>/var/spool/lp</code>	A list of current requests that are in the print queue
<code>/var/lp/logs/requests</code>	An ongoing history of print requests

Print Queue Logs

The scheduler for each system keeps a log of print requests in the directories `/var/spool/lp/tmp/system` and `/var/spool/lp/requests/system`. Each print request has two files (one in each directory) that contain information about the request. The information in the `/var/spool/lp/requests/system` directory can be accessed only by root or lp. The information in the `/var/spool/lp/tmp/system` can be accessed only by the user who submitted the request, root, 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 request is finished, the information in the files is combined and appended to the `/var/lp/logs/requests` file, which is described in the section that follows.

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 Logs

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 logs to diagnose and troubleshoot printing problems. This is an example of the contents of the `/var/lp/logs` directory:

```
# cd /var/lp/logs
# ls
lpsched.1      requests      requests.2
lpsched        lpsched.2    requests.1
#
```

The two most important log files for troubleshooting is the `lpsched` log, which contains information about local printing requests and the `requests` log, which contains information about print requests that are completed and no longer in the print queue.

The `requests` log has a simple structure, so that you can extract data using common UNIX shell commands. Requests are listed in the order they are printed, and are 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:

```
# pwd
/var/lp/logs
# tail requests.2
= print1-3, uid 0, gid 1, size 206662, Wed Nov 14 08:56:30 MST 2001
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 LP `requests` log.

TABLE 8-4 Letter Codes in the LP `requests` Log

Letter	Content of Line
=	The separator line. It 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 <code>any</code> .
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: <code>resume</code> , <code>hold</code> , and <code>immediate</code> .
N	(Optional) The type of alert used when the print request was successfully completed. The type is the letter <code>M</code> if the user was notified by email or <code>w</code> if the user was notified by a message to the terminal.

TABLE 8-4 Letter Codes in the LP `requests` Log (Continued)

Letter	Content of Line
O	(Optional) The printer-dependent <code>-o</code> options (for example, <code>nobanner</code>).
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 (the <code>lp -r</code> 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 form. Several bits are used internally by the print service. The bits and what they mean are describe in the table that follows this one.
x	(Optional) The slow filter used for the print request.
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 <code>D</code> 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 and their descriptions.

TABLE 8-5 Outcome Codes in the LP `requests` Log

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.

TABLE 8-5 Outcome Codes in the LP requests Log (Continued)

Outcome Code	Description
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 8-6 Contents of the `/var/spool/lp` Directory

File	Type	Description
SCHEDLOCK	File	Lock file for the scheduler. Check for this file if the scheduler dies and will not restart.
admins	Directory	Link to <code>/etc/lp</code> .
bin	Directory	Link to <code>/usr/lib/lp/bin</code> .
logs	Link	Link to <code>../lp/logs</code> where completed print requests are logged.
model	Link	Link to <code>/usr/lib/lp/model</code> .
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 <code>/var/spool/lp/tmp/hostname</code> , which contains the spooled requests.

TABLE 8-6 Contents of the `/var/spool/lp` Directory (Continued)

File	Type	Description
tmp	Directory	Directory for each configured printer where print requests are logged until printed. Changes to existing print requests are also recorded in this log.

LP Print Service Commands

The following table lists frequently used LP print service commands. You must be root, lp, or assume an equivalent role to use the 1M commands.

TABLE 8-7 Quick Reference to LP Print Service Commands

Command	Man Page	Task
enable	enable(1)	Activate a printer
cancel	cancel(1)	Cancel a print request
lp	lp(1)	Send one or more file(s) to a printer
lpstat	lpstat(1)	Report the status of the LP print service
disable	enable(1)	Deactivate one or more printers
accept	accept(1M)	Permit print requests to be queued for a specific destination
reject	accept(1M)	Prevent print requests from being queued for a specific destination
lpadmin	lpadmin(1M)	Set up or change printer configuration
lpfilter	lpfilter(1M)	Set up or change filter definitions
lpforms	lpforms(1M)	Set up or change preprinted forms
lpadmin	lpadmin(1M)	Mount a form
lpmove	lpmove(1M)	Move output requests from one destination to another
lpsched	lpsched(1M)	Start the LP print service scheduler
lpshut	lpshut(1M)	Stop the LP print service scheduler
lpusers	lpusers(1M)	Set or change the default priority and priority limits that can be requested by users of the LP print service

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 jobs
- Tracks forms mounted on the printer
- Tracks print wheels 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 `lp sched`. The scheduler daemon updates the LP system files with information about printer setup and configuration.

The `lp sched` daemon schedules all local print requests on a print server, as shown in the following figure, whether users 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, if there is one, in the queue on the print server.

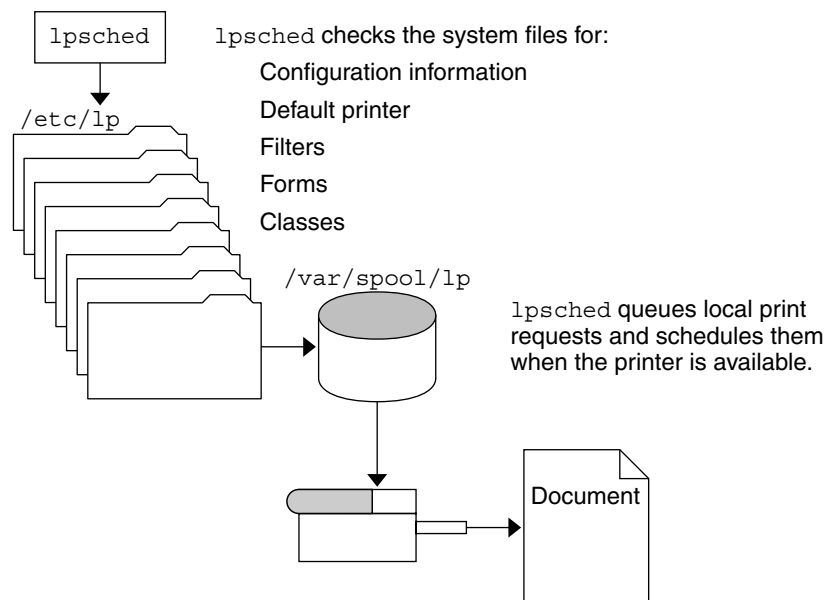


FIGURE 8-1 The `lpsched` Daemon Schedules Local Print Requests

Each print server must have *only* one LP scheduler running. The scheduler is started when a system is booted (or enters run level 2) by the `/etc/rc2.d/S80lp` control script. Without rebooting the systems, you can stop the scheduler with the `/etc/init.d/lp stop` command and restart the scheduler with the `/etc/init.d/lp start` command. The scheduler for each system manages requests issued to the system by the `lp` commands.

Scheduling Network Print Requests

Each print client communicates directly with a print sever over the network. The communication is done between the requesting command (`lp`, `lpstat`, `cancel`, `lpr`, `lpq`, or `lprm`) and the print service on the print server. Doing so, reduces the print system overhead on client only systems, improving scalability, performance and accuracy of data.

Print servers listen for print request 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, returning the results to the requester. It starts on demand and exits when it has serviced the network request. This eliminates idle system overhead for printing. It also eliminates any additional system configuration for network printing support as was the case in previous versions of Solaris printing.

Filtering Print Files

Print filters are programs on the print server that convert the content of a queued file from one format to another.

A print filter can be as simple or as complex as needed. The SunOS release 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 non-PostScript printers, you have to create the filters and add them to the systems that need them.

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), and point to the filter program (for example, `/usr/lib/lp/postscript/postdaisy`).

What the Printer Interface Program Does

The LP print service interacts with other parts of the operating system. It uses a standard printer interface program to:

- 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 one. You can create custom interface programs, but 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 of each print request that it processes and notes any errors that occur during the printing process. This log 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. It 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

Starting with the Solaris 2.6 release, 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.

1. **Become superuser or `lp` on the printer 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 which 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 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 local system does all processing. However, 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.

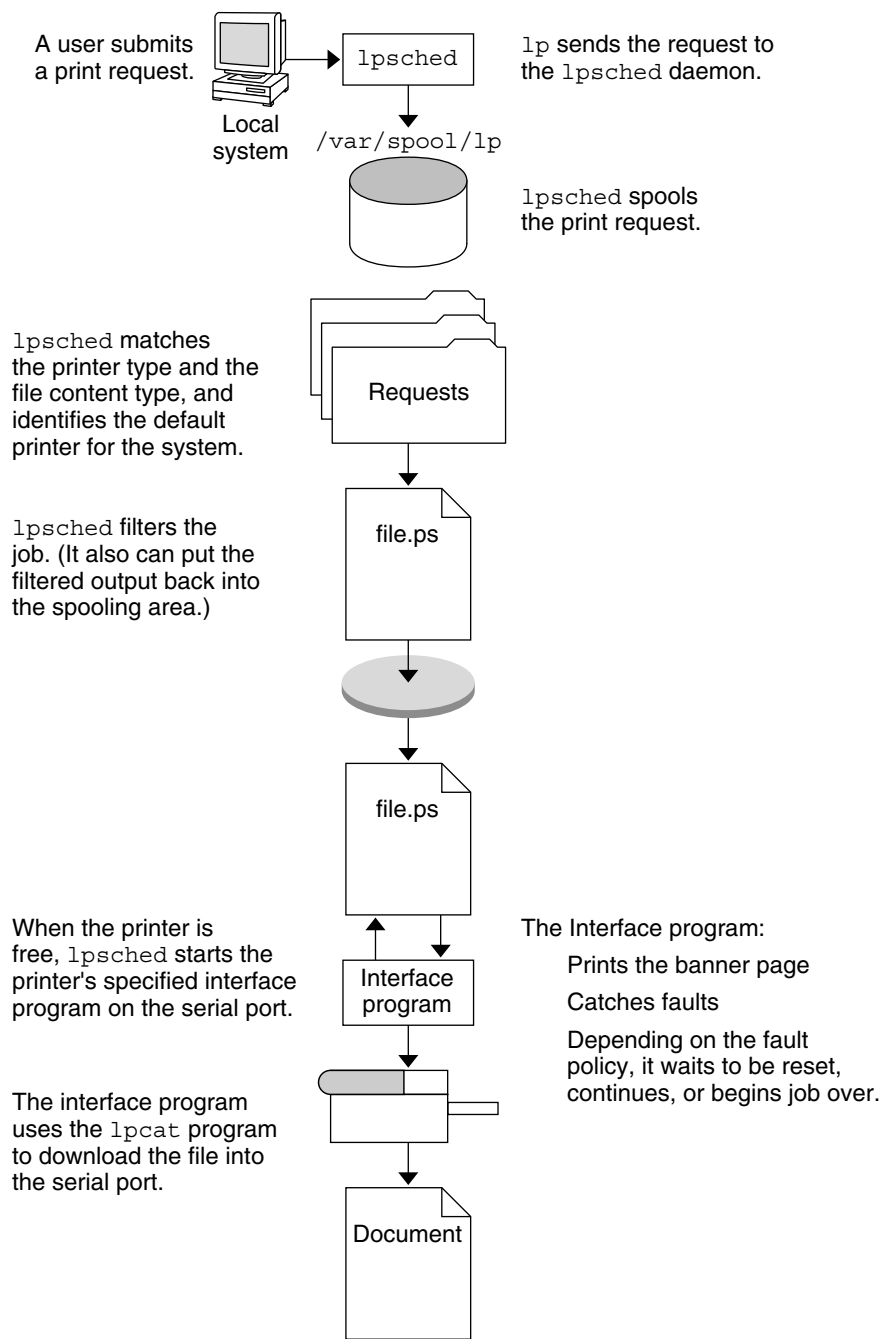


FIGURE 8-2 The Local Printing Process

How Remote Printing Works

The following figure shows what happens when a user on a Solaris print client submits a print request to a lpd-based print server. The command opens a connection and handles its own communications with the print server directly.

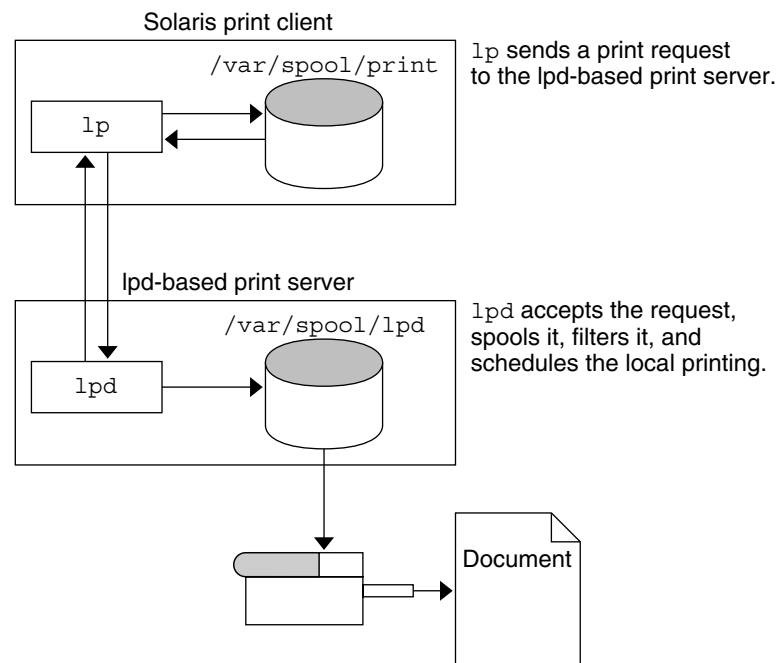


FIGURE 8-3 Printing Between a Solaris Print Client and a lpd-based Print Server

The following figure shows a 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 network listen process, `inetd`, 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.

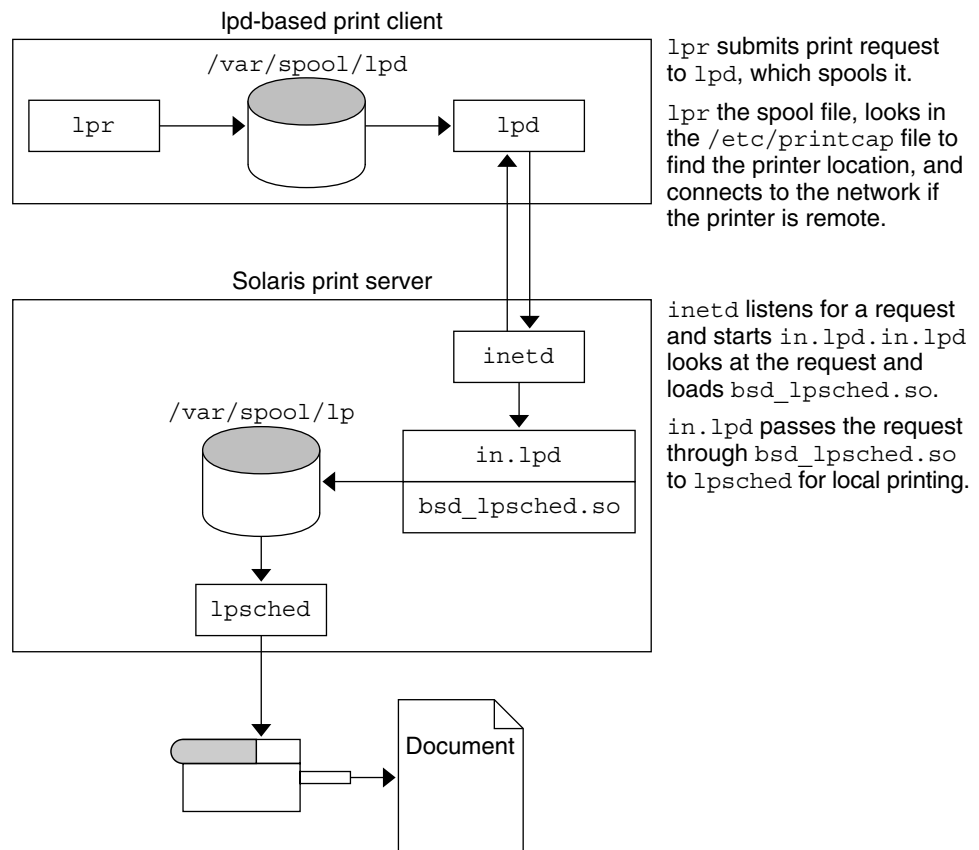


FIGURE 8-4 Printing Between a lpd-based Print Client and a Solaris Print Server

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.

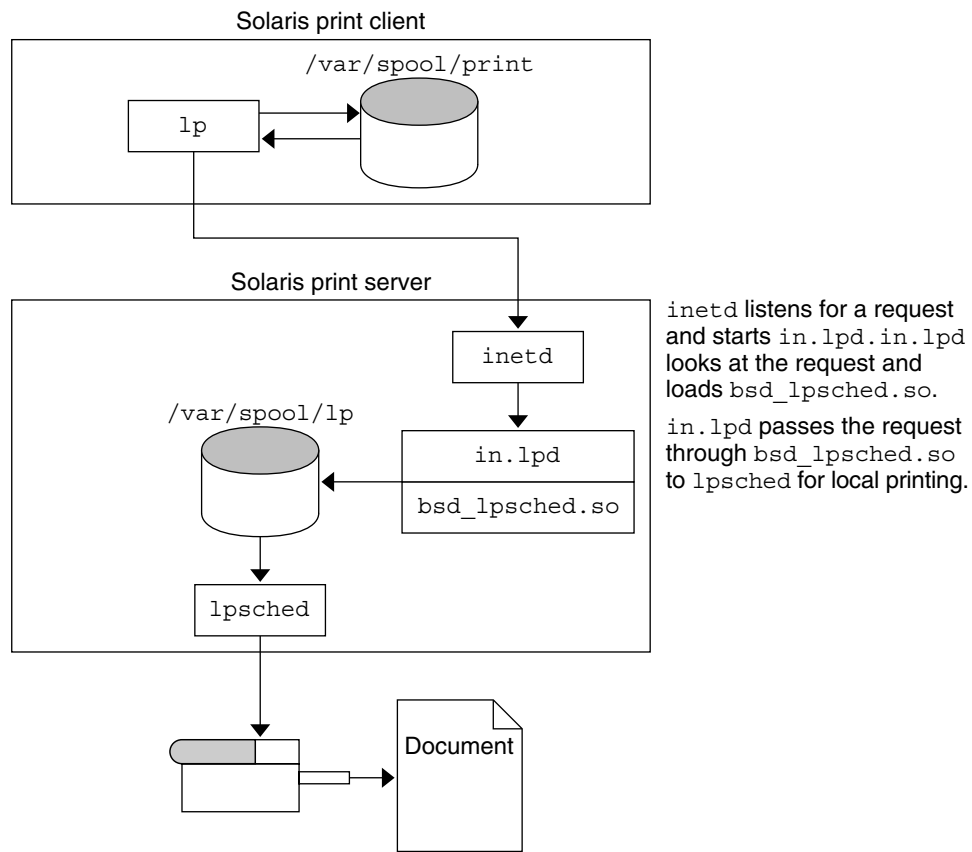


FIGURE 8-5 Printing Between a Solaris Print Client and a Solaris Print Server

Managing Terminals and Modems

Topics

This topic map lists the chapters that provide information for managing terminals and modems.

Chapter 10	Provides overview information about terminals and modems.
Chapter 11	Provides step-by-step instructions for setting up terminals and modems.
Chapter 12	Provides step-by-step instructions for using SAF commands to set up terminals and modems.

Managing Terminals and Modems (Overview)

This chapter provides the overview information for managing terminals and modems. This is a list of the overview information in this chapter.

- “Terminals, Modems, Ports, and Services” on page 172
- “Tools for Managing Terminals and Modems” on page 174
- “Serial Ports Tool” on page 174
- “Service Access Facility” on page 175

For step-by-step instructions about how to set up terminals and modems with Serial Ports Tool, see Chapter 11.

For step-by-step instructions about how to set up terminals and modems with the Service Access Facility (SAF), see Chapter 12.

What’s New in Managing Terminals and Modems?

The Solaris Management Console provides a Serial Ports Tool for setting up terminals and modems. For information on starting the Solaris Management Console, see “Starting the Solaris Management Console” in *System Administration Guide: Basic Administration*. For information on step-by-step instructions for setting up a terminal or modem, see the console online help.

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 environment.

Terminals

Your system's bit-mapped graphics display is not the same as an alphanumeric terminal, which 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, meaning you can access other computers from your own home, but nobody outside can gain access to your machine.

Dial-in service is just the opposite. It allows people to access a system from remote sites, but 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 plugged.

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 includes most modems, alphanumeric terminals, plotters, and some printers) can be plugged 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, it might be necessary 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 via the serial port software. The 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 environment.

TABLE 10-1 Port Monitor Types

Port Monitor	Man Page	Description
listen	listen(1M)	Controls access to network services, such as handling remote print requests prior to the Solaris 2.6 release. The default Solaris operating environment no longer uses this port monitor type.

TABLE 10-1 Port Monitor Types (Continued)

Port Monitor	Man Page	Description
ttymon	ttymon(1M)	Provides access to the login services needed by modems and alphanumeric terminals. 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(1M)`. The new `ttymon` is more powerful; a single `ttymon` can replace multiple occurrences of `getty`. Otherwise, these two programs serve the same function.

Tools for Managing Terminals and Modems

The following table lists the tools for managing terminals and modems.

TABLE 10-2 Tools For Managing Terminals and Modems

Tool Description	Tool	For More Information
The most comprehensive	Service Access Facility (SAF) commands	“Service Access Facility” on page 175
The quickest setup	Solaris Management Console’s Serial Ports Tool	Chapter 11 and Solaris Management Console online help
The quickest setup	Admintool	Admintool online help

Serial Ports Tool

Serial Ports Tool sets up the serial port software to work with terminals and modems by calling the `pmadm` command with the appropriate information. It also provides:

- 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, SAF enables you to set up:

- `ttymon` and `listen` port monitors (using the `sacadm` command)
- `ttymon` port monitor services (using the `pmadm` and `ttyadm` commands)
- `listen` port monitor services (using the `pmadm` and `nlsadmin` commands)
- And troubleshoot `tty` devices
- And troubleshoot incoming network requests for printing service
- And troubleshoot the Service Access Controller (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). SAF is not a program. It is a hierarchy of background processes and administrative commands.

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.

This is a list of the step-by-step instructions in this chapter.

- "How to Set Up a Terminal" on page 180
- "How to Set Up a Modem" on page 181
- "How to Initialize a Port" on page 181

For overview information about terminals and modems, see Chapter 10.

Setting Up Terminals and Modems with Serial Ports Tool

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 the preceding 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 12.

For information on using the Serial Ports command line interface, see `smserialport(1M)`.

Setting Up Terminals

The following table describes the menu items (and their default values) when setting up a terminal using Serial Ports.

TABLE 11-1 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	No
	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 available when setting up a modem using Serial Ports.

TABLE 11-2 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.

TABLE 11-2 Modem Templates *(Continued)*

Modem Configuration	Description
Dial-In and Out (Bidirectional)	Users can either dial in or out from the modem.

The following table describes the default values of each template.

TABLE 11-3 Modem Template Default Values

Detail	Item	Modem - Dial-In Only	Modem - Dial-Out Only	Modem - Dial In and Out
Basic	Port Name	—	—	—
	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	Yes	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 11-4 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:

TABLE 11-4 Initialize Only - No Connection Default Values (Continued)

Detail	Item	Default Value
Advanced	Carrier Detection	Software
	Option: Connect on Carrier	No
	Option: Bidirectional	Yes
	Option: Initialize Only	Yes
	Timeout (seconds)	Never
	Port Monitor	zsmon
	Service Program	/usr/bin/login

▼ How to Set Up a Terminal

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 appears in the Basic Detail mode.
For a description of the Terminal menu items, see Table 11-1.
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 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 tty1
```

▼ How to Set Up a Modem

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 appears in the Basic Detail mode.
For a description of the Modem menu items, see Table 11-3.
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 tty1
```

▼ How to Initialize a Port

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 appears in the Basic Detail mode.

For a description of the Initialize Only menu items, see Table 11-4.

6. Click OK.

7. To configure the advanced items, select the port configured as a 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, it is best to 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 those of the *ttylabel*.
- 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.

```
$ pmadm -l -t ttymon
```

Examine `/etc/ttydefs` and double check the label definition against the terminal configuration. Use `sacadm` 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 `pmadm` reports that the service for the terminal's port is *enabled*, and if the terminal's configuration matches the port monitor's, 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 it with two other parts that are known to be reliable.

Test all of the following:

- Serial ports
 - Modems
 - Cables
 - Connectors
- Do not use 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
```

- If you are connecting a modem to an IA based system, verify the modem is supported by viewing the *Solaris 9 (Intel Platform Edition) Hardware Compatibility List*.

Managing Serial Ports With the Service Access Facility (Tasks)

This chapter describes how to manage serial port services with the Service Access Facility (SAF).

This is a list of the step-by-step instructions in this chapter.

- “How to Add a `ttymon` Port Monitor” on page 193
- “How to View `ttymon` Port Monitor Status” on page 193
- “How to Stop a `ttymon` Port Monitor” on page 194
- “How to Start a `ttymon` Port Monitor” on page 194
- “How to Disable a `ttymon` Port Monitor” on page 195
- “How to Enable a `ttymon` Port Monitor” on page 195
- “How to Remove a `ttymon` Port Monitor” on page 195
- “How to Add a Service” on page 196
- “How to View the Status of a TTY Port Service” on page 197
- “How to Enable a Port Monitor Service” on page 199
- “How to Disable a Port Monitor Service” on page 200

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

- “Using the Service Access Facility” on page 186
- “Overall Administration: `sacadm` Command” on page 187
- “Port Monitor Service Administrator: `pmadm` Command” on page 188
- “Port Monitors: TTY Monitor and Network Listener” on page 190

For reference information about the SAF, see “Service Access Facility Administration (Reference)” on page 200.

Using the Service Access Facility

You can set up terminals and modems with the Solaris Management Console's Serial Ports Tool, `admintool`, or the SAF commands.

The SAF is the tool used for administering terminals, modems, and other network devices. The top-level SAF program is the Service Access Controller (SAC). The SAC controls port monitors which 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 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 and one instance of `listen` can provide multiple services on a network interface.

TABLE 12-1 SAF Control Hierarchy

Function	Program	Description
Overall Administration	<code>sacadm</code>	Command for adding and removing port monitors
Service Access Controller	<code>sac</code>	SAF's master program
Port Monitors	<code>ttymon</code>	Monitors serial port login requests
	<code>listen</code>	Monitors requests for network services
Port Monitor Service Administrator	<code>pmadm</code>	Command for controlling port monitors services
Services	logins; remote procedure calls; other	Services to which SAF provides access
Console Administration	<code>console login</code>	The console is automatically set up via an entry in the <code>/etc/inittab</code> file using <code>ttymon-express</code> mode. Do not use the <code>pmadm</code> or <code>sacadm</code> to manage the console directly. For more information, see "ttymon and the Console Port" on page 191.

Overall Administration: `sacadm` Command

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 SAC upon entering multiuser mode.

When SAC is invoked, it first looks for, and interprets, each system's configuration script, by which SAC customizes its environment. 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 it has interpreted the per-system configuration script, the SAC program reads its administrative file and starts the specified port monitors. For each port monitor, SAC runs a copy of itself (SAC forks a child process). Each child 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 using the command found in the SAC 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/_safconfig`, 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 Administrator: pmadm Command

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 via the `pmtag` argument, and a particular port via 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 `ttypadm` and `nlsadmin`, respectively.

A Port Monitor at Work: `ttymon`

Whenever you attempt to log in by using a directly connected modem or alphanumeric terminal, `ttymon` goes to work, as follows.

As shown in the following figure, 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.

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.

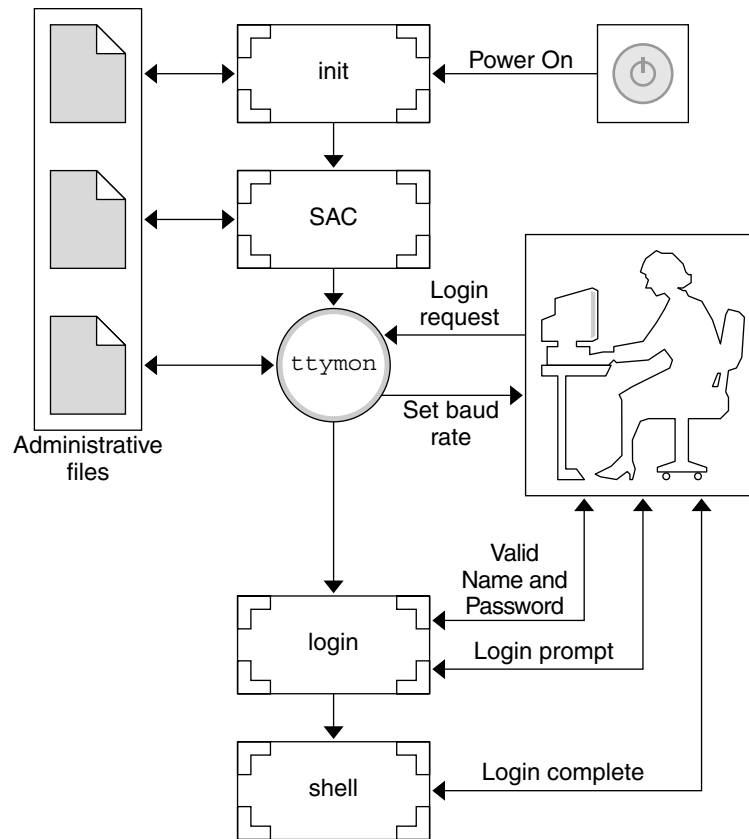


FIGURE 12-1 How `ttymon` Helps Process a Login Request

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 what data transfer rate, line discipline, and handshaking protocol are required to communicate with the device.

Having established the proper parameters for communication with the modem or terminal, 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 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 interprets the per-service configuration file for the port, creates an `/etc/utmpx` entry if required, establishes the service environment, and invokes the service associated with the port.

After the service terminates, the `ttymon` port monitor cleans up the `/etc/utmpx` entry, if one exists, and returns the port to its initial state.

Bidirectional Service

If a port is configured for bidirectional service, the `ttymon` port monitor will:

- Allow users to connect to a service
- Allow the `uucico`, `cu`, or `ct` commands to use the port for dialing out (if the port's free)
- Wait to read a character before printing a prompt
- Invoke the port's associated service, without sending the prompt message, when a connection is requested (if the connect-on-carrier flag is set)

Port Monitors: TTY Monitor and Network Listener

Though SAF provides a generic means for administering any future or third-party port monitors, only two are implemented in the Solaris environment: `ttymon` and `listen`.

TTY Port Monitor: `ttymon`

The `ttymon` port monitor is STREAMS-based. It monitors ports; sets terminal modes, baud rates, and line disciplines; and invokes the login process. (It provides Solaris users the same services that `getty` did under previous versions of SunOS 4.1 software.)

The `ttymon` port monitor runs under the SAC program. It is configured using 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 using the `pmadm` and `ttysadm` commands.

`ttymon` and the Console Port

Console services are not managed by the Service Access Controller nor any explicit `ttymon` administration file. An entry in the `/etc/inittab` file is used to manage the console port using `ttymon` in *express* mode, which 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 looks like this:

```
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`, which means 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 the 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 hangup by setting the line speed to zero before setting the default or specified speed.

`-p "`uname -n` console login:`

Identifies the prompt string for the console port.

`-t terminal_type`

Identifies 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.

Special `ttymon`-Specific Administrative Command: `ttyadm`

The `ttymon` administrative file is updated by `sacadm` and `pmadm`, as well as by the `ttyadm` command. The `ttyadm` command formats `ttymon`-specific information and writes it to the standard output, providing a means for presenting formatted `ttymon`-specific data to the `sacadm` and `pmadm` commands.

Thus, `ttyadm` does not administer `ttymon` directly; rather, it complements the generic administrative commands, `sacadm` and `pmadm`. For more information, see `ttyadm(1M)`.

Network Listener Service: `listen`

The `listen` port monitor runs under SAC. It monitors the network for service requests, accepts requests when they arrive, and invokes servers in response to those service requests.

The `listen` port monitor is configured 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 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 environment, `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 `sacadm` and `pmadm`, as well as by the `nlsadmin` command. The `nlsadmin` command formats `listen`-specific information and writes it to the standard output, providing a means of presenting formatted `listen`-specific data to the `sacadm` and `pmadm` commands.

Thus, `nlsadmin` does not administer `listen` directly. Rather, it complements the generic administrative commands, `sacadm` and `pmadm`.

Each network can have at least one instance of the network listener process associated with it. Each network is configured separately. 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 `nlsadmin(1M)`.

Administering `ttymon` Port Monitors

Use the `sacadm` command to add, list, remove, kill, start, enable, disable, enable, and remove a `ttymon` port monitor.

Note – You must be superuser to perform the following procedures.

▼ How to Add a `ttymon` Port Monitor

To add a `ttymon` port monitor, type the following:

```
# sacadm -a -p mbmon -t ttymon -c /usr/lib/saf/ttymon -v `ttyadm
-V` -y "TTY Ports a & b"
```

- a The *add* port monitor flag.
- 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

To see the status of a `ttymon` port monitor, type the following:

```
# sacadm -l -p mbmon
```

- l The *list* port monitor status flag.
- p Specifies the *pmtag* `mbmon` as the port monitor tag.

Example—Viewing `ttymon` Port Monitor Status

```
# 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.
mbmon

PMTYPE               Identifies the port monitor type, ttymon.
ttymon

FLGS                 Indicates whether the following two flags are set:
-                    d, do not enable the new port monitor, or
                     x, do not start the new port monitor. There are no flags set in
                     this example.

RCNT                 Indicates the return count value. A return count of 0 indicates
0                    that the port monitor is not to be restarted if it fails.

STATUS               Indicates the current status of the port monitor.
STARTING

COMMAND              Identifies the command used to start the port monitor.
/usr/lib/saf ...

#TTY Ports a & b     Identifies any comment used to describe the port monitor.

```

▼ How to Stop a `ttymon` Port Monitor

To kill a `ttymon` port monitor, type the following:

```
# sacadm -k -p mbmon
```

`-k` The *kill* port monitor status flag.

`-p` Specifies the *pmtag* `mbmon` as the port monitor tag.

▼ How to Start a `ttymon` Port Monitor

To start a killed `ttymon` port monitor, type the following:

```
# sacadm -s -p mbmon
```

`-s` 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.

To disable a ttymon port monitor, type the following:

```
# sacadm -d -p mbmon
```

-d 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.

To enable a ttymon port monitor, type the following:

```
# sacadm -e -p mbmon
```

-e The *enable* port monitor status flag.

-p Specifies the *pmtag* mbmon as the port monitor tag.

▼ How to Remove a ttymon Port Monitor

To remove a ttymon port monitor, type the following:

```
# sacadm -r -p mbmon
```

-r The *remove* port monitor status flag.

-p Specifies the *pmtag* mbmon as the port monitor tag.

Note – Removing a port monitor deletes all the configuration files associated with it. Port monitor configuration files cannot be updated or changed using `sacadm`. To reconfigure a port monitor, *remove* it and *add* a new one.

Administering `ttymon` Services

Use `pmadm` to add services, list the services of one or more ports associated with a port monitor, and enable or disable a service.

Note – You must be superuser to perform the following procedures.

▼ How to Add a Service

To add a standard terminal service to the `mbmon` port monitor, type the following:

```
# 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 to the next line. Do not put a Return or line feed after `contty`.

- a 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 it runs.
- v Specifies the *version* number of the port monitor.
- m Specifies the `ttymon`-specific configuration data formatted by `ttyadm`.

The above `pmadm` command contains an embedded `ttyadm` command. The options in this embedded command are as follows:

- b 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 to list the status of a TTY port, or all the ports associated with a port monitor.

To list one service of a port monitor, type the following:

```
# pmadm -l -p mbmon -s a
```

- l Lists service information.
- p Specifies the *pmtag* `mbmon` as the port monitor tag.
- s Specifies the *svctag* `a` as the port monitor *service* tag.

To list all services of all port monitors, type the following:

```
# pmadm -l
```

- l Lists service information

To list all services of a port monitor, type the following:

```
# pmadm -l -p mbmon
```

- l Lists service information
- p Specifies the *pmtag* `mbmon` as the port monitor tag

Example—Viewing the Status of a TTY Port Monitor Service

```
# 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 - y #
```

mbmon	Identifies the port monitor name, mbmon, set by using the pmadm -p command.
ttymon	Identifies the port monitor type, ttymon.
a	Indicates the service tag value set by using the pmadm -s command.
-	Identifies whether the following flags are set by using the pmadm -f command: x, which means do not enable the service; u, which means create a utmpx entry for the service. No flags are set in this example.
root	Identifies the ID assigned to the service when its started. This value is set by using the pmadm -i command.
<PMSPECIFIC>	<i>Information</i>
/dev/term/a	Indicates the TTY port pathname 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.
-	Indicates a value set by using the ttyadm -r option. This option determines when ttymon displays a prompt after receiving data from a port. If count is 0, ttymon will wait until it receives any character. If count is greater than 0, ttymon will wait until count new lines have been received. No value is set in this example.

<code>/usr/bin/login</code>	Identifies the full pathname of the service to be invoked when a connected is received. This value is set by using <code>ttyadm -s</code> command.
<code>-</code>	Identifies the <code>ttyadm -t</code> command's (timeout) value. This option specifies that <code>ttymon</code> should close a port if the open on the port succeeds, and no input data is received in timeout seconds. There is no timeout value in this example.
<code>contty</code>	Identifies the TTY label in the <code>/etc/ttydefs</code> file. This value is set by using the <code>ttyadm -l</code> command.
<code>ldterm, ttcompat</code>	Identifies the STREAMS modules to be pushed. These modules are set by using the <code>ttyadmin -m</code> command.
<code>login: Terminal disabled</code>	Identifies an inactive message to be displayed when the port is disabled. This message is set by using the <code>ttyadm -i</code> command.
<code>tvi925</code>	Identifies the terminal type, if set, by using the <code>ttyadm -T</code> command. The terminal type is <code>tvi925</code> in this example.
<code>y</code>	Identifies the software carrier value set by using the <code>ttyadm -S</code> command; <code>n</code> will turn software carrier off, <code>y</code> will turn software carrier on. Software carrier is turned on in this example.
<code>#</code>	Identifies any comment specified with the <code>pmadm -y</code> command. (There is no comment in this example).

▼ How to Enable a Port Monitor Service

To enable a disabled port monitor service, type the following:

```
# pmadm -e -p mbmon -s a
```

- `-e` 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

To disable a port monitor service, type the following:

```
# pmadm -d -p mbmon -s a
```

- d 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 SAF

SAF uses configuration files which can be modified by using the `sacadm` and `pmadm` commands. You should not need to edit them manually.

File Name	Description
<code>/etc/saf/_sysconfig</code>	Per-system configuration script.
<code>/etc/saf/_sactab</code>	SAC's administrative file that contains configuration data for the port monitors that the SAC controls.
<code>/etc/saf/pmtag</code>	Home directory for port monitor <i>pmtag</i> .
<code>/etc/saf/pmtag/_config</code>	Per-port monitor configuration script for port monitor <i>pmtag</i> if it exists.
<code>/etc/saf/pmtag/_pmtab</code>	Port monitor <i>pmtag</i> 's administrative file that contains port monitor-specific configuration data for the services <i>pmtag</i> provides.
<code>/etc/saf/pmtag/svctag</code>	Per-service configuration script for service <i>svctag</i> .
<code>/var/saf/log</code>	SAC's log file.

File Name	Description
<code>/var/saf/pmtag</code>	Directory for files created by <i>pmtag</i> , for example, log files.

The `/etc/saf/_sactab` File

The `/etc/saf/_sactab` looks like this:

```
# VERSION=1
zsmon:ttymon::0:/usr/lib/saf/ttymon      #
```

<code># VERSION=1</code>	Indicates the Service Access Facility version number.
<code>zsmon</code>	Is the name of the port monitor.
<code>ttymon</code>	Is the type of port monitor.
<code>::</code>	Indicates whether the following two flags are set: <code>d</code> , do not enable the port monitor <code>x</code> , do not start the port monitor. No flags are set in this example.
<code>0</code>	Indicates the return code value. A return count of 0 indicates that the port monitor is not be restarted if it fails.
<code>/usr/lib/saf/ttymon</code>	Indicates the port monitor pathname.

The `/etc/saf/pmtab/_pmtab` File

The `/etc/saf/pmtab/_pmtab` file, such as `/etc/saf/zsmon/_pmtab`, looks 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:#
```

<code># VERSION=1</code>	Indicates the Service Access Facility version number.
<code>ttya</code>	Indicates the service tag.

<code>x, u</code>	Identifies whether the following flags are set: <code>x</code> , which means do not enable the service. <code>u</code> , which means create a <code>utmpx</code> entry for the service.
<code>root</code>	Indicates the identity assigned to the service tag.
<code>reserved</code>	This field is reserved.
<code>reserved</code>	This field is reserved.
<code>reserved</code>	This field is reserved.
<code>/dev/term/a</code>	Indicates the TTY port pathname.
<code>/usr/bin/login</code>	Identifies the full pathname of the service to be invoked when a connection is received.
<code>:c, b, h, I, r:</code>	Indicates whether the following flags are set: <code>c</code> , sets the connect on carrier flag for the port <code>b</code> , sets the port as bidirectional, allowing both incoming and outgoing traffic <code>h</code> , suppresses an automatic hangup immediately after an incoming call is received <code>I</code> , initializes the port <code>r</code> , forces <code>ttymon</code> to wait until it receives a character from the port before it prints the <code>login:</code> message.
<code>9600</code>	Identifies the TTY label defined in <code>/etc/ttydefs</code> file.
<code>ldterm, ttcompat</code>	Identifies the STREAMS modules to be pushed.
<code>ttya login\:</code>	Identifies the prompt to be displayed.
<code>:y/n:</code>	
<code>message</code>	Identifies any inactive (disabled) response message.
<code>tvi925</code>	Identifies the terminal type.
<code>y</code>	Indicates whether software carrier is set (<code>y/n</code>).

Service States

The `sacadm` command controls the states of services. The following table describes the possible states of services.

State	Description
Enabled	<i>Default state</i> – When the port monitor is added, the service operates.
Disabled	<i>Default state</i> – 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 `ttymon` and `listen` port monitors. The following table describes the possible port monitor states.

State	Description
Started	<i>Default state</i> – When the port monitor is added, it is automatically started.
Enabled	<i>Default state</i> – When the port monitor is added, it is automatically ready to accept requests for service.
Stopped	<i>Default state</i> – When the port monitor is removed, it is automatically stopped.
Disabled	<i>Default state</i> – When the port monitor is removed, it automatically continues existing services and refuses to add new services.
Starting	<i>Intermediate state</i> – The port monitor is in the process of starting.
Stopping	<i>Intermediate state</i> – The port monitor has been manually terminated, but it has not completed its shutdown procedure. It is on the way to becoming stopped.
Notrunning	<i>Inactive state</i> – 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	<i>Inactive state</i> – The port monitor is unable to start and remain running.

To determine the state of any particular port monitor, use the following:

```
# sacadm -l -p portmon_name
```

Port States

Ports can be enabled or disabled depending on the state of the port monitor that controls them.

State	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 will send out the disabled message when it receives a connection request.

Managing System Resources Topics

This topic map lists the chapters that provide information for managing system resources.

Chapter 14	Provides an overview of Solaris commands and utilities that help you manage system resources by using disk quotas, the <code>cron</code> and <code>at</code> commands, and accounting programs.
Chapter 15	Provides step-by-step instructions for displaying and changing system information.
Chapter 16	Provides step-by-step instructions for optimizing disk space by locating unused files and large directories.
Chapter 17	Provides step-by-step instructions for setting up and administering disk quotas.
Chapter 18	Provides step-by-step instructions for scheduling routine or one-time system events by using the <code>crontab</code> and <code>at</code> commands.
Chapter 19	Provides step-by-step instructions for setting up and maintaining system accounting.
Chapter 20	Provides reference information for system accounting software, including how the different accounting features operate and how to use the various accounting programs.

Managing System Resources (Overview)

This chapter provides a brief description of new Solaris 9 features and a map of features to help you manage system resources.

Using these features, you can display general system information, monitor disk space, set disk quotas, use accounting programs, and schedule the `crontab` 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 associated with managing system resources without Solaris resource management, see “High-Level View of Managing System Resources (Task Map)” on page 209.

For information on managing system resources with Solaris resource management, see “Introduction to Solaris 9 Resource Manager” in *System Administration Guide: Resource Management and Network Services*.

What’s New in Managing System Resources?

New `df`, `du`, and `ls` Options

The `df`, `du`, and `ls` commands have a new `-h` option to display disk usage and file or file system sizes in powers of 1024. This option makes it easier to interpret the output of the `df`, `du`, and `ls -l` commands by providing disk space in kilobytes, Mbytes, Gbytes, or terabytes if the file or directory size is larger than 1024 bytes.

For more information, see `df(1B)`, `du(1)`, and `ls(1)`.

For more information in this guide, see “How to Display Information About Files and Disk Space” on page 220.

Changing System Parameters

In previous versions of this book, there was a section on changing system parameters such as increasing the number of processes per user and shared memory segments. This section has been removed. For information on changing system tunable parameters, see *Solaris Tunable Parameters Reference Manual*.

In addition, information on setting up a network time server has been moved to “Time Related Services” in *System Administration Guide: Resource Management and Network Services*.

Solaris Resource Management

Solaris resource management enables you to control how applications use available system resources.

You can also set up the extended accounting facility to monitor and record resource consumption on a task or process basis.

Extended accounting introduces a new, variable length general-purpose accounting file format that represents general groups of accounting data. It provides the ability to configure resource utilization recorded by the kernel in the various accounting files.

Solaris resource management features include:

- *Tasks* – New process collectives for tracking resource usage.
- *Projects* – New administrative databases for charging resource usage. Resource usage can be charged to a project based on tasks performed.
- `acctadm` – A new tool for configuring various attributes of the extended accounting feature. For example, the resources tracked by the system accounting can be configured on a system-wide basis.

The new default accounting configuration requires no administration. However, do not remove the `/etc/project` file, even if you do not use the extended accounting features. Users cannot log in to the system unless this file exists.

Use the following table to find more information about the extended accounting features in this release.

Solaris Resource Management Feature	For More Information
Turning on and off extended accounting	“Extended Accounting” in <i>System Administration Guide: Resource Management and Network Services</i> and <code>acctadm(1M)</code>
Description of the <code>projects</code> database	“Projects and Tasks” in <i>System Administration Guide: Resource Management and Network Services</i> and <code>project(4)</code>
Directly collecting extended accounting data	<code>libexacct(3LIB)</code> , <code>getacct(2)</code> , <code>putacct(2)</code> , and <code>wracct(2)</code>

High-Level View of Managing System Resources (Task Map)

Task	Description	Instructions
Displaying and changing system information	Use various commands to display and changing system information such as general system information, the date and time, and changing the system’s host name.	Chapter 15
Managing disk use	Identify how disk space is used and take steps to remove old and unused files.	Chapter 16
Managing quotas	Use UFS file system quotas to manage how much disk space is used by users.	Chapter 17
Scheduling system events	Use <code>cron</code> and <code>at</code> jobs to help schedule system routines that can include clean up of old and unused files.	Chapter 18
Managing system accounting	Use system accounting to identify how users and applications are using system resources.	Chapter 19

Displaying and Changing System Information (Tasks)

This chapter describes the tasks required to display and change the most common system information. This is a list of the step-by-step instructions in this chapter.

- “How to Determine If a System Can Run the 64-bit Solaris Operating Environment” on page 212
- “How to Display General System Information” on page 214
- “How to Display a System’s Host ID Number ” on page 214
- “How to Display the Date and Time” on page 215
- “How to Set a System’s Date and Time Manually” on page 216
- “How to Set Up a Message-of-the-Day” on page 216
- “How to Change a System’s Host Name” on page 217

Displaying System Information

The following table describes commands that enable you to display general system information.

TABLE 15-1 Commands for Displaying System Information

Command	Man Page	System Information Displayed
<code>psrinfo</code>	<code>psrinfo(1M)</code>	Processor type
<code>isainfo</code>	<code>isainfo(1)</code>	Supported applications and it reports the number of bits supported by <i>native</i> applications on the running system, which can be passed as a token to scripts

TABLE 15-1 Commands for Displaying System Information (Continued)

Command	Man Page	System Information Displayed
showrev	showrev(1M)	Hostname, host ID, release, kernel architecture, application architecture, hardware provider, domain, and kernel version
uname	uname(1)	Operating system name, release, and version; node name; hardware name; processor type
hostid	hostid(1)	Host ID number
prtconf	prtconf(1M)	Installed memory
date	date(1)	Date and time

▼ How to Determine If a System Can Run the 64-bit Solaris Operating Environment

Currently, the only platform capable of supporting the 64-bit Solaris operating environment 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.

If you are running the Solaris 9 release, you can verify that the machine is an UltraSPARC system by using the `psrinfo` command:

```
# psrinfo -v  
Status of processor 0 as of: 05/30/01 13:48:46  
  Processor has been on-line since 05/18/01 10:02:40.  
  The sparcv9 processor operates at 333 MHz,  
    and has a sparcv9 floating point processor.
```

If the processor type is `sparcv9`, the platform is capable of running the 64-bit Solaris operating environment. This test does not work on previous versions of the `psrinfo` command, where all platforms report the less precise `sparc` as the processor type.

▼ How to Determine If a System Has 64-bit Solaris Capabilities Enabled

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.

Examples—Determining If a System Has 64-bit Solaris Capabilities Enabled

The output for an UltraSPARC system running a 32-bit kernel appears as follows:

```
$ isainfo -v
32-bit sparc applications
```

This output means that this system is capable of supporting only 32-bit applications.

The output for an UltraSPARC system 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, IA, or UltraSPARC system running the 32-bit Solaris operating environment appears as follows:

```
$ isainfo -b
32
```

The output from a 64-bit UltraSPARC system running the 64-bit Solaris operating environment appears as follows:

```
$ isainfo -b
64
```

The command returns 64 only. Even though a 64-bit UltraSPARC system is capable of running both types of applications, 64-bit applications are the best kind of applications to run on a 64-bit system.

▼ How to Display System and Software Release Information

To display system information and software release information, use the `showrev` command.

```
$ showrev [-a]
```

The `-a` option displays all available system information.

Example—Displaying System and Software Release Information

The following example shows the `showrev` command output.

```
$ showrev -a
Hostname: starbug
Hostid: nnnnnnnn
Release: 5.9
Kernel architecture: sun4u
Application architecture: sparc
Hardware provider: Sun_Microsystems
Domain: solar.com
Kernel version: SunOS 5.9 May 2002

OpenWindows version:
X11 Version 6.4.2 11 April 2001

No patches are installed
$
```

▼ How to Display General System Information

To display system information, use the `uname` command.

```
$ uname [-a]
```

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.

Example—Displaying General System Information

The following example shows the `uname` command output.

```
$ uname
SunOS
$ uname -a
SunOS starbug 5.9 Generic sun4u sparc SUNW,Ultra-5_10
$
```

▼ How to Display a System's Host ID Number

To display the host ID number in hexadecimal format, use the `hostid` command.

```
$ hostid
```

Example—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

To display the amount of memory installed on your system, use the `prtconf` command.

```
$ prtconf [| grep Memory]
```

The `grep Memory` command selects output from the `prtconf` command to display memory information only.

Example—Displaying a System’s Installed Memory

The following example shows sample output from the `prtconf` command.

```
# prtconf | grep Memory
Memory size: 128 Megabytes
```

▼ How to Display the Date and Time

To display the current date and time according to your system clock, use the `date` command.

```
$ date
```

Example—Displaying the Date and Time

The following example shows sample output from the `date` command.

```
$ date
Thu May 31 17:44:58 MDT 2001
$
```

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

1. **Become superuser.**
2. **Enter the new date and time.**

```
# date mddHHMM[[cc]yy]
```

<i>mm</i>	Month, using two digits.
<i>dd</i>	Day of the month, using two digits.
<i>HH</i>	Hour, using two digits and a 24-hour clock.
<i>MM</i>	Minutes, using two digits.
<i>cc</i>	Century, using two digits.
<i>yy</i>	Year, using two digits.

See `date(1)` for more information.

3. **Verify that you have reset your system's date correctly by using the `date` command with no options.**

Example—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
Thu Jun 21 13:59:15 MDT 2001
# date 0621141001
Thu Jun 21 14:10:00 MDT 2001
```

▼ 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.

1. **Become superuser.**
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 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—Setting Up a Message-of-the-Day

The default message-of-the-day, provided when you install Solaris software, contains SunOS version information:

```
$ cat /etc/motd
Sun Microsystems Inc.   SunOS 5.9           Generic   May 2002
```

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 that you will need to update your name service database to reflect the new host name.

1. **Become superuser.**
2. **Change the system's host name in the following files:**
 - `/etc/nodename`
 - `/etc/hostname.xxy`
 - `/etc/inet/hosts`
 - `/etc/net/ticlts/hosts`
 - `/etc/net/ticots/hosts`
 - `/etc/net/ticotsord/hosts`
3. **(Optional) If using a name service, change the system's host name in the `hostfile`.**
4. **Reboot the system to activate the new host name.**

```
# init 6
```

Managing Disk Use (Tasks)

This chapter describes how to optimize disk space by locating unused files and large directories.

This is a list of the step-by-step instructions in this chapter.

- “How to Display Information About Files and Disk Space” on page 220
- “How to Display the Size of Files” on page 222
- “How to Find Large Files” on page 223
- “How to Find Files That Exceed a Specified Size Limit” on page 223
- “How to Display the Size of Directories, Subdirectories, and Files” on page 224
- “How to Display the User Ownership of Local UFS File Systems” on page 226
- “How to List the Newest Files” on page 227
- “How to Find and Remove Old or Inactive Files ” on page 228
- “How to Clear Out Temporary Directories” on page 229
- “How to Find and Delete `core` Files” on page 230
- “How to Delete Crash Dump Files” on page 230

Displaying Information About Files and Disk Space

This table summarizes the commands available for displaying information about file size and disk space.

Command	Man Page	Description
<code>df</code>	<code>df(1M)</code>	Reports the number of free disk blocks and files.

Command	Man Page	Description
<code>du</code>	<code>du(1)</code>	Summarizes disk space allocated to each subdirectory.
<code>find -size</code>	<code>find(1)</code>	Searches recursively through a directory based on the size specified with the <code>-size</code> option.
<code>ls -lh</code>	<code>ls(1)</code>	Lists the size of a file in the power of 1024.

▼ How to Display Information About Files and Disk Space

Display information about how disk space is used by using the `df` command.

```
$ df [directory] [-h] [-t] [-t]
```

<code>df</code>	With no options, lists all mounted file systems and their device names, the number of 512-byte blocks used, and the number of files.
<i>directory</i>	Specifies the directory whose file system you want to check.
<code>-h</code>	Displays disk space in the power of 1024.
<code>-t</code>	Displays the total blocks as well as the blocks used for all mounted file systems.

Examples—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
/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
```

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
/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
```

Note – 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 a MNTFS file system.

The following example shows a list of all mounted file systems, device names, total 512-byte blocks used, and number of files. The second line of each two-line entry displays the total number of blocks and files 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
/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

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 `ls(1)` and `find(1)`.

▼ How to Display the Size of Files

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 and directory sizes into Kbytes, Mbytes, Gbytes, or Terabytes when the file or directory size is larger than 1024 bytes. This option also modifies the output displayed by the <code>-o</code> , <code>-n</code> , <code>-@</code> , and <code>-g</code> options to display file or directory sizes in the new format. For more information, see <code>ls(1)</code> .
-s	Displays a list of the files and directories, showing the sizes in blocks.

Examples—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 exactt/
-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

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
```

```
sort -nr
```

Sorts the list of files by block size from largest to smallest.

Example—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

To locate and display the names of files that exceed a specified size, use the `find` command.

```
$ find directory -size +nnn
```

<i>directory</i>	Identifies the directory you want to search.
<i>-size +<i>nnn</i></i>	Is a number of 512-byte blocks. Files that exceed this size are listed.

Example—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.

```
$ 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 `du(1)` and `quot(1M)`.

▼ How to Display the Size of Directories, Subdirectories, and Files

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 ...]
```


<code>du</code>	Displays the size of each directory you specify, including each subdirectory beneath it.
<code>-a</code>	Displays the size of each file and subdirectory, and the total number of blocks contained in the specified directory.
<code>-s</code>	Displays the total number of blocks contained in the specified directory.
<code>-h</code>	Displays the size of each directory in 1024 bytes.
<code>-H</code>	Displays the size of each directory in 1000 bytes.
<i>directory ...</i>	Identifies one or more directories you want to check.

Examples—Displaying the Size of Directories, Subdirectories, and Files

The following example shows the total 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, all of the subdirectories and files they contain, and the total number of blocks contained in each directory.

```
$ du /var/adm /var/spool/lp
2      /var/adm/log
2      /var/adm/passwd
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
130    /var/adm
4      /var/spool/lp/admins
2      /var/spool/lp/fifos/private
2      /var/spool/lp/fifos/public
6      /var/spool/lp/fifos
2      /var/spool/lp/requests/starbug
4      /var/spool/lp/requests
2      /var/spool/lp/system
2      /var/spool/lp/tmp/starbug
2      /var/spool/lp/tmp/.net/tmp/starbug
4      /var/spool/lp/tmp/.net/tmp
2      /var/spool/lp/tmp/.net/requests/starbug
4      /var/spool/lp/tmp/.net/requests
10     /var/spool/lp/tmp/.net
14     /var/spool/lp/tmp
40     /var/spool/lp
```

The following example shows directory sizes in 1024 bytes.

```
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

1. Become superuser.
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.

Note – The `quot` command works only on local UFS file systems.

Example—Displaying the User Ownership of Local UFS File Systems

In the following example, users of the root (/) file system are displayed. Then, users of all mounted UFS file systems are displayed.

```
# quot /
/dev/rds/c0t0d0s0:
43340  root
 3142  rimmer
   47  uucp
   35  lp
   30  adm
    4  bin
    4  daemon
# quot -a
/dev/rds/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 using the `ls` or `find` commands. For more information, see `ls(1)` and `find(1)`.

Other ways to conserve disk space include emptying temporary directories such as the ones located in `/var/tmp` or `/var/spool`, and deleting `core` and crash dump files. For more information about crash dump files, refer to Chapter 28.

▼ How to List the Newest Files

List files, displaying the most recently created or changed files first, by using the `ls -t` command.

```
$ ls -t [directory]
```

<code>-t</code>	Sorts files by latest time stamp first.
<code>directory</code>	Identifies the directory you want to search.

Example—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

1. Become superuser.
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
```

<i>directory</i>	Identifies the directory you want to search. Directories below this directory are also searched.
<code>-atime +nnn</code>	Finds files that have not been accessed within the number of days (nnn) you specify.
<code>-mtime +nnn</code>	Finds files that have not been modified within the number of days (nnn) you specify.
<i>filename</i>	Identifies the file that contains the list of inactive files.

3. Remove the inactive files that you listed in the previous step.

```
# rm `cat filename`
```

filename identifies the file created in the previous step which contains the list of inactive files.

Example—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

1. Become superuser.
2. Change to the directory to clean out.

```
# cd directory
```



Caution – Be sure you are in the right 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 containing unnecessary temporary or obsolete subdirectories and files, and delete them by repeating Step 3.

### Example—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

1. Become superuser.
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—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, so if you have enabled your system to store these files, do not retain them for longer than necessary.

1. Become superuser.
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** – Be sure you are in the right directory before completing step 3. Step 3 deletes all files in the current directory.

---

3. Remove the crash dump files.
4. Verify that the crash dump files are removed.

```
rm *
```

```
ls
```

### Example—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
```





---

## Managing Quotas (Tasks)

---

This chapter describes how to set up and administer quotas for disk space and inodes. This is a list of the step-by-step instructions in this chapter.

- “How to Configure File Systems for Quotas” on page 237
- “How to Set Up Quotas for a User” on page 238
- “How to Set Up Quotas for Multiple Users” on page 239
- “How to Check Quota Consistency” on page 240
- “How to Turn On Quotas” on page 240
- “How to Check for Exceeded Quotas” on page 241
- “How to Check Quotas on a File System” on page 242
- “How to Change the Soft Time Limit Default” on page 244
- “How to Change Quotas for a User” on page 245
- “How to Disable Quotas for a User” on page 246
- “How to Turn Off Quotas” on page 247

---

### What Are Quotas?

Quotas enable system administrators to control the size of UFS file systems by limiting 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, public and `/tmp` file systems usually do not benefit as much from the establishment of quotas.

Setting up quotas involves these general steps:

1. Enable file system quotas by issuing a series of commands, ensuring that quotas are enforced each time the system is rebooted and the file system is mounted. Entries must be added to the `/etc/vfstab` file, and a `quotas` file must be created in the root of the file system.

2. After a quota is created for one user, the quota can be copied as a prototype to set up other user quotas.
3. Before quotas are turned on, another command checks for consistency by comparing the proposed quotas to the current disk usage to make sure there are no conflicts.
4. Finally, a command turns on the quotas for one or more file systems.

These steps ensure that quotas are automatically activated on a file system each time it is mounted. For step-by-step instructions, see Chapter 17.

Once quotas are in place, they can be changed to adjust the amount of disk space or number of inodes that users can consume. Additionally, quotas can be added or removed as system needs change. For instructions on how to change quotas, disable individual quotas, or remove quotas from file systems, see “Changing and Removing Quotas” on page 243.

---

## Using Quotas

Using quotas enable system administrators to control the size of UFS file systems by limiting 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.

Once they are in place, quotas can be changed to adjust the amount of disk space or number of inodes that users can consume. Additionally, quotas can be added or removed as system needs change. See “Changing and Removing Quotas” on page 243 for instructions on changing quotas or the amount of time that quotas can be exceeded, disabling individual quotas, or removing quotas from file systems.

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 241.

## Setting Soft Limits and Hard Limits for Quotas

You can set both soft and hard limits. The system will not allow a user to exceed his or her hard limit. However, a system administrator may set a soft limit (sometimes referred to as a quota), which the user can temporarily exceed. The soft limit must be less than the hard limit.

Once the user exceeds the soft limit, a timer begins. While the 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 timer is also exceeded (more than seven days), the user will not be able to 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 (usually made up of 1 Kbyte blocks).

Assuming there are no directories, 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 simply creating one file large enough to consume all the data blocks in the user's quota.

---

## Setting Up Quotas

You can set up quotas to limit the amount of disk space and number of inodes (roughly equivalent to the number of files) available to users. These quotas are activated automatically each time a file system is mounted. This section describes how to configure file systems for quotas, and how to set up and activate quotas.

Setting up quotas involves these general steps:

1. A series of commands prepares a file system to accept quotas, ensuring that quotas will be enforced each time the system is rebooted and the file system is mounted. Entries must be added to the `/etc/vfstab` file, and a `quotas` file must be created in the top-level directory of the file system.
2. After a quota is created for one user, it can be copied as a prototype to set up other user quotas.
3. Before quotas are actually turned on, another command checks for consistency by comparing the proposed quotas with the current disk usage to make sure that there are no conflicts.

4. Finally, a command turns the quotas on for one or more entire file systems.

These steps ensure that quotas are automatically activated on a file system each time it is mounted. For specific information about these procedures, see “Setting Up Quotas (Task Map)” on page 237.

The following table describes the commands you use to set up disk quotas.

**TABLE 17-1** 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, and 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 space and how many inodes to allocate to each user. If you want to be sure 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.

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, you must 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. Also, if systems are rebooted infrequently, it is a good idea to periodically run the `quotacheck` command.

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 <code>/etc/vfstab</code> file so that quotas are activated each time the file system is mounted, and create a <code>quotas</code> file. | "How to Configure File Systems for Quotas" on page 237 |
| 2. Set up quotas for a user                    | Use the <code>edquota</code> command to create disk quotas and inode quotas for a single user account.                                           | "How to Set Up Quotas for a User" on page 238          |
| 3. (Optional) Set up quotas for multiple users | Use the <code>edquota</code> command to apply prototype quotas to other user accounts.                                                           | "How to Set Up Quotas for Multiple Users" on page 239  |
| 4. Check for consistency                       | Use the <code>quotacheck</code> command to compare quotas to current disk usage for consistency across one or more file systems.                 | "How to Check Quota Consistency" on page 240           |
| 5. Turn on quotas                              | Use the <code>quotaon</code> command to initiate quotas on one or more file systems.                                                             | "How to Turn On Quotas" on page 240                    |

### ▼ How to Configure File Systems for Quotas

1. **Become superuser.**
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 root access only.

```
chmod 600 quotas
```

## Examples—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 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 `/etc/vfstab` 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/rdisk/c0t4d0s0 /work ufs 3 yes rq
```

## ▼ How to Set Up Quotas for a User

1. Become superuser.

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 0 (the default) to the quotas you specify for each file system.

4. Verify the user's quota.

```
quota -v username
```

|                       |                                                                                       |
|-----------------------|---------------------------------------------------------------------------------------|
| <code>-v</code>       | Displays the user's quota information on all mounted file systems where quotas exist. |
| <code>username</code> | Specifies the user name to view quota limits.                                         |

## Examples—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 containing 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

1. **Become superuser.**
2. **Use the quota editor to apply the quotas you already established for a prototype user to the additional users you specify.**

```
edquota -p prototype-user username ...
```

|                             |                                                            |
|-----------------------------|------------------------------------------------------------|
| <code>prototype-user</code> | User name of the account for which you have set up quotas. |
| <code>username ...</code>   | Specifies one or more user names of additional accounts.   |

## Example—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

---

**Note** – To ensure accurate disk data, the file systems being checked should be quiescent when you run the `quotacheck` command manually. The `quotacheck` command is run automatically when a system is rebooted.

---

1. **Become superuser.**
2. **Run a consistency check on UFS file systems.**

```
quotacheck [-va] filesystem
```

|                         |                                                                                             |
|-------------------------|---------------------------------------------------------------------------------------------|
| <code>-v</code>         | (Optional) Identifies the disk quotas for each user on a particular file system.            |
| <code>-a</code>         | Checks all file systems with an <code>rq</code> entry in the <code>/etc/vfstab</code> file. |
| <code>filesystem</code> | Specifies the file system to check.                                                         |

See `quotacheck(1M)` for more information.

### Example—Checking Quota Consistency

The following example shows how to check quotas for the `/export/home` file system on the `/dev/rdisk/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/rdisk/c0t0d0s7 (/export/home)
```

## ▼ How to Turn On Quotas

1. **Become superuser.**
2. **Turn on file system quotas.**

```
quotaon [-v] -a filesystem ...
```

|                 |                                                                     |
|-----------------|---------------------------------------------------------------------|
| <code>-v</code> | Displays a message for each file system after quotas are turned on. |
|-----------------|---------------------------------------------------------------------|



|                             |                                                                                                                                                         |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-a</code>             | Turns on quotas for all file systems with an <code>rq</code> entry in the <code>/etc/vfstab</code> file.                                                |
| <code>filesystem ...</code> | 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—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
```

## 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 you use to check quotas.

**TABLE 17-2** Commands for Checking Quotas

| Command               | Task                                                                                                  |
|-----------------------|-------------------------------------------------------------------------------------------------------|
| <code>quota</code>    | Displays user quotas and current disk use, and information about users who are exceeding their quotas |
| <code>repquota</code> | Displays quotas, files, and the amount of space 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.

1. **Become superuser.**
2. **Display user quotas for mounted file systems where quotas are enabled.**

```
quota [-v] username
```

|                       |                                                                                  |
|-----------------------|----------------------------------------------------------------------------------|
| <code>-v</code>       | Displays one or more users' quotas on all mounted file systems that have quotas. |
| <code>username</code> | Is the login name or UID of a user's account.                                    |

## Example—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
```

|                         |                                                          |
|-------------------------|----------------------------------------------------------|
| <code>Filesystem</code> | Is the mount point for the file system.                  |
| <code>usage</code>      | Is the current block usage.                              |
| <code>quota</code>      | Is the soft block limit.                                 |
| <code>limit</code>      | Is the hard block limit.                                 |
| <code>timeleft</code>   | Is the amount of time (in days) left on the quota timer. |
| <code>files</code>      | Is the current inode usage.                              |
| <code>quota</code>      | Is the soft inode limit.                                 |
| <code>limit</code>      | Is the hard inode limit.                                 |
| <code>timeleft</code>   | 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.

1. **Become superuser.**
2. **Display all quotas for one or more file systems, even if there is no usage.**

```
repquota [-v] -a filesystem
```

|                 |                                                                                 |
|-----------------|---------------------------------------------------------------------------------|
| <code>-v</code> | Reports on quotas for all users, even those users who do not consume resources. |
| <code>-a</code> | Reports on all file systems.                                                    |

*filesystem*

Reports on the specified file system.

## Example—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

|                       |                                                          |
|-----------------------|----------------------------------------------------------|
| <code>used</code>     | Is the current block usage.                              |
| <code>soft</code>     | Is the soft block limit.                                 |
| <code>hard</code>     | Is the hard block limit.                                 |
| <code>timeleft</code> | Is the amount of time (in days) left on the quota timer. |

### File limits

|                       |                                                          |
|-----------------------|----------------------------------------------------------|
| <code>used</code>     | Is the current inode usage.                              |
| <code>soft</code>     | Is the soft inode limit.                                 |
| <code>hard</code>     | Is the hard inode limit.                                 |
| <code>timeleft</code> | 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 users can consume. You can also remove quotas, for individual users or from entire file systems, as needed.

The following table describes the commands you use to change quotas or remove quotas.

**TABLE 17-3** Commands for Changing Quotas and Removing Quotas

| Command  | Man Page     | Description                                                                                                                                                                       |
|----------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| edquota  | edquota      | Changes the hard limits and soft limits on the number of inodes or amount of disk space for each user. Also, changes the soft quota time limit for each file system with a quota. |
| quotaoff | quotaoff(1M) | Turns off quotas for specified file systems.                                                                                                                                      |

## ▼ How to Change the Soft Time 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 may exceed their disk space quotas or inode quotas by using the `edquota` command.

1. **Become superuser.**
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 you specify by numbers and the keywords month, week, day, hour, min, or sec.**

---

**Note** – This procedure does not affect current quota violators.

---

## Examples—Changing the Soft Time 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 0 (default) value 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 two weeks, and 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

1. Become superuser.
2. Use the quota editor to open a temporary file containing 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** – Although you can specify multiple users as arguments to the `edquota` command, the information displayed does not show which user this information belongs to, which could create some confusion.

---

3. Enter 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.

## Examples—Changing 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 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 example shows the same temporary file after quotas have been changed.

```
fs /files blocks (soft = 0, hard = 500) inodes (soft = 0, hard = 100)
```

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

1. Become superuser.
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** – Although you can specify multiple users as arguments to the `edquota` command, the information displayed does not show which user this information belongs with, which could create some confusion.

---

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** – Be sure 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.

## Examples—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

1. Become superuser.
2. Turn off file system quotas.

```
quotaoff [-v] -a filesystem ...
```

|                         |                                                                                                                                                     |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-v</code>         | Displays a message from each file system when quotas are turned off.                                                                                |
| <code>-a</code>         | Turns off quotas for all file systems.                                                                                                              |
| <code>filesystem</code> | Turns off quotas for one or more file systems you specify. More than one file system is specified by separating each file system name with a space. |

### Example—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
```





---

## Scheduling System Tasks (Tasks)

---

This chapter describes how to schedule routine or single (one-time) systems tasks by using the `crontab` and `at` commands. This chapter also explains how to control access to these commands by using `cron.deny`, `cron.allow`, and `at.deny` files.

This is a list of the step-by-step instructions in this chapter.

- “How to Create or Edit a `crontab` File” on page 254
- “How to Display a `crontab` File” on page 255
- “How to Remove a `crontab` File” on page 257
- “How to Deny `crontab` Access” on page 258
- “How to Limit `crontab` Access to Specified Users” on page 259
- “How to Create an `at` Job” on page 261
- “How to Display the `at` Queue” on page 262
- “How to Display `at` Jobs” on page 263
- “How to Remove `at` Jobs” on page 263
- “How to Deny Access to the `at` Command” on page 264

---

## 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` 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 18–1** Command Summary: Scheduling System Tasks

| Command              | What It Schedules                          | Location of Files                     | Files That Control Access                                                  |
|----------------------|--------------------------------------------|---------------------------------------|----------------------------------------------------------------------------|
| <code>crontab</code> | Multiple system tasks at regular intervals | <code>/var/spool/cron/crontabs</code> | <code>/etc/cron.d/cron.allow</code> and <code>/etc/cron.d/cron.deny</code> |
| <code>at</code>      | A single system task                       | <code>/var/spool/cron/atjobs</code>   | <code>/etc/cron.d/at.deny</code>                                           |

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 "Working With the 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:

- 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:

- 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:

- 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 254.

## 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, and then are removed from their directory. Therefore, `at` is most useful for running simple commands or scripts that direct output into separate files for later examination.

Submitting an `at` job involves entering 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 260.

The `at` command stores the command or script you entered, 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 your `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 261.

---

## 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 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
15 3 * * 0 /usr/lib/fs/nfs/nfsfind
1 2 * * * [-x /usr/sbin/rtc] && /usr/sbin/rtc -c > /dev/null 2>&1
30 3 * * * [-x /usr/lib/gss/gsscred_clean] && /usr/lib/gss/gsscred_clean

```

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. 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 253.

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 18-2** Default crontab Files

| crontab File | Function                                         |
|--------------|--------------------------------------------------|
| adm          | Accounting                                       |
| lp           | Printing                                         |
| root         | General system functions and file system cleanup |
| sys          | Performance 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 checks for new crontab files or changes to existing ones, reads the execution times listed within the files, and submits the commands for execution at the proper times.

In much the same way, the `cron` daemon controls the scheduling of `at` files, which are stored in the `/var/spool/cron/atjobs` directory.

## Syntax of `crontab` File Entries

A `crontab` file consists of commands, one per line, that execute automatically at the time specified by the first five fields at the beginning of each command line. These first five fields, described in the following table, are separated by spaces.

**TABLE 18-3** 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 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, because `crontab` does not recognize extra carriage returns. For more detailed information about `crontab` entries and command options, refer to `crontab(1)`.

---

## 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 set up for your system environment, which is defined by the `EDITOR` environment variable. If this variable has not been set, `crontab` 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

1. **(Optional) Become superuser to create or edit a crontab file that belongs to root or another user.**
2. **Create a new crontab file, or edit an existing one.**

```
$ crontab -e [username]
```

Where *username* specifies the name of the user's account for which you want to create or edit a crontab file. To create or edit crontab files 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 is overwritten with an empty file.

---

3. **Add command lines to the file.**  
Follow the syntax described in “Syntax of crontab File Entries” on page 253. The crontab file will be placed in `/var/spool/cron/crontabs`.
4. **Verify your crontab file changes.**

```
crontab -l [username]
```

## Example—Creating or Editing 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` to make sure 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 a crontab File

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 display shows that crontab files exist for users `smith` and `jones`.

```
$ 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 255.

---

## Displaying crontab Files

The `crontab -l` command displays the contents of your crontab file much the way the `cat` command displays the contents of other types of files. You do not have to change directories to `/var/spool/cron/crontabs` (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

1. (Optional) Become superuser to display a crontab file that belongs to root or another user.

## 2. Display the crontab file.

```
$ crontab -l [username]
```

Where *username* specifies the name of the user's account for which you want to create or edit a crontab file. To create or edit crontab files 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 is overwritten with an empty file.

---

## Example—Displaying a crontab File

The following example shows how to use the `crontab -l` command to display the contents of the user's default crontab file, the default root crontab file, and the crontab file belonging to another user.

```
$ crontab -l
13 13 * * * chmod g+w /home1/documents/*.book > /dev/null 2>&1
$ suPassword:
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.
#
The rtc command is run to adjust the real time clock if and when
daylight savings time changes.
#
10 3 * * * /usr/sbin/logadm
15 3 * * 0 /usr/lib/fs/nfs/nfsfind
1 2 * * * [-x /usr/sbin/rtc] && /usr/sbin/rtc -c > /dev/null 2>&1
30 3 * * * [-x /usr/lib/gss/gsscred_clean] && /usr/lib/gss/gsscred_clean
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, `crontab -r` removes your own crontab file. You must be superuser to remove crontab files that belong to superuser or other users.



You do not have to change directories to `/var/spool/cron/crontabs` (where crontab files are located) to use this command.

## ▼ How to Remove a crontab File

1. (Optional) Become superuser to remove a crontab file that belongs to root or another user.
2. Remove the crontab file.

```
$ crontab -r [username]
```

Where *username* specifies the name of the user's account for which you want to create or edit a crontab file. To create or edit crontab files 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 is overwritten with an empty file.

---

3. Verify that the crontab file is removed.

```
ls /var/spool/cron/crontabs
```

### Example—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 the `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 per line. These access control files work together as follows:

- If `cron.allow` exists, only the users 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 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, 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 the ones 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 Access

1. **Become superuser.**
2. **Edit the `/etc/cron.d/cron.deny` file and add user names, one per line, who will be prevented from using crontab commands.**

```
daemon
bin
smtp
nuucp
listen
nobody
noaccess
username1
username2
```

```
username3
```

```
.
.
.
```

3. **Verify the `/etc/cron.d/cron.deny` file.**

```
cat /etc/cron.d/cron.deny
```

## ▼ How to Limit `crontab` Access to Specified Users

1. **Become superuser.**
2. **Create the `/etc/cron.d/cron.allow` file.**
3. **Enter 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. **Enter the user names, one per line, that will be allowed to use the `crontab` command.**

```
root
username1
username2
username3
.
.
.
```

## Examples—Limiting `crontab` Access to Specified Users

The following example shows a `cron.deny` file that prevents user names `visitor`, `jones`, and `temp` from accessing the `crontab` command.

```
$ cat /etc/cron.d/cron.denydaemon
bin
smtp
nuucp
listen
nobody
noaccess
jones
temp
visitor
```

The following example shows a `cron.allow` file. The users `smith`, `jones`, `lp`, and `root` are the only ones who can access the `crontab` command.

```
$ cat /etc/cron.d/cron.allow
root
jones
```

```
lp
smith
```

## ▼ How to Verify Limited crontab Access

To verify if a specific user can access crontab, use the `crontab -l` command while you are logged into the user account.

```
$ crontab -l
```

If the user can access crontab, and already has created a crontab file, the file is displayed. Otherwise, if the user can access crontab but no crontab file exists, a message such as the following is displayed:

```
crontab: can't open your crontab file
```

This user either is listed in `cron.allow` (if the file exists), or the user is not listed in `cron.deny`.

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 `cron.allow` (if the file exists), or the user is listed in `cron.deny`.

---

## Scheduling a Single System Task (at)

The following sections describe how to use the `at` command to schedule jobs (commands and scripts) for execution at a later time, how to display and remove these jobs, and 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, which becomes the job's file name.

### Description of the `at` Command

Submitting an `at` job file involves:

1. Invoking the `at` utility and specifying a command execution time.
2. Entering 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 2001
```

## 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 per line. The users listed in this file cannot access `at` commands.

The `at.deny` file, 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` access you want to restrict.

## ▼ How to Create an `at` Job

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.

|             |                                                                                                                                                                                           |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>time</i> | Hour that you want to schedule the job. Add am or pm if you do not specify the hours according to a 24-hour clock. Acceptable keywords are midnight, noon, and now. Minutes are optional. |
| <i>date</i> | 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 you want to execute, one per line. You may enter 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.

## Examples—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 2001
```

She received a mail 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 `big.file`.

```
$ at 4 am Saturday
at> sort -r /usr/dict/words > /export/home/jones/big.file
```

### ▼ How to Display the `at` Queue

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 created.

```
$ atq
```

### ▼ How to Verify an `at` Job

To verify that you have created an `at` job, use the `atq` command. 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, 2001 19:30 jones 897355800.a a stdin
 2nd Jul 14, 2001 23:45 jones 897543900.a a stdin
 3rd Jul 17, 2001 04:00 jones 897732000.a a stdin
```

## ▼ How to Display at Jobs

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—Displaying at Jobs

The following example shows output from the `at -l` command, which provides status information on all jobs submitted by a user.

```
$ at -l
897543900.a Sat Jul 14 23:45:00 2001
897355800.a Thu Jul 12 19:30:00 2001
897732000.a Tue Jul 17 04:00:00 2001
```

The following example shows the output displayed when a single job is specified with the `at -l` command.

```
$ at -l 897732000.a
897732000.a Tue Jul 17 04:00:00 2001
```

## ▼ How to Remove at Jobs

1. (Optional) Become superuser to remove an at job that belongs to root or another user.
2. 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.

3. 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—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 2001
897355800.a Thu Jul 12 19:30:00 2001
897732000.a Tue Jul 17 04:00:00 2001
$ at -r 897732000.a
$ at -l 897732000.a
at: 858142000.a: No such file or directory
```

## ▼ How to Deny Access to the at Command

1. **Become superuser.**
2. **Edit the `/etc/cron.d/at.deny` file and add the names of users, one per line, that will be prevented from using at commands.**

```
daemon
bin
smtp
nuucp
listen
nobody
noaccess
username1
username2
username3
.
.
.
```

## Example—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 Access Is Denied

To verify that a username was added correctly to `/etc/cron.d/at.deny`, use the `at -l` command while logged in as the user. If the user 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` access is allowed, the `at -l` command returns nothing.



## Managing System Accounting (Tasks)

---

This section describes how to set up and maintain system accounting.

This is a list of the step-by-step instructions in this chapter.

- “How to Set Up System Accounting” on page 273
- “How to Bill Users” on page 275
- “How to Fix a Corrupted wtmpx File” on page 276
- “How to Fix tacct Errors” on page 276
- “How to Restart the runacct Script” on page 277
- “How to Set Up System Accounting” on page 273
- “How to Permanently Disable System Accounting” on page 278

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

- “What is System Accounting?” on page 267
- “Setting Up System Accounting” on page 272

For information on using extended accounting, see “Extended Accounting” in *System Administration Guide: Resource Management and Network Services*.

For reference information on the various system accounting reports, see Chapter 20.

---

## What is System Accounting?

The SunOS 5.9 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 the accounting programs to:

- 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 illustrates 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 "The `runacct` Script" on page 281.

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 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` directories. The accounting reports reside in the `/var/adm/acct` directory.

Daily accounting can help you perform four types of auditing:

- Connect
- Process
- Disk
- Fee calculations

### Connect Accounting

Connect accounting enables you to determine:

- 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, the system stores:

- Records of time adjustments
- Boot times
- Times the accounting software was turned off and on
- Changes in run levels
- The creation of user processes (login processes and init processes), and the terminations of processes.

These records (produced from the output of system programs such as `date`, `init`, `login`, `ttymon`, and `acctwtmp`) 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
- Commands run
- The `tty` controlling 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 used by the user's files

This data is collected by the `/usr/lib/acct/dodisk` shell script at intervals 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, which gather disk usage by login.



---

**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 `di_size` value of the inode.

## Calculating User Fees

The `chargefee` utility stores charges for special services provided to a user, such as file restoration, in the `/var/adm/fee` file. 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 275.

## 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 `acctwtmpt` 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.

**TABLE 19-1** Raw Accounting Data

| File in <code>/var/adm</code> | Information Stored | Written By                             | Format |
|-------------------------------|--------------------|----------------------------------------|--------|
| <code>wtmpx</code>            | Connect sessions   | <code>login</code> , <code>init</code> | binary |
|                               | Changes            | <code>date</code>                      |        |

**TABLE 19-1** Raw Accounting Data (Continued)

| File in <code>/var/adm</code>   | Information Stored | Written By                                                                                                                        | Format |
|---------------------------------|--------------------|-----------------------------------------------------------------------------------------------------------------------------------|--------|
|                                 | Reboots            | <code>acctwtmp</code>                                                                                                             |        |
|                                 | Shutdowns          | <code>shutacct</code>                                                                                                             |        |
| <code>pacctn</code>             | Processes          | Kernel (when the process ends)<br><br><code>turnacct switch</code> (which creates a new file when the old one reaches 500 blocks) | binary |
| <code>fee</code>                | Special charges    | <code>chargefee</code>                                                                                                            | ASCII  |
| <code>acct/nite/diskacct</code> | Disk space used    | <code>dodisk</code>                                                                                                               | 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 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 past 500 blocks (default), the `turnacct switch` command is executed. (The program moves the `pacct` file to the `pacctn` file and creates a new one.) 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: `/var/adm/pacctn`, `/var/adm/wtmpx`, `/var/adm/fee`, and `/var/adm/acct/nite/diskacct`, to produce command summaries and usage summaries by user name.
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/rprtMMDD` files.
9. The `monacct` script should be executed on a monthly basis (or at intervals you determine, such as the end of every fiscal period). The `monacct` script creates a report 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 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.

---

## Setting Up System Accounting

You can set up system accounting to run while the system is in multiuser mode (system state 2). Generally, this task involves:

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 19-2** Default Accounting Scripts

| Accounting Script | Man Page    | Purpose                                                                                                                                                                                                             | Run Frequency     |
|-------------------|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| ckpacct           | ckpacct(1M) | Checks the size of the /usr/adm/pacct log file and makes sure it does not get too large.                                                                                                                            | Periodically      |
| runacct           | runacct(1M) | Processes connect, disk, and fee accounting information. You can remove the commands from this script for the accounting features you do not want processed.                                                        | Daily             |
| monacct           | monacct(1M) | Generates fiscal accounting summary reports on a monthly basis. You can determine how often this script is run. You can remove the commands from this script for the accounting features you do not want generated. | On a fiscal basis |



You can choose which accounting scripts run by default. After these entries have been added to the crontab files, accounting should run automatically.

## ▼ How to Set Up System Accounting

1. **Become superuser.**
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 and local holidays.**  
For more information, see holidays(4) and the example that follows.

8. **Reboot the system, or start accounting manually by typing:**

```
/etc/init.d/acct start
```

## Examples—Setting Up Accounting

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
```

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.
#
The rtc command is run to adjust the real time clock if and when
daylight savings time changes.
#
10 3 * * * /usr/sbin/logadm
15 3 * * 0 /usr/lib/fs/nfs/nfsfind
1 2 * * * [-x /usr/sbin/rtc] && /usr/sbin/rtc -c > /dev/null 2>&1
30 3 * * * [-x /usr/lib/gss/gsscred_clean] && /usr/lib/gss/gsscred_clean
30 22 * * 4 /usr/lib/acct/dodisk
```

The following example shows a sample /etc/acct/holidays file.

```
* @(#)holidays January 1, 2001
*
* Prime/Nonprime Table for UNIX Accounting System
*
* Curr Prime Non-Prime
* Year Start Start
*
 1999 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, such as 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 acctsh(1M) for more information.

## ▼ How to Bill Users

1. Become superuser.
2. Charge a user for special services.

```
/usr/lib/acct/chargefee username amount
```

|                 |                                                                                                                                                                                                                                                             |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>username</i> | User account you want to bill.                                                                                                                                                                                                                              |
| <i>amount</i>   | Number of units to bill the user. This is an arbitrary unit that you set to charge users based on some task like printing or restoring a file. You would have to write a script that invokes <code>chargefee</code> and charges a user for a specific task. |

### Example—Billing Users

The following example charges the user `print_customer` 10 units.

```
/usr/lib/acct/chargefee print_customer 10
```

---

## Maintaining Accounting Information

This section describes how to fix corrupted 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 of the 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 the system accounting. When the date is changed manually and the system is in multiuser mode, a set of date change records is written into 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

1. **Become superuser.**
2. **Change to the `/var/adm` directory.**
3. **Convert the wtmpx file from binary to ASCII format.**

```
/usr/lib/acct/fwtmp < wtmpx > wtmpx.ascii
```

4. **Edit `wtmpx.ascii` 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 `fwtmp(1M)` 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 it. If the contents look all right, patch the latest `/var/adm/acct/sum/tacctMMDD` file, then recreate the `/var/adm/acct/sum/tacct` file. The following steps outline a simple patch procedure.

## ▼ How to Fix tacct Errors

1. **Become superuser.**
2. **Change to the `/var/adm/acct/sum` directory.**
3. **Convert the `tacctMMDD` file from binary 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.**

```
/usr/lib/acct/acctmerg -i < xtacct > tacctMMDD
```

*MMDD* is pair of two-digit numbers that represent the month and day.
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 a variety of reasons. The most common reasons are a system crash, the /var directory running out of space, or a corrupted wtmpx file. If the active.MMDD file exists, check it first for error messages. If the active and lock files exist, check fd2log 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 runacct(1M).



---

**Caution** – When you run the runacct program manually, be sure to run it as user adm.

---

### ▼ How to Restart the runacct Script

1. **Remove the lastdate file and any lock\* files, if any.**

```
$ cd /var/adm/acct/nite
$ rm lastdate lock*
```

The lastdate file contains the date that the runacct program was last run. Restarting runacct in the next step recreates this file.

2. **Restart the runacct script.**

```
$ /usr/lib/acct/runacct MMDD [state] 2> /var/adm/acct/nite/fd2log &
```

|              |                                                                                  |
|--------------|----------------------------------------------------------------------------------|
| <i>MMDD</i>  | Month and day specified by two-digit numbers.                                    |
| <i>state</i> | Specifies a state, or starting point, where the runacct processing should begin. |

---

# Stopping and Disabling System Accounting

You can temporarily stop system accounting or disable it permanently.

## ▼ How to Temporarily Stop System Accounting

1. **Become superuser.**
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 `crontab` file for user `root` 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 accounting program.**
5. **(Optional) Remove the newly added comment symbols from the `crontab` files and restart the accounting program to re-enable system accounting.**

```
/etc/init.d/acct start
```

## ▼ How to Permanently Disable System Accounting

1. **Become superuser.**
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 accounting program.**

```
/etc/init.d/acct stop
```





## System Accounting (Reference)

---

This chapter provides reference information about system accounting.

This is a list of reference information in this chapter.

- “Daily Accounting Reports ” on page 283
- “The runacct Script” on page 281
- “System Accounting Files” on page 291

---

### The runacct Script

The main daily accounting script, `runacct`, is normally invoked by the `cron` command outside of prime 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 are used to recognize an error, provide intelligent diagnostics, and complete processing in such a way that the `runacct` script can be restarted with minimal intervention. It records its progress by writing descriptive messages into 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 into 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 `runacct`. 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, a

message is written to the console, mail is sent to `root` and `adm`, locks might be removed, diagnostic files are saved, and execution is ended. For instructions on how to start the `runacct` script again, see “How to Restart the `runacct` Script” on page 277.

To allow the `runacct` script to be restarted, processing is broken down into separate re-entrant states. The `statefile` file is used to keep track of 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 20-1** `runacct` States

| State   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SETUP   | The <code>turnacct</code> switch command is executed to create a new <code>pacct</code> file. The <code>/var/adm/pacctn</code> process accounting files (except for the <code>pacct</code> file) are moved to the <code>/var/adm/spacctn.MMDD</code> files. The <code>/var/adm/wtmpx</code> file is moved to the <code>/var/adm/acct/nite/wtmp.MMDD</code> file (with the current time record added on the end) and a new <code>/var/adm/wtmp</code> file is created. The <code>closewtmp</code> and <code>utmp2wtmp</code> programs add records to the <code>wtmp.MMDD</code> file and the new <code>wtmpx</code> file to account for users currently logged in. |
| WTMPFIX | The <code>wtmpfix</code> program checks the <code>wtmp.MMDD</code> file in the <code>nite</code> directory for accuracy. Because some date changes cause the <code>acctcon</code> program to fail, the <code>wtmpfix</code> program attempts to adjust the time stamps in the <code>wtmpx</code> file if a record of a date change appears. This program also deletes any corrupted entries from the <code>wtmpx</code> file. The fixed version of the <code>wtmp.MMDD</code> file is written to the <code>tmpwtmp</code> file.                                                                                                                                   |
| CONNECT | The <code>acctcon</code> program is used to record connect accounting records in the file <code>ctacct.MMDD</code> . These records are in <code>tacct.h</code> format. In addition, <code>acctcon</code> creates the <code>lineuse</code> and <code>reboots</code> files. The <code>reboots</code> file records all the boot records found in the <code>wtmpx</code> file.                                                                                                                                                                                                                                                                                        |
| PROCESS | The <code>acctprc</code> program is used to convert the <code>/var/adm/spacctn.MMDD</code> process accounting files into total accounting records in the <code>ptacctn.MMDD</code> files. The <code>Spacct</code> and <code>ptacct</code> files are correlated by number so that if the <code>runacct</code> script fails, the <code>Spacct</code> files are not processed.                                                                                                                                                                                                                                                                                       |
| MERGE   | The <code>acctmerg</code> program merges the process accounting records with the connect accounting records to form the <code>daytacct</code> file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| FEEES   | The <code>acctmerg</code> program merges ASCII <code>tacct</code> records from the <code>fee</code> file into the <code>daytacct</code> file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**TABLE 20-1** runacct States (Continued)

| State      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DISK       | If the <code>dodisk</code> script procedure has been run, which produces the <code>disktacct</code> file, the <code>DISK</code> program merges the file into the <code>daytacct</code> file and moves the <code>disktacct</code> file to the <code>/tmp/disktacct.MMDD</code> file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| MERGETACCT | The <code>acctmerg</code> program merges the <code>daytacct</code> file with the <code>sum/tacct</code> file, the cumulative total accounting file. Each day, the <code>daytacct</code> file is saved in the <code>sum/tacct.MMDD</code> file so that the <code>sum/tacct</code> file can be re-created if it is corrupted or lost.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| CMS        | The <code>acctcms</code> program is run several times. This program is first run to generate the command summary by using the <code>Spacctn</code> files and write the data to the <code>sum/daycms</code> file. The <code>acctcms</code> program is then run to merge the <code>sum/daycms</code> file with the <code>sum/cms</code> cumulative command summary file. Finally, the <code>acctcms</code> program is run to produce <code>nite/daycms</code> and <code>nite/cms</code> , the ASCII command summary files from the <code>sum/daycms</code> and <code>sum/cms</code> files, respectively. The <code>lastlogin</code> program is used to create the <code>/var/adm/acct/sum/loginlog</code> log file, the report of when each user last logged in. If the <code>runacct</code> 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 included at this point. The <code>runacct</code> script expects it to be called the <code>/usr/lib/acct/runacct.local</code> program.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| CLEANUP    | Cleans up temporary files, runs the <code>prdaily</code> script and saves its output in the <code>sum/rpt.MMDD</code> file, removes the locks, and then exits.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



**Caution** – When restarting the `runacct` script in the `CLEANUP` state, remove the last `ptacct` file because it will not be complete.

---

## Daily Accounting Reports

The `runacct` shell script generates five basic reports upon each invocation. The following table describes the five basic reports generated.

**TABLE 20-2** Daily Accounting Reports

| Report Type                            | Description                                                                                                                                                                                                                                         |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| "Daily Report" on page 284             | Shows terminal line utilization by tty number.                                                                                                                                                                                                      |
| "Daily Usage Report " on page 285      | Indicates usage of system resources by users (listed in order of user ID).                                                                                                                                                                          |
| "Daily Command Summary " on page 286   | 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 with the monthly command summary. |
| "Monthly Command Summary " on page 288 | A cumulative summary that reflects the data accumulated since the last invocation of the <code>monacct</code> program.                                                                                                                              |
| "Last Login Report " on page 288       | Shows the last time each user logged in (arranged in chronological order).                                                                                                                                                                          |

## Daily Report

This report gives information about each terminal line used. The following is a sample Daily Report.

```
Oct 16 02:30 2001 DAILY REPORT FOR venus Page 1
```

```
from Mon Oct 15 02:30:02 2001
to Tue Oct 16 02:30:01 2001
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, comes a log of system reboots, shutdowns, power failure recoveries, and any other record dumped into the `/var/adm/wtmpx` file by the `acctwtmp` program. For more information, see `acct(1M)`.

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 in the Daily Report.

**TABLE 20-3** Daily Report Data

| Column  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LINE    | The terminal line or access port.                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| MINUTES | The total number of minutes that the line was in use during the accounting period.                                                                                                                                                                                                                                                                                                                                                                                  |
| PERCENT | The TOTAL DURATION divided by the total 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, it listed the number of times that a line or port was used to log in a user.)                                                                                                                                                                                                                                                                                                                         |
| # OFF   | This column reflects the number of times a user logs out and any interrupts that occur on that line. Generally, interrupts occur on a port when <code>ttymon</code> 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 there is a bad connection 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 276.

## Daily Usage Report

The Daily Usage Report gives a breakdown of system resource utilization by user. A sample of this report follows.

```
Oct 16 02:30 2001 DAILY USAGE REPORT FOR skisun Page 1
```

| UID | LOGIN NAME | CPU PRIME | (MINS) NPRIME | KCORE- PRIME | MINS NPRIME | CONNECT PRIME | (MINS) NPRIME | DISK BLOCKS | # OF PROCS | # OF SESS | # OF DISK SAMPLES | FEE |
|-----|------------|-----------|---------------|--------------|-------------|---------------|---------------|-------------|------------|-----------|-------------------|-----|
| 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 in the Daily Usage Report.

**TABLE 20-4** 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 (non-prime) utilization. The accounting system's version of this data is located in the <code>/etc/acct/holidays</code> 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 into the system, or "real time." Divided into PRIME and NPRIME use. 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 <code>acctdusg</code> program, which runs the disk accounting programs and merges the accounting records ( <code>daytacct</code> ). 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 a user logged on to the system.                                                                                                                                                                                                                                                         |
| # DISK SAMPLES | Number of times 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 <code>chargefee</code> script.                                                                                                                                                                       |

## Daily Command Summary

The Daily Command Summary report shows the system resource use by command. With this report, you can identify the most heavily used commands and, based on how those commands use system resources, gain insight on how best to tune the system.

These reports are sorted by `TOTAL KCOREMIN`, which is an arbitrary gauge but often a good one 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 in the Daily Command Summary.

**TABLE 20-5** Daily Command Summary

| Column         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| COMMAND NAME   | Name of the command. Unfortunately, all shell procedures are lumped together under the name <code>sh</code> because only object modules are reported by the process accounting system. You should monitor the frequency of programs called <code>a.out</code> or <code>core</code> or any other unexpected name. You can use the <code>acctcom</code> 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 of the TOTAL KCOREMIN over the number of invocations reflected by NUMBER CMDS.                                                                                                                                                                                                                                                                                                                                                                   |
| MEAN CPU-MIN   | Mean derived between the NUMBER CMDS and TOTAL CPU-MIN.                                                                                                                                                                                                                                                                                                                                                                                               |
| HOG FACTOR     | Total CPU time divided by elapsed time. Shows the ratio of system availability to system use, providing a relative measure of total available CPU time consumed by the process during its execution.                                                                                                                                                                                                                                                  |
| CHARS TRNSFD   | Total number of characters pushed around by the read and write system calls. Might be negative due to overflow.                                                                                                                                                                                                                                                                                                                                       |

**TABLE 20-5** Daily Command Summary (Continued)

| Column      | Description                                                                   |
|-------------|-------------------------------------------------------------------------------|
| 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 report follows.

Oct 16 02:30 2001 MONTHLY TOTAL COMMAND SUMMARY Page 1

| COMMAND NAME | NUMBER CMDS | TOTAL COMMAND SUMMARY |               |                |             |              |            |              |             |
|--------------|-------------|-----------------------|---------------|----------------|-------------|--------------|------------|--------------|-------------|
|              |             | TOTAL KCOREMIN        | TOTAL CPU-MIN | TOTAL REAL-MIN | MEAN SIZE-K | MEAN CPU-MIN | HOG FACTOR | CHARS TRNSFD | BLOCKS 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 in the Monthly Command Summary, see "Daily Command Summary" on page 286.

## 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 sample report appears follows.

Oct 16 02:30 2001 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 |



## Looking at the `pacct` File With `acctcom`

At any time, you can examine the contents of the `/var/adm/pacctn` 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:

Sample `acctcom` output follows:

```
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
```

- Command name (pound (#) sign if the command was executed with superuser privileges)
- User name
- tty name (listed as ? if unknown)
- Command starting time
- Command ending time
- Real time (in seconds)
- CPU time (in seconds)
- Mean size (in Kbytes)

You can obtain the following information by using `acctcom` 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` options.

**TABLE 20-6** `acctcom` Options

| 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 <code>fork/exec</code> 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 containing the I/O counts in the output.                                                                                                                                                            |
| -k               | Shows total <code>kcore</code> minutes instead of memory size.                                                                                                                                                     |
| -m               | Shows mean core size. This 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 <i>sec</i>    | Shows only processes with total CPU time (system plus user) exceeding <i>sec</i> seconds.                                                                                                                          |
| -e <i>time</i>   | Shows processes existing at or before <i>time</i> , given in the format <code>hr[:min[:sec]]</code> .                                                                                                              |
| -E <i>time</i>   | Shows processes starting at or before <i>time</i> , given in the format <code>hr[:min[:sec]]</code> . Using the same time for both <code>-S</code> and <code>-E</code> , shows processes that existed at the time. |
| -g <i>group</i>  | Shows only processes that belong to <i>group</i> .                                                                                                                                                                 |
| -H <i>factor</i> | Shows only processes that exceed <i>factor</i> , where <i>factor</i> is the "hog factor" (see the <code>-h</code> option).                                                                                         |
| -I <i>chars</i>  | Shows only processes that transferred more characters than the cutoff number specified by <i>chars</i> .                                                                                                           |
| -l <i>line</i>   | Show only processes belonging to the terminal <code>/dev/line</code> .                                                                                                                                             |

**TABLE 20-6** `acctcom` Options (Continued)

| Option                  | Description                                                                                                       |
|-------------------------|-------------------------------------------------------------------------------------------------------------------|
| <code>-n pattern</code> | Shows only commands matching <i>pattern</i> (a regular expression except that "+" means one or more occurrences). |
| <code>-o ofile</code>   | Instead of printing the records, copies them in <code>acct .h</code> format to <i>ofile</i> .                     |
| <code>-O sec</code>     | Shows only processes with CPU system time exceeding <i>sec</i> seconds.                                           |
| <code>-s time</code>    | Show processes existing at or after <i>time</i> , given in the format <i>hr[:min[:sec]]</i> .                     |
| <code>-S time</code>    | Show processes starting at or after <i>time</i> , given in the format <i>hr[:min[:sec]]</i> .                     |
| <code>-u user</code>    | Shows only processes that belong to <i>user</i> .                                                                 |

## System Accounting Files

The `/var/adm` directory contains the active data collection files. The following table describes the accounting files in this directory.

**TABLE 20-7** Files in the `/var/adm` Directory

| File                      | Description                                                                                    |
|---------------------------|------------------------------------------------------------------------------------------------|
| <code>dtmp</code>         | Output from the <code>acctdusg</code> program                                                  |
| <code>fee</code>          | Output from the <code>chargefee</code> program, which are the ASCII <code>tacct</code> records |
| <code>pacct</code>        | Active process accounting file                                                                 |
| <code>pacctn</code>       | Process accounting files switched by running the <code>turnacct</code> script                  |
| <code>Spacctn.MMDD</code> | Process accounting files for <i>MMDD</i> during execution of the <code>runacct</code> script   |

The `/var/adm/acct` directory contains the `nite`, `sum`, and `fiscal` directories, which 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 20-8** Files in the `/var/adm/acct/nite` Directory

| File                        | Description                                                                                                                                                                               |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>active</code>         | Used by the <code>runacct</code> script to record progress and print warning and error messages                                                                                           |
| <code>active.MMDD</code>    | Same as the <code>active</code> file after the <code>runacct</code> script detects an error                                                                                               |
| <code>cms</code>            | ASCII total command summary used by <code>prdaily</code>                                                                                                                                  |
| <code>ctacct.MMDD</code>    | Connect accounting records in <code>tacct.h</code> format                                                                                                                                 |
| <code>ctmp</code>           | Output of <code>acctcon1</code> program, connect session records in <code>ctmp.h</code> format ( <code>acctcon1</code> and <code>acctcon2</code> are provided for compatibility purposes) |
| <code>daycms</code>         | ASCII daily command summary used by the <code>prdaily</code> script.                                                                                                                      |
| <code>daytacct</code>       | Total accounting records for one day in <code>tacct.h</code> format                                                                                                                       |
| <code>disktacct</code>      | Disk accounting records in <code>tacct.h</code> format, created by the <code>dodisk</code> script                                                                                         |
| <code>fd2log</code>         | Diagnostic output during execution of the <code>runacct</code> script                                                                                                                     |
| <code>lastdate</code>       | Last day the <code>runacct</code> script executed (in <code>date +%m%d</code> format)                                                                                                     |
| <code>lock</code>           | Used to control serial use of the <code>runacct</code> script                                                                                                                             |
| <code>lineuse</code>        | <code>tty</code> line usage report used by the <code>prdaily</code> script                                                                                                                |
| <code>log</code>            | Diagnostic output from the <code>acctcon</code> program                                                                                                                                   |
| <code>log.MMDD</code>       | Same as the <code>log</code> file after the <code>runacct</code> script detects an error                                                                                                  |
| <code>owtmpx</code>         | Previous day's <code>wtmpx</code> file                                                                                                                                                    |
| <code>reboots</code>        | Beginning and ending dates from the <code>wtmpx</code> file and a listing of reboots                                                                                                      |
| <code>statefile</code>      | Used to record current state during execution of the <code>runacct</code> script                                                                                                          |
| <code>tmpwtmp</code>        | <code>wtmpx</code> file corrected by the <code>wtmpfix</code> program                                                                                                                     |
| <code>wtmperror</code>      | Contains <code>wtmpfix</code> error messages                                                                                                                                              |
| <code>wtmperror.MMDD</code> | Same as the <code>wtmperror</code> file after the <code>runacct</code> script detects an error                                                                                            |
| <code>wtmp.MMDD</code>      | The <code>runacct</code> script's copy of the <code>wtmpx</code> 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 20–9** Files in the `/var/adm/acct/sum` Directory

| File                    | Description                                                                                                                       |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <code>cms</code>        | Total command summary file for current fiscal period in binary format                                                             |
| <code>cmsprev</code>    | Command summary file without latest update                                                                                        |
| <code>daycms</code>     | Command summary file for the day's usage in internal summary format                                                               |
| <code>loginlog</code>   | Record of last date each user logged on; created by the <code>lastlogin</code> script and used in the <code>prdaily</code> script |
| <code>rprt.MMDD</code>  | Saved output of <code>prdaily</code> script                                                                                       |
| <code>tacct</code>      | Cumulative total accounting file for current fiscal period                                                                        |
| <code>tacctprev</code>  | Same as <code>tacct</code> without latest update                                                                                  |
| <code>tacct.MMDD</code> | Total accounting file for <code>MMDD</code>                                                                                       |

The fiscal directory contains periodic summary files created by the `monacct` script. The following table summarizes the files in the `/var/adm/acct/fiscal` directory.

**TABLE 20–10** Files in the `/var/adm/acct/fiscal` Directory

| File                 | Description                                                                      |
|----------------------|----------------------------------------------------------------------------------|
| <code>cmsn</code>    | Total command summary file for fiscal period <i>n</i> in internal summary format |
| <code>fisrptn</code> | Report similar to <code>rprt<math>n</math></code> for fiscal period <i>n</i>     |
| <code>tacctn</code>  | Total accounting file for fiscal period <i>n</i>                                 |

## Files Produced by the `runacct` Script

The following table summarizes most useful files produced by the `runacct` script found in the `/var/adm/acct` directory.

**TABLE 20-11** Files Produced by runacct

| <b>File</b>   | <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 about three to one, the line is very likely failing. |
| nite/daytacct | The total accounting file for the day in tacct.h format.                                                                                                                                                                                                                                                                                                                                                                            |
| 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.                                                                                                                                                                                                                                                        |
| 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/cms       | This file is the accumulation of each day's command summaries. It is restarted by the execution of the monacct script. The ASCII version is the nite/cms file.                                                                                                                                                                                                                                                                      |
| sum/loginlog  | The runacct script calls 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.                                                                                                                                                                                                                                                                                                                       |

## Managing System Performance Topics

---

This topic map lists the chapters that provide information for managing system performance.

|            |                                                                                                                                                    |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Chapter 22 | Provides overview information about system performance topics.                                                                                     |
| Chapter 23 | Provides step-by-step instructions for using process commands to enhance system performance.                                                       |
| Chapter 24 | Provides step-by-step instructions for using <code>vmstat</code> , <code>sar</code> , and disk utilization commands to monitor system performance. |





## Managing System Performance (Overview)

---

Achieving good performance from a computer or network is an important part of system administration. This chapter is an overview of some factors that contribute to maintaining and managing the performance of the computer systems in your care.

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

- “Where to Find System Performance Tasks” on page 300
- “System Performance and System Resources” on page 300
- “Processes and System Performance” on page 301
- “About Monitoring Performance” on page 302

---

## What’s New in Managing System Performance?

This section describes new features for managing system performance.

### DNLC Improvements

The directory name look-up cache (DNLC) is enhanced to provide improved performance when you access files in large directories with 1000 or more files.

The DNLC is a general file system service that caches the most recently referenced directory names and their associated vnodes. UFS directory entries are stored linearly on disk, which means that locating an entry requires searching each entry for the name. The addition of a new entry requires searching the entire directory to ensure that the name does not exist. To solve this performance problem, the DNLC caches entire directories in memory.

Another feature in this release is the DNLC caching of file objects that have been looked up, but do not exist. This feature is known as *negative caching*, and is useful because some applications repeatedly test to check if a file exists.

For more information, see *Solaris Tunable Parameters Reference Manual*.

## The pargs and preap Commands

Two new commands, `pargs` and `preap`, improve process debugging. The `pargs` command can be used to print the arguments and environment variables associated with a live process or core file. The `preap` command can be used to remove defunct (zombie) processes.

The `pargs` command solves a long-standing problem of being unable to display all the arguments that are passed to a process with the `ps` command. Use the `pargs` command in combination with the `pgrep` command to display the arguments that are passed to a process, as follows:

```
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
```

Use the `pargs -e` command to display the environment variables that are associated with a process, as in the following example:

```
$ pargs -e 6763
6763: tcsh
envp[0]: DISPLAY=:0.0
```

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.

You can use the `preap` command to clean up a defunct (also called a *zombie*) process. A zombie process has not yet had its exit status reaped (or claimed) by its parent. These processes are generally harmless, but can consume system resources if they are numerous.

For information on using the `preap` command, see `preap(1)`. For information on the using the `pargs` command, see `proc(1)`.

## Performance Tool

You can monitor system performance and system resource utilization with the Solaris Management Console's Performance Tool.

For information on allocating, monitoring, and controlling system resources with the Solaris resource management, see "Introduction to Solaris 9 Resource Manager" in *System Administration Guide: Resource Management and Network Services*.

For information on using and starting the Solaris Management Console, see "Working With the Management Console (Tasks)" in *System Administration Guide: Basic Administration*.

## New Fair Share (FSS) and Fixed (FX) Schedulers

The FSS and FX schedulers are new user process schedulers. Both schedulers use the same range of priorities (0 to 59) as the time sharing (TS)/interactive (IA) scheduling classes.

FX is a fixed-priority scheduler that gives processes that run in it an assured priority. FX is not meant to be a default scheduler. TS balances loads. FX does not. You can run FX and TS on the same system, but if you set the priority too high for FX, you could starve the processes that run in TS.

The FSS guarantees application performance by explicitly allocating shares of CPU resources to projects. Additional advantages of the FSS are as follows:

- Allows the fine-grained sharing of CPU resources on a system.
- Controls the allocation of available CPU resources among workloads based on their importance. Their importance is expressed by the number of shares of CPU resources that you assign to each workload.

FSS should not be used with TS or FX on the same system unless processor sets are used. If you use processor sets, you can mix TS, IA, and FX with FSS in one system as long as all the processes that run on each processor set are in one scheduling class to prevent them from competing for the same CPUs. The FX scheduler, in particular, should not be used in conjunction with the new FSS scheduling class unless processor sets are used, to prevent applications in the FX class from using priorities that are high enough to starve applications in the FSS class.

The following table identifies where to get more information about using these schedulers and other Solaris resource management features, such as projects and tasks, to better manage system resources.

| Topic                       | For More Information                                                                                   |
|-----------------------------|--------------------------------------------------------------------------------------------------------|
| Using FSS and FX schedulers | “Fair Share Scheduler” in <i>System Administration Guide: Resource Management and Network Services</i> |
| FSS                         | FSS(7)                                                                                                 |
| FX                          | priocntl(1) and dispadmin(1M)                                                                          |
| Projects                    | “Projects and Tasks” in <i>System Administration Guide: Resource Management and Network Services</i>   |

---

## Where to Find System Performance Tasks

| System Performance Task           | For More Information                               |
|-----------------------------------|----------------------------------------------------|
| Manage processes                  | Chapter 23                                         |
| Monitor system performance        | Chapter 24                                         |
| Change Solaris tunable parameters | <i>Solaris Tunable Parameters Reference Manual</i> |

---

## System Performance and System Resources

The performance of a computer system depends upon how the system uses and allocates its resources. It is important to 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 memory (RAM) on the system.                                                                     |

Chapter 24 describes the tools that display statistics about the activity and the performance of a system.

## Processes and System Performance

The following table describes terms related to processes.

**TABLE 22-1** Process Terminology

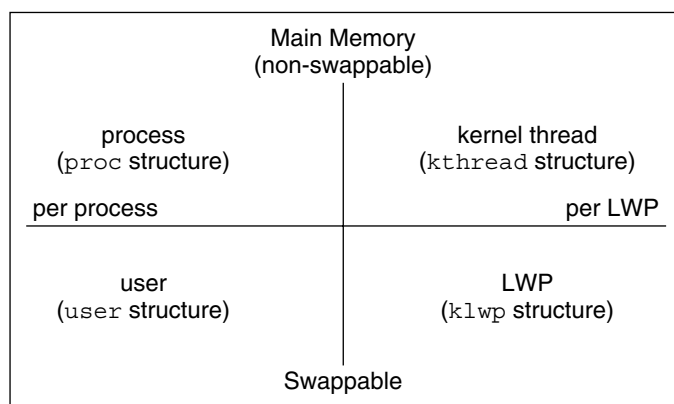
| Term                      | Description                                                                                                                                                                                                                                                                                                            |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. They 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 22-2** Process Structures

| Structure            | Description                                                                                      |
|----------------------|--------------------------------------------------------------------------------------------------|
| <code>proc</code>    | Contains information that pertains to the whole process and must be in main memory all the time. |
| <code>kthread</code> | Contains information that pertains to one LWP and must be in main memory all the time.           |
| <code>user</code>    | Contains the “per process” information that is swappable.                                        |
| <code>klwp</code>    | Contains the “per LWP process” information that is swappable.                                    |

The following figure illustrates the relationship of these process structures.



**FIGURE 22-1** 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 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 22-3** Performance Monitoring Tools

| Command                      | Description                                                                                      | For More Information                                                                                    |
|------------------------------|--------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| sar and sadc commands        | Collects and reports on system activity data                                                     | Chapter 24                                                                                              |
| ps and prstat commands       | Displays information about active processes                                                      | Chapter 23                                                                                              |
| vmstat and iostat commands   | Summarizes system activity data, such as virtual memory statistics, disk usage, and CPU activity | Chapter 24                                                                                              |
| swap command                 | Displays information about available swap space on your system                                   | “Configuring Additional Swap Space (Tasks)” in <i>System Administration Guide: Basic Administration</i> |
| netstat and nfsstat commands | Displays information about network performance                                                   | netstat(1M) and nfsstat(1M)                                                                             |
| Sun Enterprise SyMON         | Collects system activity data on Sun’s Enterprise level systems                                  | <i>Sun Enterprise SyMON 2.0.1 Software User’s Guide</i>                                                 |





---

## Managing System Processes (Tasks)

---

This chapter describes the procedures for managing system processes. This is a list of the step-by-step instructions in this chapter.

- “How to List Processes” on page 307
- “How to Display Information About Processes” on page 310
- “How to Control Processes” on page 311
- “How to Kill a Process” on page 312
- “How to Display Basic Information About Process Classes” on page 314
- “How to Display the Global Priority of a Process” on page 314
- “How to Designate a Process Priority” on page 315
- “How to Change Scheduling Parameters of a Timesharing Process” on page 315
- “How to Change the Class of a Process” on page 316
- “How to Change the Priority of a Process” on page 318

---

## Commands for Managing System Processes (Overview)

The following table describes the commands for managing system processes.

**TABLE 23-1** Commands for Managing Processes

| Command           | Man Page                        | Description                                                                                                     |
|-------------------|---------------------------------|-----------------------------------------------------------------------------------------------------------------|
| ps, pgrep, prstat | ps(1), pgrep(1), and prstat(1M) | Checks the status of active processes on a system, as well as displays detailed information about the processes |

**TABLE 23-1** Commands for Managing Processes (Continued)

| Command   | Man Page      | Description                                                                                   |
|-----------|---------------|-----------------------------------------------------------------------------------------------|
| dispadmin | dispadmin(1M) | Lists default scheduling policies                                                             |
| prionctl  | prionctl(1)   | Assigns processes to a priority class and manages process priorities                          |
| nice      | nice(1)       | Changes the priority of a timesharing process                                                 |
| psrset    | psrset(1M)    | Binds specific process groups to a group of processors rather than to just a single processor |

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 "Working With the Management Console (Tasks)" in *System Administration Guide: Basic Administration*.

## 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 such administrative tasks 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 `ps(1)`.

**TABLE 23–2** 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 utilization for scheduling. This field is not displayed when the <code>-c</code> option is used.                                     |
| CLS   | The scheduling class to which the process belongs: real-time, system, or timesharing. This field is included only with the <code>-c</code> option. |
| PRI   | The kernel thread's scheduling priority. Higher numbers indicate a higher priority.                                                                |
| NI    | The process's <code>nice</code> number, which contributes to its scheduling priority. Making a process "nicer" means lowering its priority.        |
| ADDR  | The address of the <code>proc</code> 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.                                                                                                            |

## ▼ How to List Processes

You can use the `ps` command to list all the processes on a system.

```
$ ps [-efc]
```

|                  |                                                                                          |
|------------------|------------------------------------------------------------------------------------------|
| <code>ps</code>  | Displays only the processes that are associated with your login session.                 |
| <code>-ef</code> | Displays full information about all the processes that are being executed on the system. |
| <code>-c</code>  | Displays process scheduler information.                                                  |

## Example—Listing Processes

The following example shows output from the `ps` command when no options are used.

```
$ ps
 PID TTY TIME CMD
 1664 pts/4 0:06 csh
 2081 pts/4 0:00 ps
```

The following example shows output from `ps -ef`. 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
 .
 .
 .
```

## The /proc File System and Commands

Process tools that are available in the `/usr/proc/bin` directory display highly detailed information about the processes listed in the `/proc` directory. (The `/proc` directory is also known as the process file system (PROCFS).) Images of active processes are stored here by their process ID number.

The process tools are similar to some options of the `ps` command, except that the output that is provided by the tools is more detailed. In general, the process tools 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 `/proc` Tools

You can display detailed, technical information about or control active processes by using some of the process tool commands that are contained in the `/usr/proc/bin` directory. Table 23-3 lists some of the `/proc` tools.

If a process becomes trapped in an endless loop, or if it takes too long to execute, you might want to stop (kill) the process. For more information about stopping processes using the `pkill` command, see Chapter 23.

The `/proc` file system is a directory hierarchy that contains additional subdirectories for state information and control functions.

It also provides a watchpoint 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 watchpoint facility, which means that the entire watchpoint process is faster.

The following restrictions have been removed when setting watchpoints by using the `dbx` debugging tool:

- Setting watchpoints on local variables on the stack due to SPARC register windows
- Setting watchpoints on multithreaded processes

For more information, see `proc(4)`, `core(4)`, and `mdb(1)`.

**TABLE 23-3** `/usr/proc/bin` Tools

| Process Tool        | Description                                                                                     |
|---------------------|-------------------------------------------------------------------------------------------------|
| <code>pcred</code>  | Displays process credential information                                                         |
| <code>pfiles</code> | Reports <code>fstat</code> and <code>fcntl</code> information for open files in a process       |
| <code>pflags</code> | Prints <code>/proc</code> tracing flags, pending and held signals, and other status information |

**TABLE 23-3** /usr/proc/bin Tools (Continued)

| Process Tool | Description                                                    |
|--------------|----------------------------------------------------------------|
| 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)`.

---

**Note** – To avoid typing long command names, add the process tool directory to your `PATH` variable. Then, you can run process tools by entering only the last part of each file name (for example, `pwdx` instead of `/usr/proc/bin/pwdx`).

---

## ▼ How to Display Information About Processes

1. Obtain the process ID of the process you want to display more information about.

```
pgrep process
```

*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 you need.

```
/usr/proc/bin/pcommand pid
```

*pcommand* Process tool command that you want to run. Table 23-3 lists these commands.

*pid* Indicates the process ID.

## Example—Displaying Information About Processes

The following example shows how to use process tool commands to display more information about an `lpsched` process. First, the `/usr/proc/bin` path is defined to avoid typing long commands. Next, the identification number for the `lpsched` process is obtained. Finally, output from three process tool commands is shown.

```
PATH=$PATH:/usr/proc/bin
export PATH 1
pgrep lpsched 2
213
pwdx 213 3
213: /
ptree 213 4
213 /usr/lib/lpsched
pfiles 213 5
213: /usr/lib/lpsched
Current rlimit: 4096 file descriptors
0: S_IFIFO mode:0000 dev:270,0 ino:67 uid:0 gid:0 size:0
 O_RDWR
1: S_IFIFO mode:0000 dev:270,0 ino:67 uid:0 gid:0 size:0
 O_RDWR
3: S_IFCHR mode:0666 dev:136,0 ino:35882 uid:0 gid:3 rdev:21,0
 O_WRONLY FD_CLOEXEC
4: S_IFDOOR mode:0444 dev:275,0 ino:18526 uid:0 gid:0 size:0
 O_RDONLY|O_LARGEFILE FD_CLOEXEC door to nscd[208]
5: S_IFREG mode:0664 dev:136,0 ino:64648 uid:71 gid:8 size:0
 O_WRONLY
```

1. Adds the `/usr/proc/bin` directory to the `PATH` variable
2. Obtains the process identification number for `lpsched`
3. Displays the current working directory for `lpsched`
4. Displays the process tree that contains `lpsched`
5. Displays `fstat` and `fcntl` information

## ▼ How to Control Processes

1. Obtain the process ID of the process you want to control.

```
pgrep process
```

*process* is the name of the process you want to control.

The process identification number is in the first column of the output.

2. Use the appropriate `/usr/proc/bin` command to control the process.

```
/usr/proc/bin/pcommand pid
```

|                 |                                                                        |
|-----------------|------------------------------------------------------------------------|
| <i>pcommand</i> | Process tool command you want to run. Table 23–3 lists these commands. |
| <i>pid</i>      | Identifies the process ID.                                             |

### 3. Verify the process status.

```
pgrep PID
```

## Example—Controlling Processes

The following example shows how to use process tools to stop and restart `dtpad`.

```
PATH=$PATH:/usr/proc/bin
export PATH 1
pgrep dtpad 2
2921
psstop 2921 3
prun 2921 4
```

1. Adds the `/usr/proc/bin` directory to the `PATH` variable
2. Obtains the process identification number for `dtpad`
3. Stops the `dtpad` process
4. Restarts the `dtpad` process

## Killing a Process (`kill`)

Sometimes, it is necessary 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, and 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 might crash the system.

For more information, see `pgrep(1)`.

### ▼ How to Kill a Process

1. **(Optional) To kill a process that belongs to another user, become superuser.**
2. **Obtain the of the process ID of the process you want to stop.**

```
$ pgrep process
```

*process* is the name of the process you want to display more information about.

The process identification number is displayed in the first column of the output.

3. **Stop the process.**



```
$ pkill [-9] pid
```

|     |                                               |
|-----|-----------------------------------------------|
| -9  | Ensures that the process terminates promptly. |
| pid | Process ID to stop.                           |

#### 4. Verify that the process has been stopped.

```
$ pgrep pid
```

---

## Managing Process Class Information

The following list identifies the process scheduling classes that can be configured on your system, and 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 and 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 With `priocntl`

The scheduling priority of a process is the priority it is assigned by the process scheduler, according to scheduling policies. The `dispadm` command lists the default scheduling policies.

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” on page 315.

## ▼ How to Display Basic Information About Process Classes

You can display process scheduling classes and priority ranges with the `prionctl -l` command.

```
$ prionctl -l
```

### Example—Getting Basic Information About Process Classes

The following example shows output from the `prionctl -l` command.

```
prionctl -l
CONFIGURED CLASSES
=====

SYS (System Class)

TS (Time Sharing)
 Configured TS User Priority Range: -60 through 60

IA (Interactive)
 Configured IA User Priority Range: -60 through 60

RT (Real Time)
 Maximum Configured RT Priority: 59
```

## ▼ How to Display the Global Priority of a Process

You can display the global priority of a process by using the `ps` command.

```
$ ps -ecl
```

The global priority is listed under the `PRI` column.

### Example—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 `sh` has the lowest.

```
$ ps -ecl
 F S UID PID PPID CLS PRI ADDR SZ WCHAN TY 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

1. Become superuser.
2. Start a process with a designated priority.

```
priocntl -e -c class -m userlimit -p pri command-name
```

|                                  |                                                                                                                                                                                                                  |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-e</code>                  | Executes the command.                                                                                                                                                                                            |
| <code>-c class</code>            | Specifies the class within which to run the process. The valid classes are TS (timesharing), RT (real time), IA (interactive), FSS (fair share), or FX (fixed priority).                                         |
| <code>-m userlimit</code>        | Specifies the maximum amount you can raise or lower your priority, when using the <code>-p</code> option.                                                                                                        |
| <code>-p pri command-name</code> | Lets you specify the relative priority in the RT class, for a real-time thread. For a timesharing process, the <code>-p</code> 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—Designating a Process Priority

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

1. Become superuser.
2. Change the scheduling parameters of a running timesharing process.

```
priocntl -s -m userlimit [-p userpriority] -i idtype idlist
```

|                               |                                                                                                                                                                                                                                                 |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-s</code>               | Lets you set the upper limit on the user priority range and change the current priority.                                                                                                                                                        |
| <code>-m userlimit</code>     | Specifies the maximum amount you can raise or lower your priority, when you use the <code>-p</code> option.                                                                                                                                     |
| <code>-p userpriority</code>  | Allows you to designate a priority.                                                                                                                                                                                                             |
| <code>-i idtype idlist</code> | Uses a combination of <i>idtype</i> and <i>idlist</i> to identify the process or processes. The <i>idtype</i> specifies the type of ID, such as <i>pid</i> or <i>UID</i> . Use <i>idlist</i> to identify a list of <i>pids</i> or <i>UIDs</i> . |

### 3. Verify the process status.

```
ps -ecl | grep idlist
```

## Example—Changing Scheduling Parameters of a Timesharing Process

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 -t 500 -p 20 myprog
ps -ecl | grep myprog
```

## ▼ How to Change the Class of a Process

### 1. (Optional) Become superuser.

---

**Note** – You must be superuser or working in a real-time shell to change a process from, or to, a real-time process.

---

### 2. Change the class of a process.

```
priocntl -s -c class -i idtype idlist
```

|                       |                                                                                          |
|-----------------------|------------------------------------------------------------------------------------------|
| <code>-s</code>       | Lets you set the upper limit on the user priority range and change the current priority. |
| <code>-c class</code> | Specifies the class, TS or RT, to which you are changing the process.                    |

`-i idtype idlist` Uses a combination of *idtype* and *idlist* to identify the process or processes. The *idtype* specifies the type of ID, such as *pid* or *UID*. Use *idlist* to identify a list of *pids* or *UIDs*.

### 3. Verify the process status.

```
ps -ecl | grep idlist
```

## Example—Changing the Class of a Process

The following example shows how to change all the processes belonging to user 15249 to real-time processes.

```
priocntl -s -c RT -i uid 15249
ps -ecl | grep 15249
```

---

**Note** – 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.

---

## Changing the Priority of a Timesharing Process With the `nice` Command

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, and is determined by several factors, including its 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. But 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 between 0 and +40, with 0 representing the highest priority. The default value is 20. Two versions of the command are available, the standard version, `/usr/bin/nice`, and a version that is part of the C shell.

## ▼ How to Change the Priority of a Process

You can raise or lower the priority of a command or a process by changing the `nice` number. To lower the priority of a process:

|                                             |                                                                 |
|---------------------------------------------|-----------------------------------------------------------------|
| <code>/usr/bin/nice command-name</code>     | Increase the <code>nice</code> number by 4 units (the default). |
| <code>/usr/bin/nice +4 command-name</code>  | Increase the <code>nice</code> number by 4 units.               |
| <code>/usr/bin/nice -10 command-name</code> | Increase the <code>nice</code> number by 10 units.              |

The first and second commands increase the `nice` number by 4 units (the default). The third command increases the `nice` number by 10 units. All three commands lower the priority of the process.

The following commands raise the priority of the command by lowering the `nice` number.

To raise the priority of a process, do one of the following:

|                                              |                                                                                                                                                                                      |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>/usr/bin/nice -10 command-name</code>  | Raises the priority of the command by lowering the <code>nice</code> number by 10 units.                                                                                             |
| <code>/usr/bin/nice --10 command-name</code> | Raises the priority of the command by lowering the <code>nice</code> number by 10 units. The first minus sign is the option sign. The second minus sign indicates a negative number. |

## Troubleshooting Problems With System Processes

Here are some tips on obvious problems you might find:

- 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 see 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.





## Monitoring System Performance (Tasks)

---

This chapter describes procedures for monitoring system performance by using the `vmstat`, `iostat`, `df`, and `sar` commands. This is a list of the step-by-step instructions in this chapter.

- “How to Display Virtual Memory Statistics (`vmstat`)” on page 323
- “How to Display System Event Information (`vmstat -s`)” on page 324
- “How to Display Swapping Statistics (`vmstat -S`)” on page 325
- “How to Display Cache Flushing Statistics (`vmstat -c`)” on page 325
- “How to Display Interrupts Per Device (`vmstat -i`)” on page 326
- “How to Display Disk Utilization Information (`iostat`)” on page 327
- “How to Display Extended Disk Statistics (`iostat -xtc`)” on page 328
- “How to Display Disk Space Information (`df`)” on page 329
- “How to Check File Access (`sar -a`)” on page 330
- “How to Check Buffer Activity (`sar -b`)” on page 331
- “How to Check System Call Statistics (`sar -c`)” on page 332
- “How to Check Disk Activity (`sar -d`)” on page 333
- “How to Check Page-Out and Memory (`sar -g`)” on page 335
- “How to Check Kernel Memory Allocation (`sar -k`)” on page 336
- “How to Check Interprocess Communication (`sar -m`)” on page 337
- “How to Check Page-In Activity (`sar -p`)” on page 338
- “How to Check Queue Activity (`sar -q`)” on page 340
- “How to Check Unused Memory (`sar -r`)” on page 341
- “How to Check CPU Utilization (`sar -u`)” on page 341
- “How to Check System Table Status (`sar -v`)” on page 342
- “How to Check Swapping Activity (`sar -w`)” on page 343
- “How to Check Terminal Activity (`sar -y`)” on page 345
- “How to Check Overall System Performance (`sar -A`)” on page 346
- “How to Set Up Automatic Data Collection” on page 349

---

## What's New in Monitoring System Performance?

The `vmstat` command now includes all system interrupts in the `faults-in` column.

In previous Solaris releases, this column did not report clock device interrupts.

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

---

## Displaying Virtual Memory Statistics (`vmstat`)

You can use the `vmstat` command to report virtual memory statistics and such information about system events 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` output.

**TABLE 24-1** 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                |
| memory   | w          | The number of swapped out LWPs that are waiting for processing resources to finish |
|          |            | Reports on usage of real memory and virtual memory:                                |
|          | swap       | Available swap space                                                               |
| page     | free       | Size of the free list                                                              |
|          |            | Reports on page faults and paging activity, in units per second:                   |
|          | re         | Pages reclaimed                                                                    |

**TABLE 24-1** Output From the `vmstat` Command (Continued)

| Category            | Field Name      | Description                                                                                                                        |
|---------------------|-----------------|------------------------------------------------------------------------------------------------------------------------------------|
|                     | <code>mf</code> | Minor and major faults                                                                                                             |
|                     | <code>pi</code> | Kbytes paged in                                                                                                                    |
|                     | <code>po</code> | Kbytes paged out                                                                                                                   |
|                     | <code>fr</code> | Kbytes freed                                                                                                                       |
|                     | <code>de</code> | Anticipated memory that is needed by recently swapped-in processes                                                                 |
|                     | <code>sr</code> | Pages scanned by the page daemon (not currently in use). If <code>sr</code> does not equal zero, the page daemon has been running. |
| <code>disk</code>   |                 | Reports the number of disk operations per second, showing data on up to four disks                                                 |
| <code>faults</code> |                 | Reports the trap/interrupt rates (per second):                                                                                     |
|                     | <code>in</code> | Interrupts per second                                                                                                              |
|                     | <code>sy</code> | System calls per second                                                                                                            |
|                     | <code>cs</code> | CPU context switch rate                                                                                                            |
| <code>cpu</code>    |                 | Reports on the use of CPU time:                                                                                                    |
|                     | <code>us</code> | User time                                                                                                                          |
|                     | <code>sy</code> | System time                                                                                                                        |
|                     | <code>id</code> | Idle time                                                                                                                          |

For a more detailed description of this command, see `vmstat(1M)`.

## ▼ How to Display Virtual Memory Statistics (`vmstat`)

Collect virtual memory statistics by using the `vmstat` command with a time interval in seconds.

```
$ vmstat n
```

*n* is the interval in seconds between reports.

## Example—Displaying Virtual Memory Statistics

The following example shows the `vmstat` display of statistics gathered at five-second intervals.

```
$ vmstat 5
procs memory page disk faults cpu
r b w swap free re mf pi po fr de sr f0 s3 -- -- in sy cs us sy id
0 0 8 28312 668 0 9 2 0 1 0 0 0 1 0 0 10 61 82 1 2 97
0 0 3 31940 248 0 10 20 0 26 0 27 0 4 0 0 53 189 191 6 6 88
0 0 3 32080 288 3 19 49 6 26 0 15 0 9 0 0 75 415 277 6 15 79
0 0 3 32080 256 0 26 20 6 21 0 12 1 6 0 0 163 110 138 1 3 96
0 1 3 32060 256 3 45 52 28 61 0 27 5 12 0 0 195 191 223 7 11 82
0 0 3 32056 260 0 1 0 0 0 0 0 0 0 0 0 4 52 84 0 1 99
```

## ▼ How to Display System Event Information (`vmstat -s`)

Run the `vmstat -s` command to show the total of various system events that 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
392182 total address trans. faults taken
20419 page ins
 923 page outs
30072 pages paged in
 9194 pages paged out
65167 total reclaims
65157 reclaims from free list
 0 micro (hat) faults
392182 minor (as) faults
19383 major faults
85775 copy-on-write faults
66637 zero fill page faults
46309 pages examined by the clock daemon
 6 revolutions of the clock hand
15578 pages freed by the clock daemon
 4398 forks
 352 vforks
 4267 execs
12926285 cpu context switches
109029866 device interrupts
 499296 traps
22461261 system calls
 778068 total name lookups (cache hits 97%)
 18739 user cpu
 34662 system cpu
```

```
52051435 idle cpu
25252 wait cpu
```

## ▼ How to Display Swapping Statistics (`vmstat -S`)

Run `vmstat -S` to show swapping statistics.

```
$ vmstat -S
procs memory page disk faults cpu
r b w swap free si so pi po fr de sr f0 s0 s6 -- in sy cs us sy id
0 0 0 200968 17936 0 0 0 0 0 0 0 0 0 0 0 109 43 24 0 0 100
```

The swapping statistics fields are described in the following table. For a description of the other fields, see Table 24-1.

**TABLE 24-2** Output From the `vmstat -S` Command

| Field Name | Description                                           |
|------------|-------------------------------------------------------|
| 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 both fields. Use the `sar` command to display a more accurate accounting of swap statistics.

## ▼ How to Display Cache Flushing Statistics (`vmstat -c`)

Run the `vmstat -c` command to show cache flushing statistics for a virtual cache.

```
$ vmstat -c
usr ctx rgn seg pag par
0 60714 5 134584 4486560 4718054
```

The output shows the total number of cache flushes since the last boot. The cache types are described in the following table.

**TABLE 24-3** Output From the `vmstat -c` Command

| Cache Name | Cache Type |
|------------|------------|
| usr        | User       |

**TABLE 24-3** Output From the `vmstat -c` Command (Continued)

| Cache Name | Cache Type   |
|------------|--------------|
| ctx        | Context      |
| rgn        | Region       |
| seg        | Segment      |
| pag        | Page         |
| par        | Partial-page |

## ▼ How to Display Interrupts Per Device (`vmstat -i`)

Run the `vmstat -i` command to show the number of interrupts per device.

```
$ vmstat -i
```

### Example—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 n`)

Use the `iostat` command to report statistics about disk input and output, and produces measures of throughput, utilization, queue lengths, transaction rates, and service time. For a detailed description of this command, refer to `iostat(1M)`.

## ▼ How to Display Disk Utilization Information (iostat)

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`).

The following table describes the fields in the `iostat` command output.

**TABLE 24-4** Output From the `iostat n` Command

| Device Type | Field Name        | Description                                       |
|-------------|-------------------|---------------------------------------------------|
| Terminal    |                   |                                                   |
|             | <code>tin</code>  | Number of characters in the terminal input queue  |
|             | <code>tout</code> | Number of characters in the terminal output queue |
| Disk        |                   |                                                   |
|             | <code>bps</code>  | Blocks per second                                 |
|             | <code>tps</code>  | Transactions per second                           |
|             | <code>serv</code> | Average service time, in milliseconds             |
| CPU         |                   |                                                   |
|             | <code>us</code>   | In user mode                                      |
|             | <code>sy</code>   | In system mode                                    |
|             | <code>wt</code>   | Waiting for I/O                                   |
|             | <code>id</code>   | Idle                                              |

## Example—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 0 100
0 16 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 0 100
0 16 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 0 100
0 16 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 0 100
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 0 100
0 16 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 0 100

```

## ▼ How to Display Extended Disk Statistics (`iostat -xtc`)

Run the `iostat -xtc` command to get extended disk statistics.

```

$ iostat -xtc
 extended device statistics
device r/s w/s kr/s kw/s wait actv svc_t %w %b tty cpu
 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 table.

**TABLE 24-5** Output From the `iostat -xtc` Command

| Field Name         | Description                                                                |
|--------------------|----------------------------------------------------------------------------|
| <code>r/s</code>   | Reads per second                                                           |
| <code>w/s</code>   | Writes per second                                                          |
| <code>kr/s</code>  | Kbytes read per second                                                     |
| <code>kw/s</code>  | Kbytes written per second                                                  |
| <code>wait</code>  | Average number of transactions that are waiting for service (queue length) |
| <code>actv</code>  | Average number of transactions that are actively being serviced            |
| <code>svc_t</code> | Average service time, in milliseconds                                      |



**TABLE 24-5** Output From the `iostat -xtc` Command (Continued)

| Field Name | Description                                    |
|------------|------------------------------------------------|
| %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 leave a 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, or to a tape by using the `tar` or `cpio` commands. Or, you could remove the files.

For a detailed description of this command, see `df(1M)`.

### ▼ How to Display Disk Space Information (df)

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% /
```

The following table describes the `df -k` command output.

**TABLE 24-6** Output From 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 |

**TABLE 24-6** Output From the `df -k` Command (Continued)

| Field Name | Description |
|------------|-------------|
| mounted on | Mount point |

## Example—Displaying File System Information

The following example shows the `df -k` command output.

```
$ df -k
Filesystem kbytes used avail capacity Mounted on
/dev/dsk/c0t0d0s0 384120 131596 214112 39% /
/dev/dsk/c0t0d0s6 1388419 1050390 282493 79% /usr
/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
```

---

## Monitoring System Activities (sar)

Use the `sar` command to do the following:

- 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, and special request reports to pinpoint specific performance problems. “Collecting System Activity Data Automatically (sar)” on page 346 describes these tools.

For a detailed description of this command, see `sar(1)`.

### ▼ How to Check File Access (sar -a)

Display file access operation statistics with the `sar -a` command.

```
$ sar -a
SunOS venus 5.9 Generic sun4u 06/24/2001

00:00:00 iget/s namei/s dirbk/s
01:00:00 0 0 0
02:00:02 0 0 0
03:00:00 0 1 0
```

```

04:00:00 0 0 0
05:00:01 0 0 0
06:00:00 0 0 0

Average 0 1 0

```

The following table describes the operating system routines that are reported by the `sar -a` command.

**TABLE 24-7** Output from the `sar -a` Command

| Field Name           | Description                                                                                                                                                                                                                                                     |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>iget/s</code>  | The number of requests made for inodes that were not in the directory name look-up cache (DNLC).                                                                                                                                                                |
| <code>namei/s</code> | The number of file system path searches per second. If <code>namei</code> does not find a directory name in the DNLC, it calls <code>iget</code> to get the inode for either a file or directory. Hence, most <code>igets</code> are the result of DNLC misses. |
| <code>dirbk/s</code> | The number of directory block reads issued per second.                                                                                                                                                                                                          |

The larger the reported values, 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`)

Display buffer activity statistics with the `sar -b` command.

The buffer is used to cache metadata, which 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

```

The following table describes the buffer activities that are displayed by the `-b` option.

**TABLE 24-8** Output From the `sar -b` Command

| Field Name           | Description                                                                             |
|----------------------|-----------------------------------------------------------------------------------------|
| <code>bread/s</code> | Average number of reads per second that are submitted to the buffer cache from the disk |

**TABLE 24-8** Output From the `sar -b` Command (Continued)

| Field Name           | Description                                                                                                                           |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <code>lread/s</code> | Average number of logical reads per second from the buffer cache                                                                      |
| <code>%rcache</code> | Fraction of logical reads that are found in the buffer cache (100% minus the ratio of <code>bread/s</code> to <code>lread/s</code> )  |
| <code>bwrit/s</code> | Average number of physical blocks (512 blocks) that are written from the buffer cache to disk, per second                             |
| <code>lwrit/s</code> | Average number of logical writes to the buffer cache, per second                                                                      |
| <code>%wcache</code> | Fraction of logical writes that are found in the buffer cache (100% minus the ratio of <code>bwrit/s</code> to <code>lwrit/s</code> ) |
| <code>pread/s</code> | Average number of physical reads, per second that use character device interfaces                                                     |
| <code>pwrit/s</code> | 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`, which 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.

## Example—Checking Buffer Activity

The following abbreviated example of `sar -b` output shows that the `%rcache` and `%wcache` buffers are not causing any slowdowns. All the data is within acceptable limits.

```
$ sar -b
SunOS venus 5.9 Generic sun4u 06/24/2001

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
02:00:02 0 0 100 0 0 55 0 0
03:00:00 0 0 100 0 0 72 0 0
04:00:00 0 0 100 0 0 56 0 0
05:00:01 0 0 100 0 0 55 0 0
06:00:00 0 0 100 0 0 55 0 0

Average 0 0 94 0 0 64 0 0
```

## ▼ How to Check System Call Statistics (`sar -c`)

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

```

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, although the percentage varies greatly with the activities that are being performed by the system.

**TABLE 24-9** Output From the `sar -c` Command

| Field Name           | Description                                                                                                                                                               |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>scall/s</code> | All types of system calls per second (generally about 30 per second on a system with 4 to 6 users).                                                                       |
| <code>sread/s</code> | read system calls per second.                                                                                                                                             |
| <code>swrit/s</code> | write system calls per second.                                                                                                                                            |
| <code>fork/s</code>  | fork system calls per second (about 0.5 per second on a system with 4 to 6 users). This number will increase if shell scripts are running.                                |
| <code>exec/s</code>  | <code>exec</code> system calls per second. If <code>exec/s</code> divided by <code>fork/s</code> is greater than three, look for inefficient <code>PATH</code> variables. |
| <code>rchar/s</code> | Characters (bytes) transferred by <code>read</code> system calls per second.                                                                                              |
| <code>wchar/s</code> | Characters (bytes) transferred by <code>write</code> system calls per second.                                                                                             |

## Example—Checking System Call Statistics

The following abbreviated example shows output from the `sar -c` command.

```

$ sar -c
SunOS venus 5.9 Generic sun4u 06/24/2001

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
02:00:02 38 2 2 0.00 0.00 149 120
03:00:00 42 2 2 0.05 0.05 218 147
04:00:00 39 2 2 0.01 0.00 155 123
05:00:01 38 2 2 0.00 0.00 150 120
06:00:00 38 2 2 0.01 0.00 149 120

Average 50 4 3 0.02 0.02 532 238

```

## ▼ How to Check Disk Activity (`sar -d`)

Display disk activity statistics with the `sar -d` command.

```

$ sar -d
00:00:00 device %busy avque r+w/s blks/s await avserv

01:00:00 fd0 0 0.0 0 0 0.0 0.0

```

The following table describes the disk device activities that are reported by the `-d` option.

**TABLE 24-10** Output From the `sar -d` Command

| 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                                                                                                           |
| blks/s     | Number of 512-byte blocks that are transferred to the device, per second                                                                                               |
| await      | Average time, in milliseconds, that transfer requests wait idly in the queue (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, rotational latency, and data transfer times) |

Note that queue lengths and wait times are measured when there is something 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.

## Examples—Checking Disk Activity

This abbreviated example illustrates the `sar -d` output.

```

$ sar -d
SunOS venus 5.9 Generic sun4u 06/24/2001

00:00:00 device %busy avque r+w/s blks/s await avserv

01:00:00 fd0 0 0.0 0 0 0.0 0.0
 nfs1 0 0.0 0 0 0.0 0.0
 sd0 0 0.0 0 0 0.0 39.6
 sd0,a 0 0.0 0 0 0.0 39.6
 sd0,b 0 0.0 0 0 0.0 0.0
 sd0,c 0 0.0 0 0 0.0 0.0
 sd0,f 0 0.0 0 0 0.0 0.0
 sd0,g 0 0.0 0 0 0.0 0.0

```

```

sd0,h 0 0.0 0 0 0.0 0.0
sd6 0 0.0 0 0 0.0 0.0

```

## ▼ How to Check Page-Out and Memory (sar -g)

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 may 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 `pgfree/s` and `pgscan/s`, indicates a memory shortage.

The `sar -g` command also shows whether inodes are being recycled too quickly and causing a loss of reusable pages.

The following table describes the output from the `-g` option.

**TABLE 24-11** Output From the `sar -g` Command

| Field Name            | Description                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>pgout/s</code>  | The number of page-out requests per second.                                                                                                                                                                                                                                                                                                                                                                                      |
| <code>ppgout/s</code> | The actual number of pages that are paged-out, per second. A single page-out request might involve paging-out multiple pages.                                                                                                                                                                                                                                                                                                    |
| <code>pgfree/s</code> | The number of pages, per second, that are placed on the free list.                                                                                                                                                                                                                                                                                                                                                               |
| <code>pgscan/s</code> | 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.                                                                                                                                                                                                        |
| <code>%ufs_ipf</code> | The percentage of <code>ufs</code> inodes taken off the free list by <code>iget</code> that had reusable pages associated with them. These pages are flushed and cannot be reclaimed by processes. Thus, this field represents the percentage of <code>iget</code> s with page flushes. A high value indicates that the free list of inodes is page-bound, and the number of <code>ufs</code> inodes might need to be increased. |

## Example—Checking Page-Out and Memory

The following abbreviated example shows output from the `sar -g` command.

```

$ sar -g
SunOS venus 5.9 Generic sun4u 06/24/2001

```

|          |         |          |          |          |          |
|----------|---------|----------|----------|----------|----------|
| 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     |
| 02:00:02 | 0.00    | 0.00     | 0.00     | 0.00     | 0.00     |
| 03:00:00 | 0.00    | 0.01     | 0.01     | 0.00     | 0.00     |
| 04:00:00 | 0.00    | 0.00     | 0.00     | 0.00     | 0.00     |
| 05:00:01 | 0.00    | 0.00     | 0.00     | 0.00     | 0.00     |
| 06:00:00 | 0.00    | 0.00     | 0.00     | 0.00     | 0.00     |
| Average  | 0.01    | 0.12     | 0.21     | 0.66     | 0.00     |

## ▼ How to Check Kernel Memory Allocation (sar -k)

Use the `sar -k` command to report on the following activities of the Kernel Memory Allocator (KMA).

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), and oversized (greater than 4 Kbytes). The KMA keeps two pools of memory to satisfy small and large requests. The oversized requests are satisfied by allocating memory from the system page allocator.

If you are investigating 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.

```
$ sar -k
00:00:00 sml_mem alloc fail lg_mem alloc fail ovpsz_alloc fail
01:00:00 2523136 1866512 0 18939904 14762364 0 360448 0
02:00:02 2523136 1861724 0 18939904 14778748 0 360448 0
```

The following table describes the output from the `-k` option.



**TABLE 24–12** Output From the `sar -k` Command

| Field Name              | Description                                                                                                                                                                                |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>sml_mem</code>    | The amount of memory, in bytes, that the KMA has available in the small memory request pool (a small request is less than 256 bytes).                                                      |
| <code>alloc</code>      | The amount of memory, in bytes, that the KMA has allocated from its small memory request pool to small memory requests.                                                                    |
| <code>fail</code>       | The number of requests for small amounts of memory that failed.                                                                                                                            |
| <code>lg_mem</code>     | The amount of memory, in bytes, that the KMA has available in the large memory request pool (a large request is from 512 bytes to 4 Kbytes).                                               |
| <code>alloc</code>      | The amount of memory, in bytes, that the KMA has allocated from its large memory request pool to large memory requests.                                                                    |
| <code>fail</code>       | The number of failed requests for large amounts of memory                                                                                                                                  |
| <code>ovsz_alloc</code> | The amount of memory that is allocated for oversized requests (those requests that are greater than 4 Kbytes). These requests are satisfied by the page allocator. Thus, there is no pool. |
| <code>fail</code>       | The number of failed requests for oversized amounts of memory.                                                                                                                             |

## Example—Checking Kernel Memory Allocation (`sar`)

The following is an abbreviated example of `sar -k` output.

```
$ sar -k
SunOS venus 5.9 Generic sun4u 06/24/2001

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
03:00:00 2523136 1865664 0 18939904 14745884 0 360448 0
04:00:00 2523136 1867692 0 18939904 14746616 0 360448 0
05:00:01 2523136 1867208 0 18939904 14763700 0 360448 0
06:00:00 2523136 1867772 0 18939904 14779444 0 360448 0

Average 2724096 1791806 0 20089344 15434591 0 360448 0
```

## ▼ How to Check Interprocess Communication (`sar -m`)

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 will usually be zero (0.00), unless you are running applications that use messages or semaphores.

The following table describes the output from the `-m` option.

**TABLE 24-13** Output From the `sar -m` Command

| Field Name | Description                                                      |
|------------|------------------------------------------------------------------|
| msg/s      | The number of message operations (sends and receives) per second |
| sema/s     | The number of semaphore operations per second                    |

## Example—Checking Interprocess Communication

The following abbreviated example shows output from the `sar -m` command.

```
$ sar -m
SunOS venus 5.9 Generic sun4u 06/24/2001

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`)

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
```

The following table describes the reported statistics from the `-p` option.

**TABLE 24-14** Output From the `sar -p` Command

| 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 (for example, 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 <code>slock/s</code> ), 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 are illegal access to a page and "copy-on-writes." Generally, this number consists primarily of "copy-on-writes."                                                                                                                                                  |
| vflt/s     | The number of address translation page faults, per second. These faults are known as validity faults, and 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 it cannot be claimed and used by another process.                                   |

## Example—Checking Page-In Activity

The following abbreviated example shows output from the `sar -p` command.

```
$ sar -p
SunOS venus 5.9 Generic sun4u 06/24/2001

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
02:00:02 0.07 0.00 0.00 0.21 0.39 0.00
03:00:00 0.32 0.00 0.00 1.10 2.48 0.00
04:00:00 0.09 0.00 0.00 0.32 0.57 0.00
05:00:01 0.07 0.00 0.00 0.21 0.39 0.00
06:00:00 0.07 0.00 0.00 0.21 0.39 0.00

Average 0.26 0.20 0.30 0.92 1.78 0.00
```

## ▼ How to Check Queue Activity (sar -q)

Use the `sar -q` command to report the average queue length while the queue is occupied, and the percentage of time that the queue is occupied.

```
$ sar -q
00:00:00 runq-sz %runocc swpq-sz %swpocc
```

---

**Note** – The number of LWPs swapped out might be greater than zero even if the system has an abundance of free memory. This situation happens when a sleeping LWP is swapped out and has not been awakened (for example, a process or LWP is sleeping, waiting for keyboard or mouse input).

---

The following table describes the output from the `-q` option.

**TABLE 24-15** Output From the `sar -q` Command

| Field Name           | Description                                                                                                                                                                                |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>runq-sz</code> | 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. |
| <code>%runocc</code> | The percentage of time that the dispatch queues are occupied.                                                                                                                              |
| <code>swpq-sz</code> | The average number of swapped out LWPs.                                                                                                                                                    |
| <code>%swpocc</code> | The percentage of time LWPs are swapped out.                                                                                                                                               |

## Example—Checking Queue Activity

The following abbreviated example shows output from the `sar -q` command. If `%runocc` is high (greater than 90 percent) and `runq-sz` 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 venus 5.9 Generic sun4u 06/24/2001

08:45:18 runq-sz %runocc swpq-sz %swpocc
08:45:18 unix restarts
09:00:00 1.0 0 0.0 0
09:20:00 0.0 0 0.0 0
09:40:00 0.0 0 0.0 0

Average 1.0 0 0.0 0
```

## ▼ How to Check Unused Memory (sar -r)

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 table describes the output from the `-r` option.

**TABLE 24-16** Output From the `sar -r` Command

| Field Name | Description                                                                                                                                        |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| 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—Checking Unused Memory

The following example shows output from the `sar -r` command.

```
$ sar -r
SunOS venus 5.9 Generic sun4u 06/24/2001

00:00:00 freemem freeswap
01:00:00 2135 401922
02:00:02 2137 401949
03:00:00 2137 402006
04:00:00 2139 401923
05:00:01 2138 402033
06:00:00 2137 401919

Average 2500 399914
```

## ▼ How to Check CPU Utilization (sar -u)

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 `sar -u`. 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 table describes output from the `-u` option.

**TABLE 24-17** Output From the `sar -u` Command

| Field Name | Description                                                                            |
|------------|----------------------------------------------------------------------------------------|
| %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 generally means that a disk slowdown has occurred.

## Example—Checking CPU Utilization

The following example shows output from the `sar -u` command.

```
$ sar -u
SunOS venus 5.9 Generic sun4u 06/24/2001

00:00:00 %usr %sys %wio %idle
01:00:00 0 0 0 100
02:00:02 0 0 0 100
03:00:00 0 0 0 100
04:00:00 0 0 0 100
05:00:01 0 0 0 100
06:00:00 0 0 0 100
07:00:00 0 0 0 100
08:00:01 0 0 0 100
08:20:00 0 0 0 100
08:40:00 0 0 0 100
09:00:00 0 0 0 100
09:20:00 0 0 0 100
09:40:00 0 0 0 100
10:00:00 0 0 0 100
10:20:00 0 0 0 100
10:40:01 0 0 0 100
11:00:00 5 2 10 82

Average 0 0 0 100
```

### ▼ How to Check System Table Status (`sar -v`)

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

```

Output from the `-v` option is described in the following table.

**TABLE 24-18** Output From the `sar -v` Command

| Field Name           | Description                                                                                                                                                                                                             |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>proc-sz</code> | The number of process entries ( <code>proc</code> structures) that are currently being used, or allocated in the kernel.                                                                                                |
| <code>inod-sz</code> | The total number of inodes in memory verses the maximum number of inodes that are allocated in the kernel. This number is not a strict high water mark; it can overflow.                                                |
| <code>file-sz</code> | The size of the open system file table. The <code>sz</code> is given as 0, since space is allocated dynamically for the file table.                                                                                     |
| <code>ov</code>      | The number of shared memory record table entries that are currently being used or allocated in the kernel. The <code>sz</code> is given as 0 because space is allocated dynamically for the shared memory record table. |
| <code>lock-sz</code> | The number of shared memory record table entries that are currently being used or allocated in the kernel. The <code>sz</code> is given as 0 because space is allocated dynamically for the shared memory record table. |

## Example—Checking System Table Status

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 venus 5.9 Generic sun4u 06/24/2001

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
02:00:02 43/922 0 2984/4236 0 322/322 0 0/0
03:00:00 43/922 0 2986/4236 0 323/323 0 0/0
04:00:00 43/922 0 2987/4236 0 322/322 0 0/0
05:00:01 43/922 0 2987/4236 0 322/322 0 0/0
06:00:00 43/922 0 2987/4236 0 322/322 0 0/0

```

## ▼ How to Check Swapping Activity (`sar -w`)

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 table describes target values and observations.

**TABLE 24-19** Output From the `sar -w` Command

| Field Name | Description                                                                                                                                    |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| swpin/s    | The number of LWP transfers into memory per second.                                                                                            |
| bswin/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. |
| 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—Checking Swap Activity

The following example shows output from the `sar -w` command.

```

$ sar -w
SunOS venus 5.9 Generic sun4u 06/24/2001

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
02:00:02 0.00 0.0 0.00 0.0 22
03:00:00 0.00 0.0 0.00 0.0 22
04:00:00 0.00 0.0 0.00 0.0 22
05:00:01 0.00 0.0 0.00 0.0 22
06:00:00 0.00 0.0 0.00 0.0 22
07:00:00 0.00 0.0 0.00 0.0 22
08:00:01 0.00 0.0 0.00 0.0 22
08:20:00 0.00 0.0 0.00 0.0 22
08:40:00 0.00 0.0 0.00 0.0 22
09:00:00 0.00 0.0 0.00 0.0 22
09:20:00 0.00 0.0 0.00 0.0 22
09:40:00 0.00 0.0 0.00 0.0 22
10:00:00 0.00 0.0 0.00 0.0 22
10:20:00 0.00 0.0 0.00 0.0 22

```



```

10:40:01 0.00 0.0 0.00 0.0 23
11:00:00 0.00 0.0 0.00 0.0 144

Average 0.00 0.0 0.00 0.0 24

```

## ▼ How to Check Terminal Activity (`sar -y`)

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 there are any bad lines. The activities recorded are defined in the following table.

**TABLE 24-20** Output From the `sar -y` Command

| Field Name           | Description                                                                |
|----------------------|----------------------------------------------------------------------------|
| <code>rawch/s</code> | Input characters (raw queue), per second                                   |
| <code>canch/s</code> | Input characters that are processed by canon (canonical queue), per second |
| <code>outch/s</code> | Output characters (output queue) per second                                |
| <code>rcvin/s</code> | Receiver hardware interrupts per second                                    |
| <code>xmtin/s</code> | Transmitter hardware interrupts per second                                 |
| <code>mdmin/s</code> | 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 this is not the case, check for bad lines.

## Example—Checking Terminal Activity

The following abbreviated example shows output from the `sar -y` command.

```

$ sar -y
SunOS venus 5.9 Generic sun4u 06/24/2001

00:00:00 rawch/s canch/s outch/s rcvin/s xmtin/s mdmin/s
01:00:00 0 0 0 0 0 0
02:00:02 0 0 0 0 0 0
03:00:00 0 0 0 0 0 0
04:00:00 0 0 0 0 0 0

```

|          |   |   |    |   |   |   |
|----------|---|---|----|---|---|---|
| 05:00:01 | 0 | 0 | 0  | 0 | 0 | 0 |
| 06:00:00 | 0 | 0 | 0  | 0 | 0 | 0 |
| 07:00:00 | 0 | 0 | 0  | 0 | 0 | 0 |
| 08:00:01 | 0 | 0 | 0  | 0 | 0 | 0 |
| 08:20:00 | 0 | 0 | 0  | 0 | 0 | 0 |
| 08:40:00 | 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 | 0  | 0 | 0 | 0 |
| 10:00:00 | 0 | 0 | 0  | 0 | 0 | 0 |
| 10:20:00 | 0 | 0 | 0  | 0 | 0 | 0 |
| 10:40:01 | 0 | 0 | 20 | 0 | 0 | 0 |
| Average  | 0 | 0 | 3  | 0 | 0 | 0 |

## ▼ How to Check Overall System Performance (sar -A)

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 it 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 `sadd`, 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 ( $t$  should be greater than 5 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 in order to record the statistics from when the counters are reset to zero. To make sure that `sadc` is run at boot time, the `/etc/init.d/perf` file contains a command line that 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 uncomment the following lines in the `/var/spool/cron/sys` crontab 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/sadd` 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/sardd` 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 5 seconds or greater. Otherwise, the command itself might affect the sample. You must specify a time interval between 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 is 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, either the file specified by the `-f` option or, by default, the standard daily activity file, `/var/adm/sa/sadd`, for the most recent day.
- The `-s` and `-e` options define the starting and ending times for the report. Starting and ending times are of the form `hh[:mm[:ss]]` (where  $h$ ,  $m$ , and  $s$  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 24-21** 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     |
| -p     | Checks swap and dispatch activity     |
| -q     | Checks queue activity                 |
| -r     | Checks unused memory                  |

**TABLE 24-21** Options for the `sar` Command (Continued)

| Option | Actions                                                           |
|--------|-------------------------------------------------------------------|
| -u     | Checks CPU utilization                                            |
| -nv    | Checks system table status                                        |
| -w     | Checks swapping and switching volume                              |
| -y     | Checks terminal activity                                          |
| -A     | Reports overall system performance (same as entering all options) |

If no option is used, it is equivalent to calling the command with the `-u` option.

## How to Set Up Automatic Data Collection

### 1. Become superuser.

### 2. Edit the `/etc/init.d/perf` file and uncomment the following lines:

```
if [-z "$_INIT_RUN_LEVEL"]; then
set -- ` /usr/bin/who -r `
_INIT_RUN_LEVEL="$7"
_INIT_RUN_NPREV="$8"
_INIT_PREV_LEVEL="$9"
fi
#
if [$_INIT_RUN_LEVEL -ge 2 -a $_INIT_RUN_LEVEL -le 4 -a \
$_INIT_RUN_NPREV -eq 0 -a \($_INIT_PREV_LEVEL = 1 -o \
$_INIT_PREV_LEVEL = S \)]; then
/usr/bin/su sys -c "/usr/lib/sa/sadc /var/adm/sa/sa`date +%d`"
fi
```

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
```



# Troubleshooting Solaris Software Topics

---

This topic map lists the chapters that provide information for troubleshooting Solaris software problems.

|            |                                                                                                                                    |
|------------|------------------------------------------------------------------------------------------------------------------------------------|
| Chapter 26 | Provides overview information about troubleshooting common software problems and instructions for troubleshooting a system crash.  |
| Chapter 27 | Provides step-by-step instructions for managing core files.                                                                        |
| Chapter 28 | Provides step-by-step instructions for saving crash dumps and customizing system error logging.                                    |
| Chapter 29 | Provides problem scenarios and possible solutions for general software problems such as a hung system or a system that won't boot. |
| Chapter 30 | Provides solutions for solving common file access problems such as incorrect command search paths and file permissions.            |
| Chapter 31 | Provides solutions for solving common printer problems such as no output or incorrect output.                                      |
| Chapter 32 | Provides specific <code>fsck</code> error messages and corresponding solutions for solving UFS file system inconsistency problems. |
| Chapter 33 | Provides specific error messages and possible solutions for problems encountered when adding or removing software packages.        |





## 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 Software Problems?” on page 353
- “Where to Find Software Troubleshooting Tasks” on page 355
- “Troubleshooting a System Crash” on page 355
- “Troubleshooting a System Crash Checklist” on page 357
- “Viewing System Messages” on page 358
- “Customizing System Message Logging” on page 360

---

### What’s New in Troubleshooting Software Problems?

This section describes features that are new in the Solaris 9 release.

#### New System Log Rotation

In Solaris 9 release, system log files are now rotated by the `logadm` command from an entry in the root `crontab` file. The `/usr/lib/newsyslog` script is no longer used.

The new 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*.

## New Fall Back Shell for root Account

If you changed root's shell to a non-existent shell in previous Solaris releases, you were forced to boot the system from a local CD or from the network and correct the root shell entry in the `/etc/passwd` file.

If you mistakenly provide a non-existent shell for root in the Solaris 9 release, root's shell will automatically fall back to `/sbin/sh` when one of the following occurs:

- You use the `su` command to change to the superuser (root) account.
- You are prompted for the root password at the single-user mode password prompt.

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

---

## Where to Find Software Troubleshooting Tasks

| Troubleshooting Task                                                       | For More Information |
|----------------------------------------------------------------------------|----------------------|
| Manage system crash information                                            | Chapter 28           |
| Manage core files                                                          | Chapter 27           |
| Troubleshoot software problems such as reboot failures and backup problems | Chapter 29           |
| Troubleshoot file access problems                                          | Chapter 30           |
| Troubleshoot printing problems                                             | Chapter 31           |
| Resolve UFS file system inconsistencies                                    | Chapter 32           |
| Troubleshoot software package problems                                     | Chapter 33           |

---

## Troubleshooting a System Crash

If a system running the Solaris operating environment 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 359.

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 355.

If the system fails to reboot successfully after a system crash, see Chapter 29.

2. Synchronize the disks and reboot.

```
ok sync
```

If the system fails to reboot successfully after a system crash, see Chapter 29.

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 28.

## Gathering Troubleshooting Data

Answer the following questions to help isolate the system problem. Use “Troubleshooting a System Crash Checklist” on page 357 for gathering troubleshooting data for a crashed system.

**TABLE 26-1** Identifying System Crash Data

| Question                              | Description                                                                                                                                                                                                                                      |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Can you reproduce the problem?</i> | 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. |

**TABLE 26-1** Identifying System Crash Data (Continued)

| Question                                                                       | Description                                                                                                                                                             |
|--------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Are you using any third-party drivers?</i>                                  | 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.                                  |
| <i>What was the system doing just before it crashed?</i>                       | 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.                       |
| <i>Were there any unusual console messages right before the crash?</i>         | Sometimes the system will show signs of distress before it actually crashes; this information is often useful.                                                          |
| <i>Did you add any tuning parameters to the <code>/etc/system</code> file?</i> | 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.            |
| <i>Did the problem start recently?</i>                                         | 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 <code>prtdiag</code> output for sun4u systems.<br>Include Explorer output for other systems. |           |
| Are patches installed? If so, include <code>showrev -p</code> output.                                |           |
| Is the problem reproducible?                                                                         |           |
| Does the system have any third-party drivers?                                                        |           |

| Item                                                                     | Your Data |
|--------------------------------------------------------------------------|-----------|
| 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 <code>/etc/system</code> file?         |           |
| Did the problem start recently?                                          |           |

## 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 “How to Customize System Message Logging” on page 361. 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 230 and Chapter 18.

## ▼ How to View System Messages

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
```

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

## Example—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 5.9 ...
Jan 3 08:44:41 starbug genunix: [ID 913631 kern.notice] Copyright 1983-2002 ...
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
.
.
.
```

---

## 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 359.

The `/etc/syslog.conf` file has two columns separated by tabs:

*facility.level ... action*

|                       |                                                                                                                                                                                                                                                                                       |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>facility.level</i> | A <i>facility</i> or system source of the message or condition. May be a comma-separated listed of facilities. Facility values are listed in Table 26–2. A <i>level</i> , indicates the severity or priority of the condition being logged. Priority levels are listed in Table 26–3. |
| <i>action</i>         | 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 26–3 in order of severity.

**TABLE 26–2** Source Facilities for `syslog.conf` Messages

| Source | Description    |
|--------|----------------|
| kern   | The kernel     |
| auth   | Authentication |



**TABLE 26–2** Source Facilities for `syslog.conf` Messages (Continued)

| Source | Description     |
|--------|-----------------|
| 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 26–3** 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

1. **Become superuser.**
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—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 363.)

- 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 362.

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 `consadm(1M)`.

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—Enabling an Auxiliary (Remote) Console

```
consadm -a /dev/term/a
consadm
/dev/term/a
```

## ▼ How to Display a List of Auxiliary Consoles

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

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—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

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—Disabling an Auxiliary (Remote) Console

```
consadm -d /dev/term/a
consadm
```

## 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 367.

---

## 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 <code>coreadm</code> command. | “How to Display the Current Core Dump Configuration” on page 370 |
| 2. Modify the core dump configuration          | Modify the core dump configuration to do one of the following:                         |                                                                  |
|                                                | Set a core file name pattern.                                                          | “How to Set a Core File Name Pattern” on page 370                |
|                                                | Enable a per-process core file path.                                                   | “How to Enable a Per-Process Core File Path” on page 371         |
|                                                | Enable a global core file path.                                                        | “How to Enable a Global Core File Path” on page 371              |
| 3. Examine a Core Dump File                    | Use the <code>proc</code> tools to view a core dump file.                              | “Examining Core Files” on page 372                               |

---

## 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 as in previous Solaris releases. But if the global core file path is enabled and set to `/corefiles/core`, for example, then each process that expires produce *two* core files: one in the current working directory and one in the `/corefiles` directory.

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                                              |
|---------------|------------------------------------------------------------------|
| %p            | Process ID                                                       |
| %u            | Effective user ID                                                |
| %g            | Effective group ID                                               |
| %f            | Executable file name                                             |
| %n            | System node name, equivalent to the <code>uname -n</code> output |
| %m            | Machine name, equivalent to the <code>uname -m</code> output     |
| %t            | Decimal value of <code>time(2)</code> system call                |
| %%            | 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 basis or a per-process basis, and you can specify whether you want these settings saved across a system reboot.

For example, the following `coreadm` command sets the global core file pattern for all processes started by the `init` process. This pattern will persist 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 all 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: /var/core/core.%f.%p
 init core file pattern: core
 global core dumps: enabled
 per-process core dumps: enabled
 global setid core dumps: enabled
 per-process setid core dumps: disabled
 global core dump logging: disabled
```

## How to Set a Core File Name Pattern

1. **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. Set a global file name pattern.

Become superuser first.

```
coreadm -g /var/corefiles/%f.%p
```

## ▼ How to Enable a Per-Process Core File Path

1. Become superuser.
2. Enable a per-process core file path.
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

1. Become superuser.
2. Enable a global core file path.
3. Display the current process core file path to verify the configuration.

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

---

# Troubleshooting Core File Problems

## Error Message

```
NOTICE: 'set allow_setuid_core = 1' in /etc/system is obsolete
NOTICE: Use the coreadm command instead of 'allow_setuid_core'
```

## Cause

You have an obsolete parameter that allows setuid core files in your `/etc/system` file.

## Solution

Remove `allow_setuid_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 `pcrred` 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—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
```

## Managing System Crash Information (Tasks)

---

This chapter describes how to manage system crash information in the Solaris environment.

For information on the procedures associated with managing system crash information, see “Managing System Crash Information (Task Map)” on page 373.

---

## 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 <code>dumpadm</code> command.                                                                                                                                                                                 | “How to Display the Current Crash Dump Configuration” on page 377 |
| 2. Modify the crash dump configuration          | Use the <code>dumpadm</code> 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 377            |
| 3. Examine a crash dump file                    | Use the <code>mdb</code> command to view crash dump files.                                                                                                                                                                                                              | “How to Examine a Crash Dump” on page 379                         |

| Task                                                           | Description                                                                                                                                       | For Instructions                                                         |
|----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| 4. (Optional) Recover from a full crash dump directory         | The system crashes but there is no room in the <code>savecore</code> directory, and you want to save some critical system crash dump information. | "How to Recover From a Full Crash Dump Directory (Optional)" on page 380 |
| 5. (Optional) Disable or enable the saving of crash dump files | Use the <code>dumpadm</code> 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 380                |

---

## 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 375.

## 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 `mdb(1M)`.

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 environment.

- 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 Solaris operating environment. 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 may 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, <code>savecore</code> runs in the background, which speeds up the boot process. |
| savecore directory | The directory that stores system crash dump files.                                                                                                                                             |
| dump content       | Type of data, kernel memory or all of memory, to dump.                                                                                                                                         |
| minimum free space | Minimum amount of free space required in the <code>savecore</code> 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 `/etc/dumpadm.conf` edit manually. This 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, it 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. It will also 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 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

1. **Become superuser.**
2. **Display the current crash dump configuration.**

```
dumpadm
 Dump content: kernel pages
 Dump device: /dev/dsk/c0t3d0s1 (swap)
Savecore directory: /var/pluto
 Savecore enabled: yes
```

The above 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

1. **Become superuser.**
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 is 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 <i>content</i>            | Specifies the type of data to dump. Use <code>kernel</code> to dump of all kernel memory or <code>all</code> to dump all of memory. The default dump content is kernel memory.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| -d <i>dump-device</i>        | Specifies the device that stores dump data temporarily as the system crashes. The primary swap device is the default dump device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| -m <i>nnnk   nnnm   nnn%</i> | Specifies the minimum free disk space for saving crash dump files by creating a <code>minfree</code> file in the current <code>savecore</code> directory. This parameter can be specified in Kbytes ( <code>nnnk</code> ), Mbytes ( <code>nnnm</code> ) or file system size percentage ( <code>nnn%</code> ). The <code>savecore</code> 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 <code>minfree</code> 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 380. |
| -n                           | Specifies that <code>savecore</code> 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 <code>savecore</code> 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 <code>/var/crash/hostname</code> where <code>hostname</code> is the output of the <code>uname -n</code> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## Example—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

1. **Become superuser.**
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—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: 5.9 Generic (sun4u)
> ::system
set ufs_ninode=0x9c40 [0t40000]
set ncsiz=0x4e20 [0t20000]
set pt_cnt=0x400 [0t1024]
```

## ▼ How to Recover From a Full Crash Dump Directory (Optional)

In this scenario, the system crashes but there is no room in the `savecore` directory, and you want to save some critical system crash dump information.

1. **Log in as superuser 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

1. **Become superuser.**
2. **Disable or enable the saving of crash dumps on your system.**

```
dumpadm -n | -y
```

### Example—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—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
```

---

## 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” on page 381
- “What to Do if a System Hangs” on page 385
- “What to Do if a File System Fills Up” on page 386
- “What to Do if File ACLs Are Lost After Copy or Restore” on page 387
- “Troubleshooting Backup Problems” on page 387

---

### 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 <code>/platform/'uname -m'/kernel/unix</code> . | You may need to change the <code>boot-device</code> setting in the PROM on a SPARC system. For information on changing the default boot device, see “SPARC: Booting a System (Tasks)” in <i>System Administration Guide: Basic Administration</i> . |

| Cause of System Not Booting                                                                         | How to Fix the Problem                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| There is no default boot device on an IA system. The message displayed is:<br>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 <code>/etc/passwd</code> file.                                      | For information on recovering from an invalid <code>passwd</code> file, see "SPARC: Booting a System (Tasks)" in <i>System Administration Guide: Basic Administration</i> or "IA: Booting a System (Tasks)" in <i>System Administration Guide: Basic Administration</i> .                                                                                                                    |
| There's a hardware problem with a disk or another device.                                           | Check the hardware connections: <ul style="list-style-type: none"> <li>■ Make sure the equipment is plugged in.</li> <li>■ Make sure all the switches are set properly.</li> <li>■ Look at all the connectors and cables, including the Ethernet cables.</li> <li>■ If all this fails, turn off the power to the system, wait 10 to 20 seconds, and then turn on the power again.</li> </ul> |

If none of the above suggestions solve the problem, contact your local service provider.

## SPARC: Troubleshooting 64-bit Solaris Boot Problems

After the 64-bit Solaris release is installed on an UltraSPARC system, the 64-bit kernel will be booted automatically unless any of the following conditions are true:

- A FLASH PROM upgrade might be required on an UltraSPARC system before it can successfully boot the 64-bit kernel. Refer to your hardware manufacturer's documentation to determine whether your UltraSPARC system requires a firmware upgrade.
- The Open Boot PROM `boot-file` parameter is set to `kernel/unix`. If booting the 64-bit kernel fails and this parameter is set, unset it, and reboot the system.
- On some UltraSPARC systems, the 64-bit Solaris kernel is not booted by default, even when the system is completely installed with all the 64-bit Solaris components and the correct firmware is installed. Without booting the 64-bit Solaris kernel, 64-bit applications are unable to run.

To find out more about this issue, and how to enable booting the 64-bit Solaris kernel by default, see `boot(1M)`.

You can always discover which Solaris kernel the system is currently running by using the `isainfo -kv` command.

```
$ isainfo -kv
64-bit sparcv9 kernel modules
```

This output means the system is running the 64-bit Solaris kernel.

You cannot boot the 64-bit Solaris operating environment on a 32-bit Solaris system.

---

## 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 “SPARC: Booting a System” in *System Administration Guide: Basic Administration* and “IA: 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 IA systems.

### SPARC: Example—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
```

## IA: Example—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

Solaris Intel Platform Edition Booting System

Running Configuration Assistant...

Autobooting from Boot path: /pci@0,0/pci-ide@7,1/ide@0/cmdk@0,0:a

If the system hardware has changed, or to boot from a different device, interrupt the autoboot process by pressing ESC.

Press ESCape to interrupt autoboot in 5 seconds.

.  
.  
.

Boot Solaris

Select one of the identified devices to boot the Solaris kernel and choose Continue.

To perform optional features, such as modifying the autoboot and property settings, choose Boot Tasks.

An asterisk (\*) indicates the current default boot device.

> To make a selection use the arrow keys, and press Enter to mark it [X].

```
[X] NET : DEC 21142/21143 Fast Ethernet
on Board PCI at Dev 3
[] DISK: (*) Target 0, QUANTUM FIREBALL1280A
on Bus Mastering IDE controller on Board PCI at Dev 7, Func 1
[] DISK: Target 1:ST5660A
on Bus Mastering IDE controller on Board PCI at Dev 7, Func 1
[] DISK: Target 0:Maxtor 9 0680D4
on Bus Mastering IDE controller on Board PCI at Dev 7, Func 1
[] CD : Target 1:TOSHIBA CD-ROM XM-5602B 1546
on Bus Mastering IDE controller on Board PCI at Dev 7, Func 1

F2_Continue F3_Back F4_Boot Tasks F6_Help
```



```

.
.
.
 <<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci-ide@7,1/ide@0/cmdk@0,0:a
Boot args: kernel/unix -r

Select the type of installation you want to perform:

1 Solaris Interactive
2 Custom JumpStart
3 Solaris Web Start

Enter the number of your choice followed by <ENTER> the key.

If you enter anything else, or if you wait for 30 seconds,
an interactive installation will be started.

Select type of installation: b -s
.
.
.
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 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 “SPARC: Forcing a Crash Dump and Rebooting the System” in *System Administration Guide: Basic Administration* or “IA: 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 16.

### 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 and use the <code>ls -tl</code> 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 230. |

## 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 tmpfs(7FS). |

---

## 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 <code>-f</code> option, the <code>ufsdump</code> command wrote to a file in the <code>/dev</code> directory of the root ( <code>/</code> ) file system, filling it up. For example, if you typed <code>/dev/rmt/st0</code> instead of <code>/dev/rmt/0</code> , the backup file <code>/dev/rmt/st0</code> was created on the disk rather than being sent to the tape drive. | Use the <code>ls -t1</code> command in the <code>/dev</code> 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: Thu Aug 30 10:13:46 2001
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 “UFS Backup and Restore Commands (Reference)” in *System Administration Guide: Basic Administration* to find files, create a list of files to be restored, and restore them.



## 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.

- “Solving Problems With Search Paths (Command not found)” on page 391
- “Solving File Access Problems” on page 394
- “Recognizing Problems With Network Access” on page 394

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 that 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

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 | <code>\$HOME/.profile</code>                                   | <code>\$<br/>PATH=\$HOME/bin:/sbin:/usr/local/bin ...<br/>\$ export PATH</code> | A colon separates path names.       |
| C               | <code>\$HOME/.cshrc</code><br>or<br><code>\$HOME/.login</code> | <code>hostname% set path=(~bin /sbin /usr/local/bin ...)</code>                 | 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 | <code>.profile</code>      | <code>\$ . ./profile</code>           |
| C               | <code>.cshrc</code>        | <code>hostname% source .cshrc</code>  |
|                 | <code>.login</code>        | <code>hostname% source .login</code>  |

#### 5. Verify the new path.

```
$ which command
```

## Example—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), `lpsched(1M)` 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 "Securing Files (Tasks)" in *System Administration Guide: Security Services*.

---

## Recognizing Problems With Network Access

If users have problems using the `rpc` 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: Resource Management and Network Services* for information about problems with network access and problems with accessing systems through AutoFS.

## 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.

- “How to Troubleshoot No Printer Output” on page 401
- “How to Troubleshoot Incorrect Output” on page 407
- “How to Unhang the LP Print Service” on page 411
- “How to Troubleshoot an Idle (Hung) Printer” on page 411
- “How to Resolve Conflicting Printer Status Messages” on page 413

For information about printing and the LP print service, see Chapter 2.

---

## 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 397.

### 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 (FNS and 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 77.

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 401 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 `lpssystem`, `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 76 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 94 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 401).
- Check the network (“How To Check the Printer Network Connections” on page 402).
- Check the LP print service basic functions (“How to Check the Basic Functions of the LP Print Service” on page 403).
- 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 406).

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 407.

## ▼ How to Check the Printer Hardware

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.

**TABLE 31-1** 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 407.

## ▼ How To Check the Printer Network Connections

**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.

1. On the print server, make sure that the LP print service is running.
  - a. Check whether the LP scheduler is running.

```
lpstat -r
scheduler is running
```
  - b. If the scheduler is not running, become superuser or `lp`, 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 411.
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 Jul 12 14:23 2001
luna not accepting requests since Jul 12 14:23 2001
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 or `lp`, 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 Jul 12 14:25 2001.
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 or lp, 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 or lp.**

**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 or lp 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 Jul 12 14:24 2001. 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

**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 403. 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 Jul 12 14:26 2001
printer luna now printing luna-314. enabled since Jul 12 14:26 2001.
available.
```

```
luna-129 root 488 Jul 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

1. Log in as superuser or lp.
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 44.

- 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 132.

**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 istryp -inlcr -igncr icrnl -iuclc
ixon -ixany -ixoff imaxbel
isig icanon -xcase echo echoe echok -echonl -noflsh
-tostop echoctl -echoprnt 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.

**TABLE 31-2** Default stty Settings Used by the Standard 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                                            |

**TABLE 31-2** Default `stty` Settings Used by the Standard Interface Program  
(Continued)

| Option              | Meaning                                        |
|---------------------|------------------------------------------------|
| <code>-onlcr</code> | Change line feed to carriage return/line feed  |
| <code>-ocrnl</code> | Do not change carriage returns into line feeds |
| <code>-onocr</code> | Output carriage returns even at column 0       |
| <code>-nl0</code>   | No delay after line feeds                      |
| <code>-cr0</code>   | No delay after carriage returns                |
| <code>-tab0</code>  | No delay after tabs                            |
| <code>-bs0</code>   | No delay after backspaces                      |
| <code>-vt0</code>   | No delay after vertical tabs                   |
| <code>-ff0</code>   | 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.

**TABLE 31-3** `stty` Options to Correct Print Output Problems

| <code>stty</code> 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 |
| <code>oddp</code>                                         | Sets odd parity                                                  | Missing or incorrect characters appear randomly                                             |
| <code>evenp</code>                                        | Sets even parity                                                 |                                                                                             |
| <code>-parity</code>                                      | Sets no parity                                                   |                                                                                             |
| <code>-tabs</code>                                        | Sets no tabs                                                     | Text is jammed against right margin                                                         |
| <code>tabs</code>                                         | Sets tabs every eight spaces                                     | Text has no left margin, is run together, or is jammed together                             |
| <code>-onlcr</code>                                       | Sets no carriage return at the beginning of line(s)              | Incorrect double spacing                                                                    |
| <code>onlcr</code>                                        | 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

1. Log in as superuser or `lp`.
2. Stop the LP print service.

```
/etc/init.d/lp stop
```

3. Restart the LP print service.

```
/etc/init.d/lp start
```

The LP print service should restart. If you are having trouble restarting the scheduler, see “How to Restart the Print Scheduler” on page 77.

## ▼ 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

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 411.

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

### 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

### 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

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 `lp sched` 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

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.



## 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.

- “General `fsck` Error Messages” on page 417
- “Initialization Phase `fsck` Messages” on page 418
- “Phase 1: Check Blocks and Sizes Messages” on page 421
- “Phase 1B: Rescan for More DUPS Messages” on page 424
- “Phase 2: Check Path Names Messages” on page 424
- “Phase 3: Check Connectivity Messages” on page 430
- “Phase 4: Check Reference Counts Messages” on page 432
- “Phase 5: Check Cylinder Groups Messages” on page 435
- “Cleanup Phase Messages” on page 436

For information about the `fsck` command and how to use it to check file system integrity, see “Checking File System Integrity (Tasks)” in *System Administration Guide: Basic Administration*.

---

### `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 417. 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 32–1** 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 (`-op` 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.

---

## 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.

- “Specific Software Package Installation Errors” on page 438
- “General Software Package Installation Problems” on page 439

For information about managing software packages, see “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 <code>pkgadd</code> to install a package on a client. In this case, <code>pkgadd</code> attempts to install a package on a file system that is mounted from a server, but <code>pkgadd</code> 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 "Managing Software (Overview)" in <i>System Administration Guide: Basic Administration</i> 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.<br><code>NONABI_SCRIPTS=TRUE</code> |





# Index

---

## Numbers and Symbols

- 64-bit Solaris
  - determining if capabilities are enabled, 212
  - determining if the system can run, 212

## A

- accept command, 93
- accepting print requests (how to), 94
- accepting print requests (overview), 93
- access
  - to forms (how to)
    - limiting for printers, 122
    - limiting for users, 121
  - to printers
    - deleting (how to), 72
- accounting, 277, 293
  - billing users, 271, 274, 286
  - connect, 268, 282, 285, 291
  - daily, 270, 293
    - reports, 283, 290
    - step-by-step summary of, 270, 272
  - disabling, 278
  - disk, 269, 286
  - files for, 291, 293
  - fixing corrupted files
    - tacct file, 276
    - wtmpx file, 275, 282
  - maintaining, 276
  - overview, 267
  - process, 269, 271, 285
  - raw data, 270
- accounting (*continued*)
  - reports, 283, 290
    - daily command summary, 282, 286, 291
    - daily report (tty line utilization), 284
    - daily usage report, 285
    - last login report, 288
    - overview, 283
    - total command summary (monthly), 288, 292
  - set up to run automatically (how to), 273
  - starting, 273
  - stopping, 278
  - types of, 274
  - user fee calculation, 271, 274, 286
- acct.h format files, 289
- acctcms command, 282, 293
- acctcom command, 289
- acctcon command, 275, 282, 291
- acctdusg command, 270, 286, 291
- acctprc command, 282
- acctwtmp command, 269, 271, 284
- active file, 281
- active file, 277, 291
- active.MMDD file, 277, 291
- adapter board (serial port), 173
- adding
  - access to remote printers, 57
  - forms (how to), 116
  - local or attached printers (how to), 55
  - print filters (how to), 111
  - printer, new (how to)
    - with Solaris Print Manager, 55
  - printer description (how to), 77

- adding (*continued*)
  - terminfo entry (overview), 129
- address space map, 309
- adjusting printer port characteristics (overview), 127
- administering
  - forms, 122
- alert message priority (for `syslogd`), 361
- alerts
  - for mounting character sets, 106
  - for mounting forms (how to), 119
  - for mounting forms (overview), 115
  - for printer faults (how to), 84
- alias for selectable character set, 102, 108
- alignment pattern
  - defining, 150
  - printing, 118
  - protection of, 115
- allow list
  - for printer access to forms, 122
  - for user access to forms, 121
  - for user access to printers (overview), 87
- alphanumeric terminal, *See* terminals
- application threads, 301
- ASCII file
  - file content type of, 46
- at command, 260, 264
  - l option (list), 263
  - m option (mail), 261
  - automatic scheduling of, 253
  - controlling access to, 261, 264
    - overview, 250
  - denying access, 264
  - error messages, 265
  - overview, 250, 260
- at .deny file, 261, 264
  - description, 250
- at job files, 260, 264
  - creating, 261
  - deleting, 263
  - description, 251
  - displaying, 263
  - displaying queue of, 262
  - location of, 251
  - submitting, 260
  - verifying, 262
- atjobs directory, 253
  - description, 250

- atq command, 262
- automatic system activity data collection, 346, 349
- automatic system activity reporting, 346
- automatic system task execution
  - repetitive tasks, 258
  - single tasks, 260, 264
- automatically turning on quotas, 235
- automating system task execution, 250
- auxiliary (remote) console, 362

## B

- banner option, 80
- banner pages
  - making optional (how to), 80
    - with `lpadmin` command, 80
  - nothing else prints, 408
  - reasons to turn off, 79
  - setting
    - with `lpadmin` command, 79
    - with Solaris Print Manager, 42
  - troubleshooting incorrect output, 397
  - turning off (how to), 81
    - with -o banner option, 81
- baud settings, 398
- bidirectional modem service, 172, 190
- billing users, 271, 274, 286
- booting
  - displaying messages generated during, 359
  - running `sadc` command when, 347

## C

- cancel command, 95
- canceled
  - print requests (how to), 95
    - by disabling printer, 92
    - for specific user, 96
  - print requests (overview), 95
- cartridges, *See* font cartridges
- centralized
  - print configuration, 38
- changing
  - crontab files, 254
  - date, 216

- changing (*continued*)
  - forms paper, 118
  - message of the day, 217
  - priority, 315, 318
    - timesharing processes, 315, 317
  - priority of print requests (how to), 99
  - priority of print requests (overview), 89
  - quotas for individual users, 245
  - scheduling classes, 316
  - soft limit time, 244
  - system's host name, 217
  - time, 216
- character sets
  - hardware, 102
  - number, 102
  - selectable, 102
  - software, 102
- chargefee script, 271, 274, 286
- checking quota consistency, 240
- ckpacct script, 271
- class (printer), 81
  - checking status for, 90
  - defining with `lpadmin` command (how to), 82
  - not valid for enabling/disabling printer, 92
- `closewtmp` command, 282
- `cmsprev` file, 292
- Command not found error message, 391
- commands
  - monitoring usage of, 291
- configuring
  - printer ports, 43
    - for IA systems, 43
- connect accounting, 268, 285, 291
- `consadm` command, 364
  - disabling an auxiliary console, 365
  - displaying list of auxiliary consoles (how to), 365
  - enabling an auxiliary console, 364
    - across system reboots, 365
- console
  - auxiliary
    - enabling across system reboots, 365
- controlling
  - access to at command, 250, 261, 264
  - access to `crontab` command, 258
    - overview, 250
  - printer access to forms (how to), 122
- controlling (*continued*)
  - processes, 311
    - user access to forms (how to), 121
- core dump configuration
  - displaying with `coreadm`, 370
- core file name pattern
  - setting with `coreadm`, 369
- core files
  - automatically deleting, 261
- core files
  - examining with `proc` tools, 372
  - finding and deleting, 230
- core files
  - managing with `coreadm`, 368
- `coreadm` command, 368
  - displaying core dump configuration, 370
  - enabling a global core file path, 371
  - enabling a per-process core file path, 371
  - managing core files, 368
  - setting a core file name pattern, 370
- CPU (central processing unit)
  - displaying information on
    - time usage, 286, 306, 318
  - high-usage processes, 318
- crash dump configuration
  - displaying with `dumpadm`, 377
  - modifying with `dumpadm`, 377
- crash dump directory
  - recovering from a full, 380
- crash dumps
  - disable saving, 380
  - enable saving, 380
- crashes, 360, 386
  - customer service and, 356, 375
  - deleting crash dump files, 230
  - displaying system information generated
    - by, 358, 379
  - examining crash dumps, 379
  - procedure following, 356, 386
  - rebooting fails after, 381
  - saving crash dump information, 374
  - saving other system information, 359
- creating
  - at jobs, 261
  - at jobs, 262
  - `crontab` files, 254
  - form definitions (how to), 150
  - forms, 148

- creating (*continued*)
  - print filters (how to), 146
  - print filters (overview), 137
- cron.allow file, 257
- cron daemon, 251
- cron.deny file, 257
  - defaults, 258
- crontab command, 258
  - accounting scripts run by, 272, 274
  - controlling access to, 257
    - denying access, 257
    - limiting access to specific users, 258
    - overview, 250, 257
  - cron daemon and, 252
  - e option (edit), 254
  - l option (list), 255
  - r option (remove), 256
  - /var/adm maintenance and, 359
  - daily tasks, 250
  - error messages, 260
  - files used by, 252
  - overview, 249
  - quitting without saving changes, 254
  - scheduling of, 252
- crontab files
  - creating, 254
  - defaults, 252
  - deleting, 256
  - denying access, 258
  - description, 252
  - displaying, 255
  - editing, 254
  - location of, 252
  - removing, 257
  - syntax, 253
  - verifying existence of, 255
- ctacct.MMDD file, 282, 291
- ctmp file, 291
- customer service
  - sending crash information, 356
- customizing
  - exit codes, printer, 134
  - LP print service, 150
  - printer interface program (how to), 136
  - printer interface program (overview), 133
  - stty modes, 134
  - system message logging, 360
  - system message logging (how to), 361

## D

- daemons
  - lpsched, 161, 163
  - lpsched (reference), 167
  - print, 155
- daily accounting, 270, 293
  - reports, 283, 290
    - daily command summary, 282, 286, 291
    - daily report (tty line utilization), 284
    - daily usage report, 285
    - last login report, 288
    - overview, 283
    - step-by-step summary of, 270, 272
- daily tasks (scheduling with crontab), 250
- date
  - changing, 216
  - displaying, 215
- date command
  - accounting data and, 269
  - described, 211
- daytacct file, 282, 286, 291, 293
- default printer
  - setting with lpadmin command (how to), 78
  - setting with lpadmin command (overview), 78
  - setting with Solaris Print Manager, 42
- defaults
  - message of the day, 217
  - nice number, 318
  - scheduling classes, 315
  - soft limit time, 244
- defining
  - font cartridges (how to), 104
  - print wheels (how to), 104
- deleting
  - access to printers (how to), 72
  - at jobs, 263
  - core files, 230
  - crash dump files, 230
  - crontab files, 256
  - forms (how to), 117
  - log files, 255
  - old/inactive files, 227, 230, 250, 255
  - print filters (how to), 112
  - temporary files, 229
- deny list
  - for printer access to forms, 122

- deny list (*continued*)
  - for user access to forms, 121
  - for user access to printers (overview), 87
- df command, 329
  - F option (unmounted file systems), 220
  - h option, 221
  - k option (kilobytes), 220, 329
  - t option (total blocks), 220
  - examples, 220, 329
  - overview, 219, 329
- dial-in modem service, 172
- dial-out modem service, 172
- direct printing, 45
- directories
  - current working directory for processes, 309
  - displaying information about, 222, 226
  - size of, 224, 226
  - temporary, clearing out, 227, 229
- disable command, 91
- disable command, example of, 92
- disabling
  - an auxiliary console with the `consadm` command, 365
  - printers (how to), 73, 92
  - quotas for individual users, 246
  - saving crash dumps, 380
  - system accounting, 278
- disk accounting, 269, 286
- disk drives
  - displaying information about free disk space, 329
  - finding and deleting old/inactive files, 255
- disk space
  - amount of free, 220, 329
  - displaying information about, 220
    - df command, 220, 329
    - directory sizes, 224, 226
    - disk space owned per user, 226
    - file sizes, 222
    - mount point, 329
    - user ownership of, 226
  - file system usage, 220, 329
  - finding and deleting old/inactive files, 227, 230
  - finding files exceeding a size limit, 223
  - finding large files, 223
  - optimizing, 219
- disk space for print queue, 40
- disktacct file, 270, 282, 291
- disktacct.MMDD file, 282
- dispadm command
  - overview, 313
- displaying
  - acct.h format files, 289
  - at jobs, 263
  - at jobs, 263
  - at queue, 262
  - booting messages, 359
  - core dump configuration with `coreadm`, 370
  - crash dump configuration with `dumpadm`, 377
  - crash information, 358, 379
  - crontab files, 255
  - date, 211, 215
  - directory information, 222, 226
  - disk utilization with `iostat`, 327
  - file information, 222, 226
  - file system information, 220, 329
  - host ID, 211, 214
  - linked libraries, 309
  - LWP information, 309
  - operating system information, 213
  - pacctn file, 289
  - priority information, 306, 314
  - process information (how to), 310
  - quota information, 236, 241
  - scheduling class information, 306, 314
  - size of files, 222
  - status of forms (how to), 120
  - system activity information, 330, 347
  - system information
    - commands for, 211, 215
    - time, 211, 215
    - virtual memory with `vmstat`, 323
- dmesg command, 359
- DNLC improvements, 297
- dodisk script, 269
  - caution, 270
  - crontab entry that runs, 274
  - files created by, 270, 282, 291
  - overview, 269
- downloaded PostScript fonts
  - installing (how to), 125

- downloading
  - fonts, 409
  - fonts (overview), 123
  - host-resident fonts (overview), 124
- dtmp file, 291
- du command, 224, 226
  - displaying size of directories, 224
- dumpadm
  - displaying crash dump configuration, 377
  - managing system crash information, 375
  - modifying crash dump configuration, 377

**E**

- editing
  - crontab files, 254
- edquota command
  - changing quotas for individual users, 245
  - disabling quotas for individual users, 246
  - p option (prototype), 239
  - t option (time limit), 244
  - overview, 236, 243
  - setting up user quotas, 238
- enable command, 91
- enable command, example of, 92
- enabling
  - a global core file path with coreadm, 371
  - a per-process core file path with coreadm, 371
  - an auxiliary console with consadm command, 364
  - auxiliary console across system reboots, 365
  - printers (how to), 92
  - saving crash dumps, 380
- environment variables
  - LPDEST, 79
  - PRINTER, 79
- error messages
  - at command, 265
  - crash messages, 359
  - crash related, 358
  - crontab command, 260
  - customizing logging of, 360
  - log file for, 356, 358
  - priorities for, 361
  - runacct script, 277
  - sources of, 360

- error messages (*continued*)
  - specifying storage location for, 358, 360
  - /etc/acct/holidays file, 273
  - /etc/cron.d/at.deny file, 261, 264
  - /etc/cron.d/cron.allow file, 257
  - /etc/cron.d/cron.deny file, 257
  - /etc/init.d/acct file, 273
  - /etc/init.d/perf file, 347, 349
  - /etc/inittab file, 188
  - /etc/lp/classes/printer-class file, 82
  - /etc/lp/default file, 79
  - /etc/lp directory, 153
  - /etc/lp/fd directory, 110, 163
  - /etc/lp/filter.table file
    - filter deleted from, 112, 110
    - filter added in, 112
  - /etc/lp/forms directory, 115
  - /etc/lp/forms/form-name/describe file, 116
  - /etc/lp/forms/form-name file, 117
  - /etc/lp/printers directory, 154
    - of print client, 73
    - of print server, 73
  - /etc/lp/Systems file, 73
  - /etc/motd file, 217
  - /etc/printcap file, 153
  - /etc/saf/\_sactab file, 188
  - /etc/syslog.conf file, 360
  - /etc/utmpx file, 190
  - /etc/vfstab file, 237, 240
- examining a core file
  - with proc tools, 372
- executing routine tasks automatically (overview), 249
- exit codes (printer interface), 134
  - standard, 134
  - table of, 134

**F**

- fast print filters (overview), 124
- fault notification (printer)
  - setting with lpadmin command (how to), 84
  - setting with lpadmin command (overview), 83
  - setting with Solaris Print Manager, 42

- fault notification (printer) *(continued)*
  - values for alerts, 83
- fault recovery (printer), 42
- fcntl information, 309, 311
- fd2log file, 277, 281, 291
- fee file, 271, 275, 282, 291
- fees (user), 271, 274, 286
- file content type, 45
  - converted by print filters, 109, 138
  - for common printers, 47
  - menu in Solaris Print Manager, 46
  - non-PostScript printers, 47
  - PostScript, 46
  - setting with Solaris Print Manager, 42
  - simple, 46
  - troubleshooting incorrect output, 397
- file or group ownership
  - solving file access problems, 394
- file systems
  - disk space usage, 220, 329
  - displaying information about, 220, 329
  - mount point, 329
  - restoring, 274, 286
- files
  - accounting, 291, 293
  - checking access operations, 330
  - deleting old/inactive, 227, 230, 250, 255
  - displaying information about
    - listing, 222
    - listing newest, 227
    - size, 222, 226
  - displaying size of, 222
  - finding and deleting old/inactive, 227, 230, 255, 262
  - finding files exceeding a size limit, 223
  - fixing corrupted
    - wtmpx file, 282
  - for setting search path, 392
  - fstat and fcntl information display, 309, 311
  - size of, 222, 226
  - usage monitoring, 269, 286
  - used by LP print service, 155
- filtering, 45
  - printing without, 45
- filters
  - download, 123
  - downloading (overview), 124
- find command, 223, 227, 230
- finding
  - and deleting, old/inactive files, 262
  - files exceeding a size limit, 223
  - large files, 223
  - old/inactive files, 227, 230, 255, 262
- fiscrptn file, 293
- fixing
  - corrupted tacct file, 276
  - corrupted wtmpx file, 275
- font cartridges, 103
  - alerts for mounting, 106
  - defining (how to), 104
  - mounting (how to), 105
  - naming, 103
  - unmounting (how to), 105
- fonts
  - downloaded PostScript, 125
  - downloading, 409
  - downloading (overview), 123
  - host-resident, 123
  - host-resident (overview), 124
  - installing (how to), 125
  - permanently downloaded, 123
  - PostScript, 123
  - printer-resident, 123
  - printer-resident (overview), 124
  - styles, 102
- forcing programs to quit, 386
- forget root password
  - IA, 384
  - SPARC, 383
- forms
  - adding (how to), 116
  - adding (overview), 114
  - alerts for mounting (overview), 115
  - allowing user access (how to), 121
  - changing (overview), 114
  - controlling access to (overview), 116
  - creating, 148
  - default values for, 149
  - definition, creating (how to), 150
  - deleting (how to), 117
  - deleting (overview), 114
  - denying user access (how to), 121
  - displaying attributes of (overview), 115
  - limiting printer access to (how to), 122
  - limiting user access to (how to), 121

## forms (*continued*)

- mounting (how to), 117
- mounting (overview), 114
- paper (loading and removing), 118
- setting alerts for mounting (how to), 119
- tracking forms mounted (overview), 115
- unmounting (how to), 117
- viewing information (how to), 120
- viewing status of (how to), 120

fsck command, 250

fstat information, 309, 311

## G

getty, 174

global core file path

- setting with `coreadm`, 368

global priorities

- defined, 313
- displaying, 314

## H

hard disk

- recommended for print server, 41

history log (print requests), 156

holidays file, 274

host name

- changing, 217

host-resident fonts

- downloading (overview), 124

hostid command, 211, 214

## I

init process, 188

initializing quotas, 237, 240

installing

- host-resident PostScript fonts (how to), 125
- PostScript fonts (how to), 125

interface program (printer)

- customizing (how to), 136
- customizing (overview), 133
- standard, 136

interrupting programs, 386

iostat command

- basic information display, 327
- displaying disk utilization, 327
- xtc option (extended), 328
- overview, 326

## K

kernel thread

- scheduling and, 306
- structures, 301, 306

killing processes, 309, 312

kilobytes

- file system disk usage in, 329

klwp structure, 301

kthread structure, 301

## L

large files, 223

last login report, 288

lastdate file, 282, 291

lastlogin command, 282

limiting

- printer access to forms (how to), 122
- user access to forms (how to), 121

line discipline, 189

line usage monitoring, 268, 284, 293

lineuse file, 293

lineuse file, 282

listing

- files and directories, 222, 227
- processes, 307
- processes being executed, 308

local or attached printer

- adding by using Solaris Print Manager (how to), 55

local printer

- defined, 38
- task map for setting up, 50

local printing, 164

lock file, 277, 282

lock1 file, 282

log file, 291

log files, 164

- cleaning out, 164



- log files (*continued*)
    - codes in request log, 157
    - deleting automatically, 255
    - for LP print service, 155, 397
    - print queue, 156
    - print request history log, 156
    - requests, 164
  - log.MMDD file, 291
  - login monitoring
    - last login, 282, 288, 293
    - number of logins, 286
    - time usage, 268, 271, 286
  - loginlog file, 282, 292
  - LP commands, 411
  - LP print service
    - checking basic functions of, 403, 406
    - configuration files in, 153
    - customizing, 150
    - daemons, 155
    - defining printer characteristics to, 38
    - definition, 152
    - directories in, 152
    - files used by, 155
    - hung LP commands, 411
    - interface program, 163
    - log files, 155, 397
    - reference information, 152
    - structure of, 152
    - tracking forms (overview), 115
    - tracking print wheels, 104
    - troubleshooting, 396, 399
  - LP print service scheduler, *See* print scheduler
  - lpadmin command
    - adding printer description with (how to), 77
    - adjusting printer port characteristics with (how to), 128
    - defining font cartridges with, 104
    - defining print wheels with, 104
    - defining printer class with (how to), 82
    - limiting access to printers with (how to), 88
    - limiting printer access to forms with (how to), 122
    - making banner pages optional with (how to), 80
    - mounting font cartridge with (how to), 105
    - mounting forms with (how to), 117
    - mounting print wheel with (how to), 105
    - lpadmin command (*continued*)
      - o banner option, 80
      - setting alerts to mount forms with (how to), 119
      - setting alerts to mount print wheels with (how to), 106
      - setting default printer with (how to), 79
      - setting printer fault alerts with (how to), 84
      - setting printer fault recovery with (how to), 86
      - unmounting forms with (how to), 117
  - LPDEST environment variable, 79
  - lpfilter command, 110
  - lpsched
    - daemon, 161, 163
    - daemon (reference), 167
    - log files, 164
  - ls command
    - checking directory sizes, 222, 227
    - l option (size in bytes), 223
    - s option (size in blocks), 223
    - t option (newest files), 227
  - LWPs (lightweight processes)
    - defined, 301
    - displaying information on, 309
    - processes and, 301
    - structures for, 301
- ## M
- managing
    - character sets, 101
    - fonts (overview), 123
    - forms, 114
    - print filters (overview), 109
  - managing system crash information
    - with dumpadm, 375
  - maximums
    - finding files exceeding maximum size, 223
    - nice number, 318
    - priority, 315
  - mdb utility, 379
  - memory
    - command for displaying information on, 211
    - displaying information on, 215
    - example of displaying information on, 215

- memory (*continued*)
  - process structures and, 301
  - shared
    - process virtual memory, 302
  - virtual
    - process, 302
- message of the day (MOTD) facility, 216
- messages file, 356, 360
- messages .n file, 358
- minimums
  - nice number, 318
  - priority, 315
- modems
  - bidirectional service, 172, 190
  - defined, 172
  - dial-in service, 172
  - dial-out service, 172
  - different ways to use, 172
  - overview of Serial Ports Tool, 177
  - Serial Ports Tool modem templates, 178
  - setting up (how to), 181
  - tools for managing, 174
- modifying
  - crash dump configuration with
    - dumpadm, 377
- monacct script
  - crontab entry that runs, 273
  - files used/produced by, 293
  - monthly command summary and, 286, 288
  - runacct script and, 272, 281
  - scheduling running of, 272
- monthly command summary, 288
- monthly tasks (scheduling with
  - crontab), 250
- MOTD (message of the day) facility, 216
- motd file, 216
- motd file, 217
- mounting
  - font cartridges (how to), 105
  - forms (how to), 117
  - forms (overview), 114
  - print wheels (how to), 105
- moving print requests (how to), 97
- moving print requests (overview), 97

## N

- network printer
  - adding, 59
  - adding (how to)
    - with lp commands, 64
    - with Vendor Supplied Tools, 63
  - defined, 59
  - task map for setting up, 50
- networks
  - recognizing access problems, 394
- nice command, 317
- nice number, 306, 318
- nlsadmin command, 192
- non-PostScript printers, 47, 102

## O

- operating system
  - displaying information about, 213
- optimizing disk space, 219
- owtmpx file, 292

## P

- pacctn file
  - displaying, 289
  - monitoring size of, 271, 281
  - overview, 271, 282, 291
- panic messages, 358
- parallel printer, 43
- parity bit
  - printers, 398
- per-process core file path
  - setting with coreadm, 368
- perf file, 347
- performance
  - activities that are tracked, 302
  - automatic collection of activity data, 346, 349
  - file access, 330
  - manual collection of activity data, 330, 347
  - process management, 301, 309, 318
  - reports on, 330
  - system activity monitoring, 302, 330, 346
  - tools for monitoring, 303
- pfiles command, 309, 311

- pflags command, 309
- pin configuration (NULL modem cable for printers), 401
- pkill command, 309, 312
- pldd command, 309
- pmadm command
  - adding a ttymon service with, 196
  - described, 188
  - disabling a ttymon service with, 200
  - enabling a ttymon service with, 199
  - listing a ttymon service with, 197
- pmap command, 309
- port, 43
  - defined, 172
  - initialization process of, 190
  - initializing (how to), 181
  - states of (table), 204
- port monitor
  - definition, 173
  - states of (table), 203
  - ttymon and listen (defined), 173, 190
- PostScript fonts, 123
  - installing (how to), 125
- PostScript printers, 102
  - character sets for, 103
  - default print filters, 110
  - file content type for, 46
  - printer type for, 45
- PostScript Reverse printer, *See* Reverse PostScript printer
- power cycling, 386
- power failure recoveries, 284
- prdaily script
  - files used by, 291
  - line usage reporting and, 293
  - overview, 281
  - runacct script and, 281, 293
- print client
  - checking configuration of, 404
  - defined, 38
  - deleting access to printers (how to), 72
  - freeing jobs in, 413
- print configuration
  - centralized, 38
  - using Solaris and lpd-based print systems, 39
- print daemons, 155
- print filters
  - adding (how to), 111
  - adding (overview), 110
  - bypassing, 408
  - changing (overview), 110
  - characteristics of, 144
  - converting file content type, 109
  - creating (how to), 146
  - creating (overview), 110, 137
  - creating definitions, 140
  - defined, 163
  - definition, 109
  - deleting (how to), 112
  - download (overview), 124
  - fast (overview), 124
  - handling special modes, 138
  - managing (overview), 109
  - options keywords, 144
  - PostScript, 110
  - removing (overview), 110
  - required for printer fault recovery, 85
  - required for printer fault recovery (overview), 85
  - requirements for, 139
  - restoring (overview), 110
  - slow (overview), 124
  - templates to define options, 143
  - to convert from troff to PostScript, 147
  - TranScript, 110
  - types of, 138
  - used to convert files, 138
  - used to put request on hold, 139
  - viewing information about (how to), 112
- print forms, *See* forms
- print jobs, *See* print requests
- print queue
  - log of, 156
- print requests
  - accepting (how to), 94
  - canceling (how to), 95
    - by disabling printer, 92
    - for specific user, 96
  - canceling (overview), 95
  - changing priority of (how to), 99
  - changing priority of (overview), 89, 99
  - checking status of (how to), 90
  - cleaning out from log file, 164
  - IDs, 90

- print requests (*continued*)
  - canceling print requests by, 95
  - components of, 95
  - in banner page, 79
  - in status of print request, 91
- IDs (how to)
  - canceling print requests by, 96
  - changing priority of print requests using, 99
  - moving print requests by, 98
- log, 157
- moving to another printer (how to), 97
- moving to another printer (overview), 97
- moving to head of queue (overview), 99
- putting on hold (overview), 99
- rejecting (how to), 73, 94
- scheduling, 162
- setting printer to accept or reject (overview), 93
- print scheduler
  - if not running, 403
  - restarting (how to), 76
  - starting (how to), 77
  - stopping (how to), 76
  - updating LP system files, 161
- print server
  - checking access to, 406
  - checking configuration of, 404
  - checking connections of, 404
  - defined, 38, 54
  - deleting printer from (how to), 73
  - hard disk requirements for, 41
  - setting up (overview), 54
  - spooling space required for, 40
  - system resource requirements for, 39
- print service, *See* LP print service
- print wheels, 103
  - alerts for mounting, 106
  - defining (how to), 104
  - mounting (how to), 105
  - naming, 103
  - tracking, 104
  - unmounting (how to), 105
- printer
  - adding a new (how to)
    - with Solaris Print Manager, 55
- printer class, 81
  - checking status for, 90
- printer class (*continued*)
  - defining with lpadmin command (how to), 82
  - not valid for enabling/disabling printer, 92
  - setting, 42
- printer description
  - adding with lpadmin command (how to), 77
  - setting with Solaris Print Manager, 41
- printer destination
  - setting with lpadmin command (overview), 78
  - setting with Solaris Print Manager, 42
- PRINTER environment variable, 79
- printer information
  - conversion from Solaris 2.5.1, 69
- printer interface program, 163
- printer name, 41
- printer port
  - adjusting characteristics of (how to), 128
  - characteristics, adjusting (overview), 127
  - configuring, 43
  - enabling multiple ports, 43
  - parallel, 43
  - serial, 43
  - setting with Solaris Print Manager, 41
- printer-resident fonts (overview), 124
- printer(s)
  - adding an unsupported (overview), 129
  - customizing interface program (how to), 136
  - customizing interface program (overview), 133
  - faults
    - detecting, 139
    - messages, 135
  - interface program
    - exit codes, 134
    - standard, 133, 136
  - local, 164
  - non-PostScript, 137
  - remote, 166
  - settings
    - stty, 127
    - stty settings, 134
- printer status, 75
- printer type
  - defined in terminfo database, 154

- printer type (*continued*)
  - not in `terminfo` database, 45
  - setting with Solaris Print Manager, 42
  - troubleshooting setting of, 407
- printers
  - accepting print requests (overview), 93
  - access to forms (example of), 116
  - adding a network printer (how to)
    - with `lp` commands, 64
    - with Vendor Supplied Tools, 63
  - adding local or attached (how to)
    - adding by using Solaris Print Manager, 55
  - allow list (overview), 87
  - allowing user access (how to), 88
  - baud settings, 398
  - controlling availability of (overview), 93
  - default, 78
  - deleting client access to (overview), 72
  - deleting from print server (how to), 72
  - deny list (overview), 87
  - denying user access (how to), 88
  - description, adding (how to), 77
  - disabling (how to), 73, 92
  - disabling (overview), 91
  - enabling (how to), 92
  - enabling (overview), 91
  - fault alerts (how to), 84
  - fault notification, 83
  - fault recovery (overview), 85
  - faults
    - restarting printing, 412
  - file content types
    - effect on filtering, 408
    - incorrect output, 408
  - font cartridges, 102
  - local, 38, 50
  - network, 50, 59
  - non-PostScript, 102
  - parity bit, 398
  - PostScript
    - character sets for, 103
    - treatment of text, 102
  - print wheels, 102
  - rejecting print requests (overview), 93
  - setting up, 49
    - with Solaris Print Manager, 52
- printers (*continued*)
  - settings
    - baud, 398
    - parity, 398
    - return, 399
    - stty, 410
    - tab, 398
  - status
    - conflicting messages, 400
  - status (how to)
    - checking, 75, 90
  - stty settings, 410
  - troubleshooting
    - backup in printer queue, 400
    - conflicting status messages, 400, 413
    - hung LP commands, 411
    - hung printers, 399, 411
    - incorrect output, 397, 407
    - incorrect printer type, 407
    - not accepting requests, 403
    - only banner page prints, 408
    - print service commands, 399
  - printers, alerts, *See* alerts
  - printers, class, *See* printer class
  - printers, type, *See* printer type
  - `.printers` file
    - setting up (how to), 58
  - printing
    - banner pages, 79
    - local (diagram), 164
    - remote, 166
    - special modes, 138
    - spooling directory, 159
    - status messages, 413
    - stopping (overview), 91
    - user fee calculation for, 274
  - `priocntl` command
    - overview, 313
    - c option (scheduling class designation), 315
    - e option (execute), 315
    - i option (ID type), 315
    - l option (scheduling class display), 314
    - m option (max/min priority), 315
    - p option (priority designation), 315
    - s option (priority upper limit/change priority), 315

- priority (process)
  - changing, 315, 318
    - timesharing processes, 315, 317
  - designating, 315
  - displaying information on, 306, 314
  - global
    - defined, 313
    - displaying, 314
  - maximums, 315
  - overview, 313, 318
  - scheduling classes and, 315
  - user-mode priority, 313
- priority of print requests, 89, 99
- private interfaces, 153
- /proc directory, 308
- proc structure, 301, 306
- proc tools
  - examining a core file, 372
- process accounting, 269, 271, 285
- process file system (PROCFS), 308
- processes
  - accounting utilities for, 269, 271, 285
  - address space map, 309
  - application threads and, 301
  - controlling, 311
  - current working directory for, 309, 311
  - defined, 301
  - displaying information (how to), 310
  - displaying information on, 306
    - acctcom command, 289
    - daily usage report, 285
    - dead processes, 289
    - listing processes, 307
    - listing processes being executed, 308
    - LWPs, 309
    - priocntl command, 314
    - ps command, 306, 308, 314
  - displaying information with proc tool
    - commands, 309
  - displaying information with proc tools, 308
  - fstat and fcntl information for open
    - files, 309, 311
  - killing, 309, 312
  - libraries linked into, 309
  - nice number of, 306, 317
  - priority, 318
    - changing, 315, 318

- processes, priority (*continued*)
  - changing timesharing process
    - priority, 315, 317
  - designating, 315
  - displaying information on, 306, 314
  - global priorities, 313
  - maximums, 315
  - overview, 313, 318
  - scheduling classes and, 313, 315
  - user-mode priority, 313
- proc tool commands, 308
- restarting, 309
- runaway, 319
- scheduling classes, 313
  - changing, 316
  - changing priority of, 315, 318
  - defaults, 315
  - designating, 315
  - displaying information on, 306, 314
  - priority levels and, 313, 315
- signal actions, 309
- stack trace, 309
- stopping temporarily, 309
- structures for, 301, 306
- terminology, 301
- tool commands, 309
- tracing flags, 309
- trees, 309, 311
- troubleshooting, 318
- PROCFS (process file system), 308
- programs
  - disk-dependency of, 331
  - forcing to quit running, 386
  - interrupting, 386
- prtconf command, 211, 215
- ps command, 306, 308
  - fields in reports from, 306, 317
  - overview, 306
    - c option (scheduling class), 306
    - ecl option (global priority), 314
    - ef option (full information), 307
- PS printer type, 45
- psig command, 309
- PSR printer type, 45
- pstack command, 309
- ptacctn.MMDD file, 283
- ptime command, 309
- ptree command, 309, 311

pwait command, 309  
pwdx command, 309, 311

## Q

queue  
  displaying at queue, 262  
quitting  
  forcing programs to quit, 386  
quot command, 226  
quota command, 236, 241, 245  
quotacheck command, 236, 240  
quotaon command, 236, 240  
quotas, 243  
  administering, 243  
  changing, 243  
  changing for individual users, 245  
  checking, 241  
  checking for exceeded user quotas, 241  
  checking on file systems, 242  
  configuring file systems for, 237  
  consistency checking, 240  
  disabling for individual users, 246  
  displaying information on, 236, 241  
  initializing, 237, 240  
  overview, 233, 236  
  prototype for multiple users, 239  
  removing, 243  
  requirements, 236  
  setting up, 235  
  soft limit time  
    changing, 244  
    exceeding, 244  
  turning off, 237, 247  
  turning on, 235, 237, 240  
  turning on, example of, 241  
  user  
    changing for individual users, 245  
    checking for exceeded, 241  
    disabling for individual users, 246  
    setting up, 238  
    soft limit time, 244  
  verifying, 236, 241, 245  
quotas file, 235, 238

## R

real-time processes  
  changing class of, 316  
  priority of  
    specifying relative, 315  
rebooting  
  fails after crash, 381  
  monitoring, 268, 271, 282, 284  
reboots file, 282, 291  
recognizing network access problems, 394  
recover root password  
  IA, 384  
  SPARC, 383  
recovering from a full crash dump  
  directory, 380  
reject command, 93  
rejecting print requests (how to), 73, 94  
rejecting print requests (overview), 93  
remote printing  
  diagram of process, 166  
  user fee calculation for, 274  
removing  
  crontab files, 257  
repetitive system tasks, 258  
repquota command, 241  
requests log, 164  
residentfonts file, 124  
restarting  
  print scheduler (how to), 76  
  processes, 309  
  runacct script, 277, 282  
restore  
  using matching commands, 388  
return settings, printers, 399  
Reverse PostScript printer  
  file content type for, 47  
  method of printing, 45  
  printer type for, 45, 47  
rm command, 228  
root crontab file, 270  
root password, forget  
  IA, 384  
  SPARC, 383  
rprt.MMDD file, 271, 293  
rpt.MMDD file, 282, 292  
RS-232-C, *See* serial port  
runacct script, 281  
  crontab entry that runs, 281

runacct script (*continued*)  
 diagnostics file, 281  
 error messages, 277  
 error protection, 281  
 failure of, 277  
 files used/produced by, 291, 293  
 fixing corrupted files, 275, 282  
 last time executed, 291  
 monacct script and, 281  
 overview, 271  
 prdaily script and, 281, 293  
 progress file, 281  
 restarting, 277, 282  
 scheduling running of, 272  
 states of, 282  
 user fee calculation and, 274, 286  
 runaway processes, 319

## S

sa1 command, 346  
 sa2 command, 346  
 SAC, *See* Service Access Controller  
 sacadm command  
 adding a ttymon port monitor with, 193  
 described, 187  
 disabling a ttymon port monitor with, 195  
 enabling a ttymon port monitor with, 195  
 killing a ttymon port monitor with, 194  
 listing a ttymon port monitor with, 193  
 removing a ttymon port monitor with, 195  
 starting a ttymon port monitor with, 194  
 sadc command, 346, 349  
 sadd file, 347  
 SAF, *See* Service Access Facility  
 sar command, 330, 347  
 description of all options, 348  
 options listed, 348  
 overview, 330, 347  
 -A option (overall performance), 346, 348  
 -a option (file access), 330  
 -b option (buffers), 331  
 -c option (system calls), 332  
 -d option (block devices), 333  
 -e option (ending time), 348  
 -f option (file to extract data from), 348  
 -g option (page-out/memory freeing), 335

sar command (*continued*)  
 -i option (interval), 348  
 -k option (kernel memory), 336  
 -m option (interprocess communication), 337  
 -p option (page-in/page faults), 338  
 -q option (queue), 340  
 -r option (unused memory), 341  
 -s option (starting time), 348  
 -u option (CPU usage), 341  
 -v option (system tables), 342  
 -w option (swapping/switching volume), 343  
 -y option (terminal devices), 345  
 saving crash dump information, 375  
 scheduler, *See* print scheduler  
 scheduling  
 one-time system tasks, 250, 260, 264  
 repetitive system tasks, 249, 258  
 scheduling classes, 313  
 changing, 316  
 changing priority of, 315, 318  
 defaults, 315  
 designating, 315  
 displaying information on, 306, 314  
 priority levels and, 313, 315  
 search path  
 files for setting, 392  
 security  
 at command, 261  
 crontab command, 258  
 selectable character sets, 102, 108  
 serial port  
 adapter board, 173  
 defined, 173  
 Serial Ports Tool  
 terminals and modems, 174  
 serial printer, adding to system, 43  
 Service Access Controller, 187  
 Service Access Facility  
 description, 174  
 files associated with (table), 200  
 overview of, 175, 186  
 programs associated with (table), 186  
 services controlled by  
 states of (table), 203  
 uses for, 175, 186  
 when to use, 174



- setting
  - a core file name pattern with `coreadm`, 370
  - alerts to mount font cartridges (how to), 106
  - alerts to mount print wheels (how to), 106
  - alias for selectable character set (example of), 109
  - alias for selectable character set (how to), 108
- setting up
  - character sets (example of), 109
- shared memory
  - process virtual memory, 302
- `shutacct` script, 271
- shutdown command, 272
- shutdowns
  - monitoring, 271, 284
- simple file content type, 46
- size
  - directories
    - `du` command, 224
    - directory, 224, 226
    - file, 222, 226
- slow print filters, 124
- soft limit time
  - changing, 244
  - exceeding, 244
- software packages
  - troubleshooting installation of, 437
- Solaris and lpd-based print systems, 39
- Solaris Print Manager
  - ability to define printer with (overview), 41
  - adding a new printer (how to), 55
  - adding printer access, 57
  - overview, 28
  - prerequisites for using, 52
  - starting (how to), 52
- `Spacctn.MMDD` file, 282, 291
- space (disk)
  - amount of free, 220, 329
  - optimizing, 219
- spooling directory
  - print service, 159
- spooling space required for print server, 40
- standard printer interface program, 136
- starting
  - print scheduler (how to), 77
  - Solaris Print Manager (how to), 52
- startup command
  - `acct`, 270
- `statefile` file, 277, 282, 291
- states, (`runacct` script), 282
- status
  - checking of print requests, 90
  - getting printer (how to), 75
- stopping
  - a printer (overview), 91
  - print scheduler (how to), 76
  - processes temporarily, 309
  - system accounting, 278
- stty settings
  - customizing, 134
  - defaults, 127, 410
  - suggestions for, 410
  - troubleshooting, 398
- superuser (root) password, forget
  - IA, 384
  - SPARC, 383
- `sys crontab`, 347
- `syslog.conf` file, 360
- `syslogd` daemon, 358
- system activities
  - automatic collection of data on, 346, 349
  - list of activities tracked, 302
  - manual collection of data on, 330, 347
- system crash information
  - managing with `dumpadm`, 375
- system message logging (customizing), 360
- system messages
  - customizing logging (how to), 361
  - specifying storage location for, 358
- system resources
  - accounting
    - overview, 267
  - allocating for print server, 39
  - monitoring, 261
    - accounting, 277
    - accounting system for, 293
    - automatic, 261
    - crashes, 360, 386
    - quotas, 243
    - overview, 300
- system tasks
  - scheduling
    - one-time tasks, 250, 260, 264
    - repetitive tasks, 258

system tasks (*continued*)  
scheduling automatically, 249

## T

tab settings, 398  
tacct file, 276, 282, 292  
tacct.MMDD file, 276, 282, 292  
tacctn file, 293  
tacctprev file, 292  
technical support  
  crash dump analysis, 375  
  sending crash information, 356  
templates (print filters), 143  
temporary directories, 227, 229  
terminal characteristics, 134  
terminals  
  alphanumeric, 172  
  defined, 172  
  distinctions between types of, 172  
  line usage monitoring, 268, 284, 293  
  overview of Serial Ports Tool, 177  
  process controlling, 306  
  Serial Ports Tool item descriptions, 178  
  setting up (how to), 180  
  tools for managing, 174  
  troubleshooting bad lines, 285  
terminfo database  
  character set names, 102  
terminfo entry  
  adding (overview), 129  
  adding for unsupported printer (how to), 132  
  list of required items, 129  
  selectable character sets in, 102  
time  
  changing, 216  
  CPU usage, 286, 306, 318  
  displaying, 211, 215  
  processes accumulating large amounts of  
    CPU time, 318  
  soft limit, 244  
timesharing processes  
  changing scheduling parameters, 315  
  priority of  
    changing, 315, 317  
    overview, 313

timesharing processes, priority of (*continued*)  
  range of, 313  
  /tmp/disktacct.MMDD file, 282  
tmpwtmp file, 282, 291, 293  
tools  
  for displaying process information, 308  
  process, 309  
  system performance monitoring, 303  
total command summary, 288, 292  
tracing flags, 309  
TranScript filters, 110  
troff filter to convert to PostScript, 147  
troubleshooting, 399  
  printing problems, 395, 413  
  processes, 318  
  software package installation/removal, 437  
  tty lines, 285  
tty lines  
  troubleshooting bad lines, 285  
tty lines  
  usage monitoring, 284  
tty lines  
  usage monitoring, 268, 285, 293  
ttyadm command, 192  
ttymon port monitor  
  (*figure*), 189  
  adding, 193  
  bidirectional modem service and, 190  
  disabling, 195  
  enabling, 195  
  killing, 194  
  listing, 193  
  overview of function of, 188  
  removing, 195  
  starting, 194  
ttymon service  
  adding, 196  
  disabling, 200  
  enabling, 199  
  listing, 197  
tuning  
  daily command summary and, 286  
turnacct switch script, 271  
turnacct switch script, 282  
turning off quotas, 237, 247  
turning on quotas, 237, 240  
turning on quotas, example of, 241

## U

- UFS file systems
  - displaying information about, 226
  - user ownership of, 226
- uname command, 213
- UNIX systems (crash information), 374
- unmounting
  - font cartridges (how to), 105
  - forms (how to), 117
  - print wheels (how to), 105
- unsupported printers, 47, 129
- user
  - access to printer
    - ability to limit with Solaris Print Manager, 42
  - access to printer (overview)
    - limiting with lpadmin command, 87
    - limiting with Solaris Print Manager, 87
  - canceling print requests for (how to), 97
- user fees, 271, 275, 286
- user logins
  - last login monitoring, 282, 288, 293
  - number of logins, 286
  - time monitoring, 268, 282, 286
- user-mode priority, 313
- user ownership of disk space, 226
- user processes
  - changing priority, 317
  - CPU usage by, 286
  - priority of, 313
- user quotas
  - changing for individual users, 245
  - checking for exceeded, 241
  - disabling for individual users, 246
  - setting up, 238
  - soft limit time, 244
- user structure, 301
- /usr/adm/messages file, 356
- /usr/bin/mdb utility, 379
- /usr/lib/lp directory, 110
- /usr/lib/lp/model directory, 163
- /usr/lib/lp/postscript directory, 110, 163
- /usr/proc/bin directory, 308
- utmp2wtmp command, 282

## V

- /var/adm/acct directory, 291
- /var/adm/acct/fiscal directory, 291
- /var/adm/acct/nite/active file, 277, 281, 291
- /var/adm/acct/nite/active.MMDD file, 281, 291
- /var/adm/acct/nite/cms file, 282, 291
- /var/adm/acct/nite/ctacct.MMDD file, 282, 291
- /var/adm/acct/nite/ctmp file, 291
- /var/adm/acct/nite/daycms file, 282, 291, 293
- /var/adm/acct/nite/daytacct file, 282, 286, 291, 293
- /var/adm/acct/nite directory, 291
- /var/adm/acct/nite/disktacct file, 270, 282, 291
- /var/adm/acct/nite/disktacct.MMDD file, 282
- /var/adm/acct/nite/fd2log file, 277, 281, 291
- /var/adm/acct/nite/lastdate file, 282, 291
- /var/adm/acct/nite/lineuse file, 282, 291, 293
- /var/adm/acct/nite/lock file, 277, 282, 291
- /var/adm/acct/nite/lock1 file, 282
- /var/adm/acct/nite/log file, 291
- /var/adm/acct/nite/log.MMDD file, 291
- /var/adm/acct/nite/owtmpx file, 292
- /var/adm/acct/nite/reboots file, 282, 291
- /var/adm/acct/nite/statefile file, 277, 282, 291
- /var/adm/acct/nite/tmpwtmp file, 282, 291, 293
- /var/adm/acct/nite/wtmp.MMDD file, 282, 292
- /var/adm/acct/nite/wtmperror file, 291
- /var/adm/acct/nite/wtmperror.MMDD file, 291
- /var/adm/acct/sum/cms file, 282, 292
- /var/adm/acct/sum/cmsprev file, 292
- /var/adm/acct/sum/daycms file, 292
- /var/adm/acct/sum/daycmsfile, 282

- /var/adm/acct/sum directory, 271, 291
- /var/adm/acct/sum/loginlog file, 282, 292
- /var/adm/acct/sum/rprtMMDD file, 271
- /var/adm/acct/sum/rprt.MMDD file, 293
- /var/adm/acct/sum/rpt.MMDD file, 282
- /var/adm/acct/sum/tacct file, 276, 292
- /var/adm/acct/sum/tacct file, 282
- /var/adm/acct/sum/tacctMMDD file, 276
- /var/adm/acct/sum/tacct.MMDD file, 282, 292
- /var/adm/acct/sum/tacctprev file, 276, 292
- /var/adm directory
  - described, 291
- /var/adm directory
  - controlling size of, 228
- /var/adm directory
  - raw accounting data in, 270
- /var/adm/dtmp file, 291
- /var/adm/fee file, 271, 275, 282, 291
- /var/adm/messages file, 356, 360
- /var/adm/messages.n file, 358
- /var/adm/sa/sadd file, 347
- /var/adm/Spacctn.MMDD file, 282, 291
- /var/lp/logs directory, 164
- /var/lp/logs/lpsched file, 163, 397
- /var file system, 40
- /var/spool/cron/atjobs directory, 250, 253
- /var/spool/cron/crontabs directory, 252
- /var/spool/cron/crontabs/lp file, 164
- /var/spool/cron/crontabs/root file, 251, 270
- /var/spool/cron/crontabs/sys crontab, 347
- /var/spool/lp directory, 159
- /var/spool/lp directory, 40
- /var/spool/lp/requests directory, 156
- /var/spool/lp/tmp directory, 156
- verifying
  - at jobs, 262
  - crontab file existence, 255
  - quotas, 236, 241, 245
- vfstab file, quotas and, 237, 240

- vmstat command
  - displaying virtual memory, 323
  - fields in reports from, 322
  - overview, 322
  - S option (swapping statistics), 325
  - c option (cache flushing), 325
  - i option (interrupts), 326
  - s option (system events), 324

## W

- Watchdog reset ! message, 358
- weekly tasks (scheduling with crontab), 250
- wtmp.MMDD file, 282, 292
- wtmperror file, 291
- wtmperror.MMDD file, 291
- wtmpfix command, 275, 282, 291
- wtmpx file
  - daily report and, 284
  - fixing corrupted, 275, 282
  - overview, 270, 275, 282
  - shutdowns and, 272